

IBM TotalStorage SAN Volume Controller
Storage Software for Cisco MDS 9000



Configuration Guide

Version 1.1.8

IBM TotalStorage SAN Volume Controller
Storage Software for Cisco MDS 9000



Configuration Guide

Version 1.1.8

Note

Before using this information and the product it supports, read the information in "Notices" on page 269.

Second Edition (July 2004)

© Copyright International Business Machines Corporation 2003, 2004. All rights reserved.

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

Contents

About this guide	xi
Who should use this guide	xi
Related publications	xi
How to order IBM publications	xiii
Related Web sites	xiv
Emphasis	xiv
Numbering conventions	xiv
How to send your comments	xv
Summary of changes.	xv
Summary of Changes for the Configuration Guide for the IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000	xvi

Part 1. Overview 1

Chapter 1. SAN Volume Controller Storage Software for Cisco MDS 9000	3
Common terms for the SAN Volume Controller Storage Software for Cisco MDS 9000	5
Virtualization	6
Asymmetric virtualization	8
Symmetric virtualization	9
Chapter 2. Object overview	11
Nodes and clusters	12
Clusters	12
Nodes	13
I/O groups and Internal Battery Backup	14
Input/Output (I/O) groups	14
Internal battery backup overview	16
Storage subsystems and managed disks	16
Storage subsystems	16
Managed disks	18
Managed disk groups and virtual disks (VDisks).	20
Managed disk groups	20
Virtual disks	22
Hosts and virtual (VDisk) mappings	24
Host objects	24
Virtual disk-to-host mapping	25
Chapter 3. Copy Services	29
FlashCopy	29
FlashCopy mappings.	29
FlashCopy consistency groups	33
FlashCopy applications	35
FlashCopy indirection layer	35
Background copy	36
Host considerations for FlashCopy integrity	37
Remote Copy	38
Synchronous Remote Copy	39
Remote Copy partnerships	39
Remote Copy relationships	39
Remote Copy consistency groups	41
Chapter 4. Configuration rules and requirements	43

Configuration rules	43
Storage subsystems	43
Host bus adapters.	45
Nodes	46
Fibre-channel switches	46
Configuration requirements	47

Part 2. Preparing to configure the SAN Volume Controller Storage Software for Cisco MDS 9000 51

Chapter 5. Centralized management functions	53
Secure Shell (SSH)	53
Configuring the Secure Shell (SSH) client system	54
Configuring the PuTTY session for the command-line interface	54
Upgrading software on the console workstation	55
Remote Service	55
Troubleshooting console workstation problems	56
Recovering from Console for SAN Volume Controller for Cisco MDS 9000 signing off	56
Installing antivirus software	56

Part 3. SAN Volume Controller Console 57

Chapter 6. Console for SAN Volume Controller for Cisco MDS 9000	59
Accessing the Console for SAN Volume Controller for Cisco MDS 9000	59
Console for SAN Volume Controller for Cisco MDS 9000 layout	60
Console for SAN Volume Controller for Cisco MDS 9000 banner area	60
Console for SAN Volume Controller for Cisco MDS 9000 task bar	60
Console for SAN Volume Controller for Cisco MDS 9000 portfolio	60
Console for SAN Volume Controller for Cisco MDS 9000 work area	61

Chapter 7. Overview of creating a cluster using the Console for SAN Volume Controller for Cisco MDS 9000 63

Browser requirements for the Console for SAN Volume Controller for Cisco MDS 9000.	63
Preparing the Cisco MDS switch	63
Setting up the Cisco MDS switch	64
Creating SAN Volume Controller Storage Software for Cisco MDS 9000 interfaces	67
Configuring a cluster using the Console for SAN Volume Controller for Cisco MDS 9000.	67
Launching the Console for SAN Volume Controller for Cisco MDS 9000	70
Displaying cluster properties using the Console for SAN Volume Controller for Cisco MDS 9000	70
Changing the password for the Console for SAN Volume Controller for Cisco MDS 9000.	71
Managing SAN Volume Controller Users	71

Chapter 8. Scenario: typical usage for the Console for SAN Volume Controller for Cisco MDS 9000 73

Adding nodes to a cluster	74
Displaying node properties using the Console for SAN Volume Controller for Cisco MDS 9000	79
Creating managed disk groups	80
Creating virtual disks.	81

Creating hosts	82
Creating virtual disk-to-host mappings	82
Showing VDisks mapped to a host	83
Creating consistency groups	83
Creating FlashCopy mappings	84
Chapter 9. Advanced function FlashCopy overview	85
Starting FlashCopy mappings	85
Stopping FlashCopy mappings	85
Deleting FlashCopy mappings	85
Preparing FlashCopy consistency groups	86
Starting FlashCopy consistency groups	86
Stopping FlashCopy consistency groups	86
Deleting FlashCopy consistency groups	87
Chapter 10. Advanced functions overview for the SAN Volume Controller	
Console	89
Determining the WWPNs for a node using the Console for SAN Volume Controller for Cisco MDS 9000	89
Determining a storage subsystem name from its SAN Volume Controller Storage Software for Cisco MDS 9000 name using the Console for SAN Volume Controller for Cisco MDS 9000	89
Determining the relationship between VDisks and MDisks using the Console for SAN Volume Controller for Cisco MDS 9000	89
Determining the relationship between managed disks and RAID arrays or LUNs using the Console for SAN Volume Controller for Cisco MDS 9000	90
Virtual disk-to-host mappings	90
Increasing the size of your cluster using the Console for SAN Volume Controller for Cisco MDS 9000	91
Adding a node to increase the size of your cluster	91
Migrating a VDisk to a new I/O group	92
Replacing a faulty node with a spare node using the Console for SAN Volume Controller for Cisco MDS 9000	93
Recovering from offline VDisks after a node or an I/O group failed	94
Recovering a node and including it back into the cluster.	95
Moving offline VDisks to the recovery I/O group	96
Moving offline VDisks to their original I/O group	97
Replacing an HBA in a host using the Console for SAN Volume Controller for Cisco MDS 9000	97
Resetting the Node	98
Deleting hosts	98
Adding a new storage controller to a running configuration using the Console for SAN Volume Controller for Cisco MDS 9000	99
Removing a storage subsystem using the Console for SAN Volume Controller for Cisco MDS 9000	100
Shrinking virtual disks	101
Migrating virtual disks	102
Creating image-mode virtual disks	103
Advanced function Remote Copy overview	104
Advanced function cluster overview	105
Removing a node from a cluster	105
Viewing feature settings and log	107
Configuring error notification settings	107
Resetting the SSH fingerprint	107
Modifying Internet Protocol (IP) addresses	108
Listing and saving log and dump files	109

Viewing the feature log	109
Analyzing the error log.	110
Shutting down a cluster or node	110

Part 4. Installing or upgrading the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 for Windows 113

Chapter 11. Installation overview for the Console for SAN Volume Controller for Cisco MDS 9000	115
Chapter 12. Console for SAN Volume Controller for Cisco MDS 9000 hardware installation requirements	117
Chapter 13. Console for SAN Volume Controller for Cisco MDS 9000 workstation space requirements	119
Chapter 14. Console for SAN Volume Controller for Cisco MDS 9000 software installation requirements	121
Chapter 15. Installing or upgrading the Console for SAN Volume Controller for Cisco MDS 9000 in graphical mode	123
Chapter 16. Installing or upgrading the Console for SAN Volume Controller for Cisco MDS 9000 in unattended (silent) mode.	129
Chapter 17. Verifying the Windows services associated with the Console for SAN Volume Controller for Cisco MDS 9000	135
Chapter 18. Post installation tasks	137
Chapter 19. Removing the Console for SAN Volume Controller for Cisco MDS 9000	141

Part 5. Command-Line Interface 145

Chapter 20. Getting started with the Command-Line Interface.	147
Issuing CLI commands from a PuTTY SSH Client system	147
Running the PuTTY and plink utilities	148
Configuring the cluster using the CLI	150
Reviewing and setting the cluster features using the CLI	150
Displaying cluster properties using the CLI	151
Chapter 21. Scenario: typical usage for the command-line interface	153
Adding nodes to a cluster using the IBM TotalStorage CLI	154
Displaying node properties using the CLI	156
Removing managed disks that represent de-configured LUs.	157
Discovering MDisks using the CLI	158
Creating managed disk (MDisk) groups using the CLI	159
Adding MDisks to MDisk groups using the CLI.	161
Create virtual disks (VDisks)	162
Creating host objects using the CLI	165
Create VDisk-to-host mappings using the CLI	166
Create FlashCopy mappings using the CLI	167
Creating a FlashCopy consistency group and adding mappings using the CLI	168
Preparing and triggering a FlashCopy Consistency Group using the CLI	169

Chapter 22. Advanced functions with the CLI	171
Determining a nodes WWPNs using the CLI	171
Determining a storage subsystem name from its SAN Volume Controller	
Storage Software for Cisco MDS 9000 name using the CLI	171
Determining the VDisk name from the vpath number on the host	171
Determining the host that a VDisk is mapped to	172
Determining the relationship between VDIsks and MDIsks using the CLI	173
Determining the relationship between MDIsks and RAID arrays or LUNs using the CLI	174
Increasing the size of your cluster using the CLI	174
Adding a node to increase the size of your cluster using the CLI	175
Migrating a VDisk to a new I/O group	176
Replacing a faulty node in the cluster using the CLI	177
Recovering from offline VDIsks after a node or an I/O group failed using the CLI	178
Recovering a node and including it back into the cluster	179
Moving offline VDIsks to the recovery I/O group	180
Moving offline VDIsks to their original I/O group using the CLI	181
Replacing an HBA in a host using the CLI	182
Adding a new storage controller to a running configuration using the CLI	183
Removing a storage subsystem using the CLI	184
Expanding VDIsks	185
Expanding a Virtual disk that is mapped to an AIX host	186
Expanding a Virtual disk that is mapped to a Windows 2000 host	187
Shrinking a VDisk using the CLI	188
Migrating extents using the CLI	189
Migrating VDIsks between MDisk groups using the CLI	191
Migrating a VDisk between I/O groups using the CLI	192
Creating an image mode VDisk from an unmanaged MDisk using the CLI	193
Advanced function FlashCopy and Remote Copy overview for CLI	194
Advanced function cluster overview using the CLI	194
Deleting a node from a cluster using the IBM TotalStorage CLI	194
Performing the cluster maintenance procedure using the CLI	196
Setting up error notifications using the CLI	196
Listing log or dump files using the CLI	196
Viewing the feature log using the CLI	197
Analyzing the error log using the CLI	198
Shutting down a cluster or single node using the CLI	198

Part 6. Software upgrade strategy 201

Chapter 23. Upgrading the SAN Volume Controller Storage Software for Cisco MDS 9000 firmware using the Console for SAN Volume Controller for Cisco MDS 9000	203
Chapter 24. Automatic upgrade	205
Chapter 25. Automatic recovery from upgrade problems	207
Chapter 26. Installing the upgrade using the CLI	209
Chapter 27. Installing the software	211
Chapter 28. Manual recovery from software upgrade problems	213

**Part 7. Configuring other SAN devices for use with the SAN Volume Controller
Storage Software for Cisco MDS 9000 215**

Chapter 29. Configuring and servicing storage subsystems 217
Configuring a balanced storage subsystem 217

Chapter 30. Configuring the EMC CLARiiON controller 221
Configuring the EMC CLARiiON controller with Access Logix installed 221
 Registering the SAN Volume Controller Storage Software for Cisco MDS
 9000 Ports with your EMC CLARiiON 221
 Configuring your storage groups 222
Configuring the EMC CLARiiON controller (Access Logix not installed) 223
Supported models of the EMC CLARiiON 223
Supported firmware levels for the EMC CLARiiON 223
Concurrent maintenance on the EMC CLARiiON 224
Sharing the EMC CLARiiON between a host and the SAN Volume Controller
 Storage Software for Cisco MDS 9000 224
Switch zoning limitations for the EMC CLARiiON 225
Quorum disks on the EMC CLARiiON 225
Advanced functions for the EMC CLARiiON 225
Logical unit creation and deletion on the EMC CLARiiON 226
Configuring settings for the EMC CLARiiON 226
 Global settings for the EMC CLARiiON 226
 Controller settings for the EMC CLARiiON 227
 Port settings for the EMC CLARiiON 227
 LU settings for the EMC CLARiiON 228

Chapter 31. Configuring the Enterprise Storage Server 231
Configuring the Enterprise Storage Server (ESS) 231
Supported models of the ESS 232
Supported firmware levels for the ESS 232
Concurrent maintenance on the ESS 232
Sharing the ESS between a host and the SAN Volume Controller Storage
 Software for Cisco MDS 9000 232
Switch zoning limitations for the ESS 233
Quorum disks on the ESS 233
Advanced functions for the ESS 233
Logical unit creation and deletion on the ESS 233

Chapter 32. Configuring the FASTt disk controller system 235
Configuring FASTt disk controllers for the storage server 235
Support actions for the FASTt controller 236
Supported models of the IBM FASTt controller 237
Supported firmware levels for the FASTt 237
Concurrent maintenance on the IBM FASTt 237
Sharing the IBM FASTt controller between a host and the SAN Volume
 Controller Storage Software for Cisco MDS 9000 237
Quorum disks on the IBM FASTt 238
Advanced functions for the IBM FASTt 238
 Data migration on an existing FASTt installation which contains arrays 238
Logical unit creation and deletion on the IBM FASTt 239
Configuration interface for the IBM FASTt 239
Controller settings for the IBM FASTt 240
 Configuring settings for the IBM FASTt 241
 Global settings for the IBM FASTt 241

	LU settings for the IBM FAStT	242
	Miscellaneous settings for the IBM FAStT	242
	Mapping and virtualization settings for IBM FAStT	242
	Chapter 33. Configuring the HDS Lightning disk controller system.	243
	Supported models of the HDS Lightning	243
	Supported firmware levels for HDS Lightning	243
	Concurrent maintenance on the HDS Lightning 99xxV	243
	Sharing the HDS Lightning 99xxV between host and the SAN Volume Controller	243
	Quorum disks on HDS Lightning 99xxV	244
	Advanced functions for HDS Lightning	244
	Chapter 34. Configuring the HP StorageWorks subsystem	247
	HP StorageWorks definitions	247
	Configuring the HP StorageWorks controller.	249
	Supported models of the HP StorageWorks controller	251
	Supported firmware levels for the HP StorageWorks controller	251
	Concurrent maintenance on the HP StorageWorks	251
	Sharing the HP StorageWorks controller between a host and the SAN Volume Controller Storage Software for Cisco MDS 9000	252
	Switch zoning limitations for the HP StorageWorks subsystem	252
	Quorum disks on the HP StorageWorks subsystem	252
	Support for HP StorageWorks advanced functions	253
	HP StorageWorks advanced functions	253
	Logical unit creation and deletion on the HP StorageWorks	253
	Configuration interface for the HP StorageWorks	254
	Configuring settings for the HP StorageWorks	255
	Global settings for the HP StorageWorks	255
	Controller settings for the HP StorageWorks.	255
	Port settings for the HP StorageWorks.	256
	LU settings for the HP StorageWorks	257
	Connection settings for the HP StorageWorks	258
	Mapping and virtualization settings for the HP StorageWorks	258
	Appendix A. Valid combinations of FlashCopy and Remote Copy functions	259
	Appendix B. Object types	261
	Appendix C. Event codes	263
	Configuration event codes	263
	Appendix D. Accessibility	267
	Notices	269
	Trademarks.	270
	Glossary	271
	Index	277

About this guide

The IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000 Configuration guide also describes the configuration tools, both command-line and Web based, that you can use to define, expand, and maintain the storage of the IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000.

Related topics:

- “Who should use this guide”
- “Numbering conventions” on page xiv

Who should use this guide

Before using the IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000, you should have an understanding of storage area networks (SANs), the storage requirements of your enterprise, and the capabilities of your storage units.

Related topics:

- “About this guide”
- “Numbering conventions” on page xiv

Related publications

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

- The publications that make up the library for the IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000
- Other IBM publications that relate to the SAN Volume Controller Storage Software for Cisco MDS 9000
- Cisco publications that relate to the SAN Volume Controller Storage Software for Cisco MDS 9000

SAN Volume Controller Storage Software for Cisco MDS 9000 library:

Table 1 on page xii lists and describes the publications that make up the SAN Volume Controller Storage Software for Cisco MDS 9000 library. Unless otherwise noted, these publications are available in Adobe portable document format (PDF) on a compact disc (CD) that comes with the SAN Volume Controller Storage Software for Cisco MDS 9000. If you need additional copies of this CD, the order number is SK2T-8813. These publications are also available as PDF files from the following Web site:

<http://www.ibm.com/storage/support/2062-2300>

Table 1. IBM Publications in the SAN Volume Controller Storage Software for Cisco MDS 9000 library

Title	Description	Order number
<i>IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000: Command-Line Interface User's Guide</i>	This guide describes the commands that you can use from the SAN Volume Controller Storage Software for Cisco MDS 9000 command-line interface (CLI).	SC26-7555
<i>IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000: Configuration Guide</i>	This guide provides guidelines for configuring your SAN Volume Controller Storage Software for Cisco MDS 9000.	SC26-7554
<i>IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000: Getting Started Guide</i>	This guide introduces the SAN Volume Controller Storage Software for Cisco MDS 9000 and lists the features you can order. It also provides guidelines for planning the installation and configuration of the SAN Volume Controller Storage Software for Cisco MDS 9000.	GA22-1055

Other IBM publications:

Table 2 lists and describes other IBM publications that contain additional information related to the SAN Volume Controller Storage Software for Cisco MDS 9000.

Table 2. Other IBM publications

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

Cisco Systems publications:

The following list identifies publications from Cisco Systems that contain additional information that is related to the Cisco MDS 9000 Family.

- *Regulatory Compliance and Safety Information for the Cisco MDS 9000 Family*
- *Cisco MDS 9100 Series Quick Start Guide*
- *Cisco MDS 9500 Series and Cisco MDS 9216 Quick Start Guide*

- *Cisco MDS 9100 Series Hardware Installation Guide*
- *Cisco MDS 9216 Switch Hardware Installation Guide*
- *Cisco MDS 9500 Series Hardware Installation Guide*
- *Cisco SAN Volume Controller Configuration Guide*
- *Cisco MDS 9000 Family Command Reference*
- *Cisco MDS 9000 Family Configuration Guide*
- *Cisco MDS 9000 Family Fabric Manager User Guide*
- *Cisco MDS 9000 Family Troubleshooting Guide*
- *Cisco MDS 9000 Family System Messages Guide*
- *Cisco MDS 9000 Family MIB Reference Guide*

You can access the most current technical documentation at the following Web site:
<http://www.cisco.com/>

Click:

- **Products and Services**
- **Storage Networking Products**
- **Cisco MDS 9000 Series Multilayer Switches**

Ordering technical documentation from Cisco Systems:

You can find instructions for ordering documentation at the following Web site:

http://www.cisco.com/univercd/cc/td/doc/es_inpk/pdi.htm

You can order Cisco Systems documentation in the following ways:

- Registered Cisco.com users (Cisco direct customers) can order Cisco product documentation from the Networking Products MarketPlace at:
<http://www.cisco.com/en/US/partner/ordering/index.shtml>
- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco Systems Corporate Headquarters (California, USA) at 408-526-7208 or, elsewhere in North America, by calling 800-553-NETS (800-553-6387).

Related topics:

- “How to order IBM publications”
- “How to send your comments” on page xv

How to order IBM publications

The IBM publications center:

The 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:

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

Publications notification system:

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

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

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

Related topics:

- “Related publications” on page xi

Related Web sites

Table 3. Web sites

Type of information	Web site
SAN Volume Controller Storage Software for Cisco MDS 9000 support	http://www.ibm.com/storage/support/2062-2300/
Technical support for IBM storage products	http://www.ibm.com/storage/support/
Cisco Systems documentation	http://www.cisco.com/univercd/home/home.htm
Cisco Systems product site	http://www.cisco.com/
Cisco Systems international site	http://www.cisco.com/public/countries_languages.shtml

Emphasis

The following typefaces are used to show emphasis:

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

Numbering conventions

The following numbering conventions are used in this guide and in the product:

- 1 kilobyte (KB) is equal to 1024 bytes
- 1 megabyte (MB) is equal to 1 048 576 bytes
- 1 gigabyte (GB) is equal to 1 073 741 824 bytes

- 1 terabyte (TB) is equal to 1 099 511 627 776 bytes
- 1 petabyte (PB) is equal to 1 125 899 906 842 624 bytes

Related topics:

- “About this guide” on page xi
- “Who should use this guide” on page xi

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. Return it by mail or give it to an IBM representative. 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.

Related topics:

- “Related publications” on page xi

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 details the revisions added to this release.

Related topics:

- “Summary of Changes for the Configuration Guide for the IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000” on page xvi

Summary of Changes for the Configuration Guide for the IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000

New information:

The following new information was added to this guide.

- Preparing the Cisco MDS switch
- Setting up the Cisco MDS switch
- Chapter 30. Configuring the EMC CLARiiON controller
- Chapter 31. Configuring the Enterprise Storage Server
- Most sections in Chapter 32. Configuring the FASTT disk controller system
- Chapter 33. Configuring the HDS Lightning disk controller system .
- Chapter 34. Configuring the HP StorageWorks subsystem

Changed information:

The following information was changed in this guide.

- Replacing a faulty node with a spare node using the Console for SAN Volume Controller for Cisco MDS 9000
- Replacing a faulty node in the cluster using the CLI .

Deleted information:

The following information was deleted from this guide.

- Enhanced remote support configuration
- Configuring the console workstation host name

Related topics:

- “Summary of changes” on page xv

Part 1. Overview

Chapter 1. SAN Volume Controller Storage Software for Cisco MDS 9000

The SAN Volume Controller Storage Software for Cisco MDS 9000 provides virtualization software on a Cisco MDS 9000 Family switch. The IBM® TotalStorage® SAN Volume Controller Storage Software for Cisco MDS 9000 provides symmetric virtualization by creating a pool of managed disks from the attached storage subsystems, which are then mapped to a set of virtual disks for use by attached host computer systems. System administrators can view and access a common pool of storage on the SAN, which enables them to use storage resources more efficiently and provides a common base for advanced functions.

The SAN Volume Controller Storage Software for Cisco MDS 9000 performs the following functions for the SAN storage that it is controlling:

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

A *node* is a single storage engine. The nodes come in pairs on a specialized MDS 9000 line card called a Cisco MDS 9000 Caching Services Module. One or two pairs of nodes constitute a *cluster*. A minimum of two Cisco MDS 9000 Caching Services Modules are required to create a redundant configuration. All I/O operations handled by the nodes in an I/O group are cached on both nodes for redundancy. Two nodes on separate line cards can form an I/O group. Each virtual volume is defined to an I/O group. Each node has a battery that allows it to dump its cache information to the local hard disk drive when line power fails.

The SAN Volume Controller Storage Software for Cisco MDS 9000 I/O groups see the storage presented to the SAN by the backend controllers as a number of disks known as *managed disks*. The application services do not see these managed disks. Instead they see a number of logical disks, known as *virtual disks*, that are presented to the SAN by the SAN Volume Controller Storage Software for Cisco MDS 9000. Each node must only be in one I/O group and provide access to the virtual disks in the I/O group.

The caching service module appears as another device to the SAN to which the Cisco MDS 9000 switch is attached. It has a worldwide node name (WWNN), worldwide port names (WWPNs), a fibre-channel connection through the backplane of the Cisco MDS 9000 switch, and a set of node ports on the MDS 9000 switch. The node ports are virtual ports created by the supervisor line card on the MDS 9000 switch. There are three node ports. The intended use of the node ports is:

- One node port for communication with hosts

- One node port for communication with the storage devices
- One node port for communication between the nodes

These node ports can attach to multiple virtual SANs (VSANs), up to a maximum of 64 VSANs.

The SAN Volume Controller Storage Software for Cisco MDS 9000 helps to provide continuous operations and can also optimize the data path to ensure performance levels are maintained.

The SAN fabric contains three distinct zones, or distinct VSANs, that correspond to the node ports. Typically, there is a host zone, a disk zone, and a management zone. In the host zone, the host systems can identify and address the nodes. You can have more than one host zone. Generally, you will create one host zone per operating system type. In the disk zone, the nodes can identify the disk drives. Host systems cannot operate on the disk drives directly; all data transfer occurs through the nodes. In the management zone, the nodes can communicate with each other to facilitate functions such as remote copy. The port that is used for communication between the nodes can be placed in a maximum of four VSANs. Several host systems can be connected to a SAN fabric. A cluster of SAN Volume Controller Storage Software for Cisco MDS 9000s is connected to the same fabric and presents virtual disks to the host systems. You configure these virtual disks using the disks located on the RAID controllers.

Note: You can have more than one host zone. Generally, you create one host zone per operating system type, because some operating systems will not tolerate other operating systems in the same zone.

When hardware service or maintenance is required, you can remove the nodes on the Cisco MDS 9000 Caching Services Module from their clusters. After you remove the nodes, you can replace the Cisco MDS 9000 Caching Services Module field replaceable unit (FRU). All disk drive communication and communication between nodes is performed through the SAN. All SAN Volume Controller Storage Software for Cisco MDS 9000 configuration and service commands are sent to the cluster through an Ethernet network.

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

Configuration information is stored on every node that is in the cluster to allow replacement of FRUs. When a new Cisco MDS 9000 Caching Services Module is installed and when the node is added back into the cluster, configuration information that is required by that node is ready from other nodes in the cluster.

SAN Volume Controller Storage Software for Cisco MDS 9000 operating environment:

- Contains a minimum of four SAN Volume Controller Storage Software for Cisco MDS 9000 nodes

Features of the Cisco MDS 9000 Caching Services Module:

- Two Sibyte BCM 1250 dual-processor SOC (System On a Chip)
 - Each Sibyte consists of two Broadcom SB-1 MIPS64 CPUs running at 750 Mhz
- 4 GB of memory per Sibyte

- 20 GB IDE hard disk drive per Sibyte
- 1 Flamingo ASIC per Sibyte
 - Each Flamingo connects via 2 GB/s channels to the forwarding ASIC
- Forwarding and Queuing ASICs
- Control Plan Processor (Pentium CPU)
- Battery module

Supported hosts:

For a list of supported operating systems, see the IBM TotalStorage SAN Volume Controller Web site at <http://www-1.ibm.com/servers/storage/support/virtual/2062-2300.html> and click **Supported software levels**.

Multipathing software:

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

Note: The multipath drivers, SDD and RDAC, can coexist on a host for certain operating systems.

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

www.ibm.com/storage/support/2062-2300

User interfaces:

The SAN Volume Controller Storage Software for Cisco MDS 9000 provides the following user interfaces:

- IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000, a graphical user interface (GUI) that supports flexible and rapid access to storage management information
- The IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000 command-line interface (CLI)
- The Cisco MDS 9000 CLI that supports commands for managing the switch in which the Caching Services Module is housed, as well as commands for cluster configurations and service functions.

Related topics:

- “Virtual disks” on page 22
- “Virtualization” on page 6

Common terms for the SAN Volume Controller Storage Software for Cisco MDS 9000

The following list defines terms that you will need to know to work with the SAN Volume Controller Storage Software for Cisco MDS 9000:

Properties:

Console for SAN Volume Controller for Cisco MDS 9000

The graphical user interface for the SAN Volume Controller Storage Software for Cisco MDS 9000.

SAN Volume Controller Storage Software for Cisco MDS 9000

A version of the SAN Volume Controller Storage Software that operates on the Cisco MDS 9000 Caching Services Module hardware platform for the Cisco MDS 9000 series of fibre-channel switches and integrates with Cisco-provided services.

Cisco MDS 9000 Caching Services Module

A specialized Cisco MDS 9000 fibre-channel line card that contains a pair of storage engines, or nodes on which the SAN Volume Controller Storage Software for Cisco MDS 9000 operates.

SAN Volume Controller for Cisco MDS 9000

The solution or system provided by the combination of the software (Storage Software for Cisco MDS 9000) and the hardware (Cisco MDS 9000 Caching Services Module).

Console workstation

The workstation where the Console for SAN Volume Controller for Cisco MDS 9000 is locally installed.

Virtualization

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

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

- Virtualization at the *server* level is performed by managing volumes on the operating systems servers. An increase in the amount of logical storage over physical storage is suitable for environments that do not have storage networks.
- Virtualization at the *storage device* level is in common use. Striping, mirroring, and redundant arrays of independent disks (RAIDs) are used by almost all disk subsystems. This type of virtualization can range from simple RAID controllers to advanced volume management such as that provided by the IBM TotalStorage Enterprise Storage Server (ESS) or by Log Structured Arrays (LSA). The Virtual Tape Server (VTS) is another example of virtualization at the device level.
- Virtualization at the *fabric* level enables storage pools to be independent of the servers and the physical components that make up the storage pools. One management interface can be used to manage different storage systems without affecting the servers. The SAN Volume Controller Storage Software for Cisco MDS 9000 is used to perform virtualization at the fabric level.
- Virtualization at the *file system* level provides the highest benefit because data is shared, allocated, and protected, not volumes.

Virtualization is a radical departure from traditional storage management. In traditional storage management, storage is attached directly to a host system, which

controls storage management. SANs introduced the principle of networks of storage, but storage is still primarily created and maintained at the RAID subsystem level. Multiple RAID controllers of different types require knowledge of, and software that is specific to, the given hardware. Virtualization brings a central point of control for disk creation and maintenance. It brings new ways of handling storage maintenance.

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

Types of virtualization:

Virtualization can be performed either asymmetrically or symmetrically:

Asymmetric

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

Symmetric

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

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

Note: The SAN Volume Controller Storage Software for Cisco MDS 9000 implements fabric-level *virtualization*. Within the context of the SAN Volume Controller Storage Software for Cisco MDS 9000 and throughout this document, *virtualization* refers to symmetric fabric-level virtualization.




Figure 1. Levels of virtualization

Related topics:

- Chapter 29, “Configuring and servicing storage subsystems,” on page 217
- “Virtual disks” on page 22

Asymmetric virtualization

In asymmetric virtual storage networks, the data flow, (2) in the figure below, is separated from the control flow, (1). A separate network or SAN link is used for control purposes. The metadata server contains all the mapping and locking tables while the storage devices contain only data. Because the flow of control is separated from the flow of data, I/O operations can use the full bandwidth of the SAN. A separate network or SAN link is used for control purposes. There are disadvantages, however, to asymmetric virtualization.

The disadvantages to asymmetric virtualization include:

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




Figure 2. Asymmetrical virtualization

Related topics:

- “Virtualization” on page 6
- “Symmetric virtualization”

Symmetric virtualization

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

In symmetric virtual storage networks (see Figure 3), data and control both flow over the same path. Because the separation of the control from the data occurs in the data path, the storage can be pooled under the control of the virtualization engine. The virtualization engine performs the logical-to-physical mapping.




Figure 3. Symmetrical virtualization

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

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

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

The IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000 uses symmetric virtualization. Single virtualization engines, which are known as nodes, are combined to create clusters. Each cluster can contain between two and four nodes.

Related topics:

- “Virtualization” on page 6
- “Asymmetric virtualization” on page 8

Chapter 2. Object overview

Nodes are deployed in pairs known as **I/O groups**. Each node may be in only one I/O group. A **cluster** may include one or two I/O groups.

Virtual disks (VDisks) are logical disks that are presented to the SAN by nodes. Virtual disks are also associated with an I/O group. The nodes in the I/O group provide access to the virtual disks in the I/O group. When an application server performs I/O to a virtual disk, it has the choice of accessing the virtual disk via either of the nodes in the I/O group. As each I/O group only has two nodes, the distributed cache the SAN Volume Controller Storage Software for Cisco MDS 9000 provides is only 2-way.

The nodes are part of an MDS 9000 line card that is seated within the switch. This line card, which is called a caching services module, contains a battery module that allows cached data to be dumped to a hard disk in the event of line power failure.

The nodes in a cluster see the storage presented by SAN-attached **storage subsystems** as a number of disks, known as **managed disks (MDisks)**. Because the SAN Volume Controller Storage Software for Cisco MDS 9000 does not attempt to provide recovery from physical disk failures within the backend disk controllers, a managed disk is usually, but not necessarily, a redundant array of independent disks (RAID) array.

Each managed disk is divided into a number of **extents** which are numbered from 0, sequentially, from the start to the end of the managed disk. The extent size must be specified when creating an MDisk group.

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

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

The SAN Volume Controller Storage Software for Cisco MDS 9000 detects the Fibre Channel ports that are connected to the SAN. These correspond to the Host Bus Adapter (HBA) Fibre Channels worldwide port names (WWPNs) that are present in the application servers. The SAN Volume Controller Storage Software for Cisco MDS 9000 allows you to create logical host objects that group together WWPNs belonging to a single application server or multiple application servers.

Attention: Host Bus Adapters of different vendors, models, or operating systems should not be placed in the same zone or the same host object.

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

Nodes and clusters

Nodes are deployed in pairs called I/O groups. One node in the cluster is designated the configuration node but each node in the cluster holds a copy of the cluster state information.

Related topics:

- Chapter 1, “SAN Volume Controller Storage Software for Cisco MDS 9000,” on page 3
- “Virtualization” on page 6

Clusters

A cluster can consist of at least two nodes with a maximum of four nodes. Therefore, you can assign up to four SAN Volume Controller Storage Software for Cisco MDS 9000 nodes to one cluster. Some service actions can be performed at node level, but all configuration is replicated across all nodes in the cluster. Because configuration is performed at the cluster level, an IP address is assigned to the cluster instead of to each of the nodes.

Because all of your configuration and service actions are replicated across all nodes in a cluster, you can take advantage of the virtualization and the advanced features of the SAN Volume Controller Storage Software for Cisco MDS 9000 after configuring the cluster.

A cluster can span more than one CISCO MDS 9000 switch.

A switch can contain more than one cluster.

The nodes on a Cisco MDS 9000 Caching Services Module do not have to be in the same cluster.

Cluster state and the configuration node:

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

Figure 4 on page 13 shows an example cluster containing four nodes. The cluster state shown in the grey box does not actually exist, instead each node holds a copy of the entire cluster state.

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

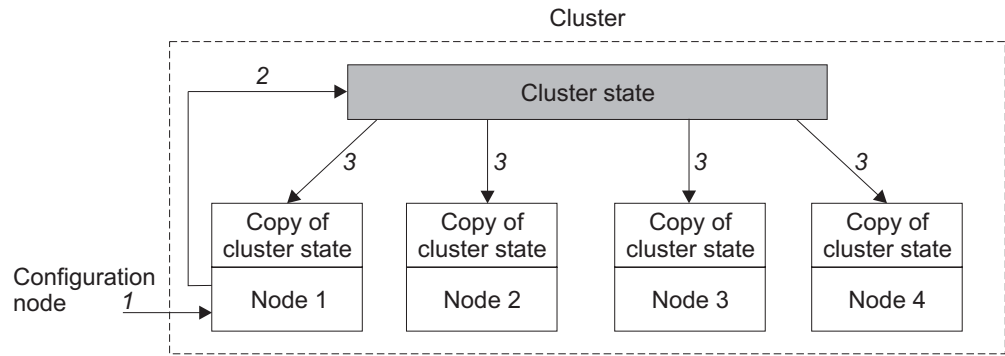


Figure 4. Cluster, nodes, and cluster state.

Related topics:

- “Nodes and clusters” on page 12

Nodes

Nodes are deployed in pairs known as I/O groups. Each node may be in only one I/O group. Two I/O groups comprise a cluster.

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

Table 4 describes the operational states of a node.

Table 4. Node state

State	Description
Adding	The node was added to the cluster but is not yet synchronized with the cluster state (see Note).
Deleting	The node is in the process of being deleted from the cluster.
Online	The node is operational, assigned to a cluster, and has access to the fibre-channel SAN fabric.
Offline	The node is not operational. The node was assigned to a cluster but is not available on the fibre-channel SAN fabric.
Pending	The node is transitioning between states and, in a few seconds, will move to one of the other states.
<p>Note: It is possible that a node can stay in the Adding state for a long time. If this is the case, delete the node and then re-add it. However, you should wait for at least 30 minutes before doing this. If the node that has been added is at a lower code level than the rest of the cluster, the node will be upgraded to the cluster code level, which can take up to 20 minutes. During this time the node will be shown as adding.</p>	

Configuration node

The configuration node is a focal point for configuration commands, and it manages the data that describes the cluster configuration.

If the configuration node fails, the cluster chooses a new configuration node. This action is called configuration node failover. The switch that contains the new node takes over the cluster IP address. Thus you can access the cluster through the same IP address although the original configuration node has failed. During the failover, there is a short period when you cannot use the command line tools or Console for SAN Volume Controller for Cisco MDS 9000.

The figure below shows an example cluster containing four nodes. Node 1 has been designated the configuration node. User requests (1) are targeted at Node 1. This may result in requests (2) being targeted at the other nodes in the cluster, and data being returned to Node 1.




Figure 5. Configuration node

Related topics:

- Chapter 1, “SAN Volume Controller Storage Software for Cisco MDS 9000,” on page 3
- “Virtualization” on page 6

I/O groups and Internal Battery Backup

Virtual disks are logical disks that are presented to the SAN by SAN Volume Controller Storage Software for Cisco MDS 9000 nodes. Virtual disks are also associated with an I/O group. The SAN Volume Controller Storage Software for Cisco MDS 9000 contains internal battery backup units to provide data integrity in the event of a cluster wide power failure.

Input/Output (I/O) groups

It usually contains two SAN Volume Controller Storage Software for Cisco MDS 9000 nodes for availability purposes. However, depending on the configuration, an I/O group may be empty or just contain a single node. Each node is associated with only one I/O group, and each virtual disk (VDisk) is associated with only one I/O group. The nodes in the I/O group provide access to the VDIs in the I/O group.

When an application server performs I/O to a virtual disk, it has the choice of accessing the virtual disk via either of the nodes in the I/O group. A virtual disk can specify a preferred node. This is specified when the virtual disk is created. This is the node through which a virtual disk should normally be accessed. As each I/O group only has two nodes, the distributed cache in the SAN Volume Controller Storage Software for Cisco MDS 9000 need only be 2-way. When I/O is performed to a virtual disk, the node that processes the I/O duplicates the data onto the partner node that is in the I/O group.

I/O traffic for a particular virtual disk is, at any one time, handled exclusively by the nodes in a single I/O group.

The figure below shows an example I/O group. A write operation from a host is shown (item 1), that is targeted for virtual disk A. This write is targeted at the preferred node, Node 1 (item 2). The write is cached and a copy of the data is made in the partner node, Node 2's cache (item 3). The write is now complete, so far as the host is concerned. At some later time, the data is written, or destaged, to storage (item 4). The nodes in an I/O group should not be on the same Cisco MDS 9000 Caching Services Module in order to avoid a single point of failure.

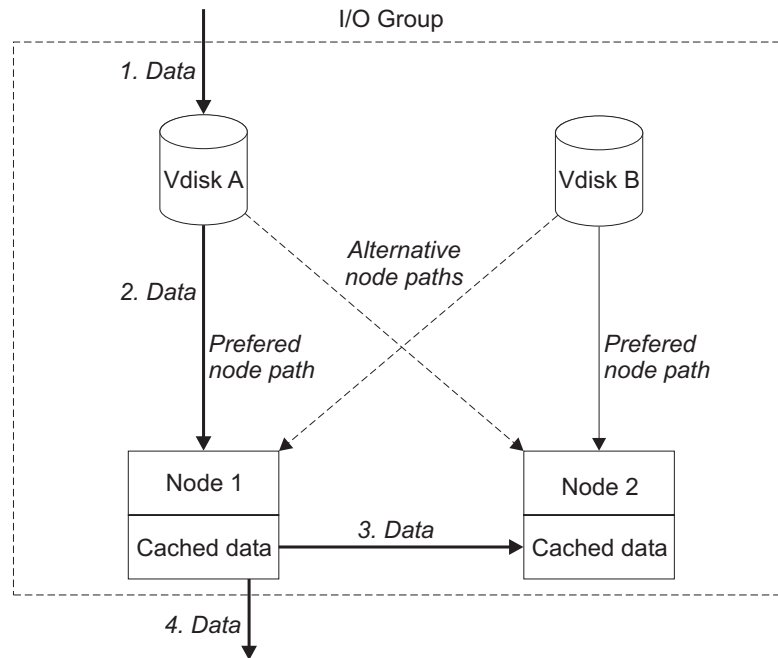


Figure 6. I/O group example

When a node fails within an I/O group, the other node in the I/O group will take 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 a node has failed in an I/O group, the cache is flushed to the disk and then goes into write-through mode. Therefore, any writes for the virtual disks that are assigned to this I/O group are not cached; it is sent directly to the storage device. If both nodes in an I/O group go offline, the virtual disks that are assigned to the I/O group cannot be accessed.

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

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

Related topics:

- “Nodes and clusters” on page 12
- “Virtual disks” on page 22

- “Related publications” on page xi

Internal battery backup overview

Note: The Cisco MDS 9000 Caching Services Module internal backup batteries are an integral part of the SAN Volume Controller Storage Software for Cisco MDS 9000 solution. The SAN Volume Controller Storage Software for Cisco MDS 9000 internal backup batteries must be used in accordance with documented guidelines.

To provide full redundancy and concurrent maintenance, the SAN Volume Controller Storage Software for Cisco MDS 9000 must be installed on at least two different Cisco MDS 9000 Caching Services Modules.

Each node on the Cisco MDS 9000 Caching Services Module is connected to its own battery. The batteries are mounted on the Cisco MDS 9000 Caching Services Module itself, as are two hard disk drives. In the event of a power failure, the batteries allow the cached memory on the nodes to dump to their associated hard disk.

A SAN Volume Controller Storage Software for Cisco MDS 9000 is not fully operational until the internal batteries' charge state is sufficient to power the Cisco MDS 9000 Caching Services Module nodes for long enough to permit all memory to be saved to the hard disk. The internal battery backup has sufficient capacity to save all the data on the hard disk SAN Volume Controller Storage Software for Cisco MDS 9000 at least twice. For a fully-charged internal battery backup, even after battery capacity has been used to power the SAN Volume Controller Storage Software for Cisco MDS 9000s while they save DRAM data, sufficient battery capacity will remain to let the SAN Volume Controller Storage Software for Cisco MDS 9000s become fully operational as soon as input power is restored.

It is very important that each of the two nodes in the I/O group be on different Cisco MDS 9000 Caching Services Modules.

Note: You cannot add nodes in the same Cisco MDS 9000 Caching Services Module to the same I/O group.

This configuration ensures that the cache and cluster state information is protected against the failure of a Cisco MDS 9000 Caching Services Module.

When nodes are added to the cluster, you must specify the I/O group they will join.

Storage subsystems and managed disks

Storage subsystems

A storage subsystem is a device that coordinates and controls the operation of one or more disk drives and synchronizes the operation of the drives with the operation of the system as a whole.

Storage subsystems attached to the SAN fabric provide the physical storage devices that the cluster detects as managed disks. These are usually RAID arrays, because the SAN Volume Controller Storage Software for Cisco MDS 9000 does

not attempt to provide recovery from physical disk failures within the storage subsystem. The nodes in the cluster can be in only one fibre-channel SAN or in up to four virtual SANs (VSANs).

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

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

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

Many storage subsystems allow the storage provided by a RAID array to be divided up into many SCSI logical units (LUs) that are presented on the SAN. With the SAN Volume Controller Storage Software for Cisco MDS 9000 it is recommended that storage subsystems are configured to present each RAID array as a single SCSI LU that will be recognized by the SAN Volume Controller Storage Software for Cisco MDS 9000 as a single managed disk. The virtualization features of the SAN Volume Controller Storage Software for Cisco MDS 9000 can then be used to divide up the storage into virtual disks.

Some storage subsystems allow the exported storage to be increased in size. The SAN Volume Controller Storage Software for Cisco MDS 9000 will not use this extra capacity. Instead of increasing the size of an existing managed disk, a new managed disk should be added to the managed disk group and the extra capacity will be available for the SAN Volume Controller Storage Software for Cisco MDS 9000 to use.

Attention: If you delete a RAID that is being used by the SAN Volume Controller Storage Software for Cisco MDS 9000, the MDisk group will go offline and the data in that group will be lost.

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

The cluster detects and provides a view of the storage subsystems that the SAN Volume Controller Storage Software for Cisco MDS 9000 supports. The cluster can also determine which MDisks each subsystem has and can provide a view of MDisks filtered by device. This view enables you to associate the MDisks with the RAID arrays that the subsystem presents.

Note: The SAN Volume Controller Storage Software for Cisco MDS 9000 supports storage that is internally configured as a RAID array. However, it is possible to configure a storage subsystem as a non-RAID device. RAID provides redundancy at the disk level. For RAID devices, a single physical disk failure does not cause an MDisk failure, an MDisk group failure, or a failure in the virtual disks (VDisks) that were created from the MDisk group.

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

Related topics:

- “Managed disks”
- “Managed disk groups” on page 20
- “Virtual disks” on page 22

Managed disks

A managed disk (MDisk) is a logical disk (typically a RAID array or partition thereof) that a storage subsystem has exported to the SAN fabric to which the nodes in the cluster are attached. A managed disk might, therefore, consist of multiple physical disks that are presented as a single logical disk to the SAN. A managed disk always provides usable blocks of physical storage to the cluster even if it does not have a one-to-one correspondence with a physical disk.

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

Access modes:

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

Unmanaged

The MDisk is not used by the cluster.

Managed

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

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

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

Table 5 describes the operational states of a managed disk:

Table 5. Managed disk status

Status	Description
--------	-------------

Table 5. Managed disk status (continued)

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

Extents:

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

Attention: If your fabric is undergoing transient link breaks or you have been replacing cables or connections in your fabric, you might see one or more MDisks change to the degraded status. If an I/O operation was attempted during the link breaks and the same I/O failed several times, the MDisk will be partially excluded and will change to a status of degraded. You should include the MDisk to resolve the problem. You can include the MDisk by either selecting the Include MDisk task from the Work with Managed Disks - Managed Disk panel in the SAN Volume Controller for Cisco MDS 9000, or issue the following command:

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

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

Related topics:

- “Storage subsystems” on page 16

Managed disk groups and virtual disks (VDisks)

Virtual disks are created from the extents of managed disks. Only managed disks that are in the same managed disk group can contribute extents to a virtual disk.

Managed disk groups

An *MDisk group* is a collection of MDisks that jointly contain all the data for a specified set of virtual disks (VDisks). All MDisks in a group are split into extents of the same size. VDIs are created from the extents that are available in the group.

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

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

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

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

Table 6 describes the operational states of an MDisk group.

Table 6. Managed disk group status

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

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

Consider the following guidelines when you create MDisk groups:

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

Extent:

To track the space that is available, the SAN Volume Controller Storage Software for Cisco MDS 9000 divides each MDisk in an MDisk group into chunks of equal size. These chunks are called *extents*, and are indexed internally. Extent sizes can be 16, 32, 64, 128, 256, or 512 MB.

You must specify the extent size when you create a new MDisk group. You cannot change the extent size later; it must remain constant throughout the lifetime of the MDisk group. MDisk groups can have different extent sizes. However, different extent sizes can place restrictions on the use of data migration. The choice of extent size affects the total amount of storage that can be managed by a SAN Volume Controller Storage Software for Cisco MDS 9000 cluster. Table 7 shows the maximum amount of storage that can be managed by a cluster for each extent size. Because the SAN Volume Controller Storage Software for Cisco MDS 9000 allocates a whole number of extents to each virtual disk that is created, using a larger extent size can increase the amount of wasted storage at the end of each virtual disk. Larger extent sizes also reduce the ability of the SAN Volume Controller Storage Software for Cisco MDS 9000 to distribute sequential I/O workloads across many managed disks. Therefore, larger extent sizes might reduce the performance benefits of virtualization.

Table 7. Capacities of the cluster given extent size

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

Table 7. Capacities of the cluster given extent size (continued)

Extent size	Maximum storage capacity of cluster
512 MB	2 PB

Figure 7 shows an MDisk group containing four MDisks.




Figure 7. MDisk group

Related topics:

- “Managed disks” on page 18
- “Virtual disks”

Virtual disks

A *VDisk* is a logical disk that the cluster presents to the storage area network (SAN). Application servers on the SAN access VDisks, not managed disks (MDisks). VDisks are created from a set of extents in an MDisk group. There are three types of VDisks: striped, sequential, and image.

Types:

You can create the following types of VDisks:

Striped

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

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

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

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

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

Figure 8 shows an example of a managed disk group containing three MDisks. This figure also shows a striped virtual disk created from the extents available in the group.

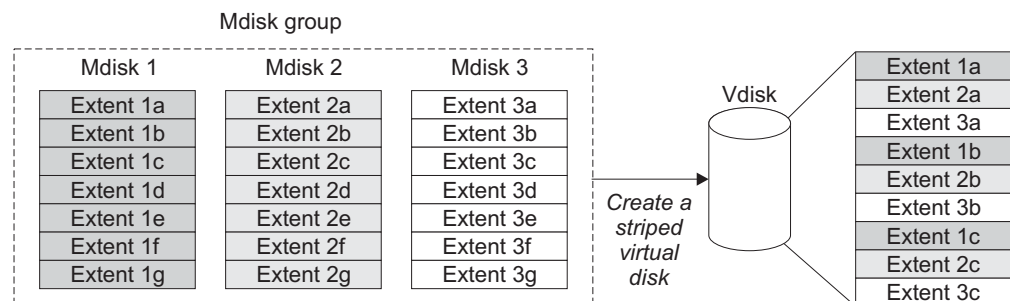


Figure 8. Managed disk groups and VDIs

Sequential

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

Image

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

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

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

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

MDisks that contain existing data have an initial mode of unmanaged, and the cluster cannot determine whether it contains a LUN that contains a complete disk or a new disk ready to be used.

A virtual disk can have one of three states. Table 8 describes the different states of a virtual disk:

Table 8. Virtual disk status

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

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

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

Related topics:

- “Virtualization” on page 6

Hosts and virtual (VDisk) mappings

Host objects

Creating a host in a cluster results in the creation of a logical host object. A logical host object has one or more worldwide port names (WWPNs) assigned to it.

Generally, a logical host object is associated with a physical host system. However, a single logical host object can have WWPNs from multiple physical host systems that are assigned to it.

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

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

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

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

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

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

Node Login Counts:

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

Virtual disk-to-host mapping

Virtual disk-to-host mapping is similar in concept to logical unit number (LUN) mapping or masking. LUN mapping is the process of controlling which hosts have access to specific logical units (LUs) within the disk controllers. LUN mapping is typically done at the disk controller level. Virtual disk-to-host mapping is done at the SAN Volume Controller Storage Software for Cisco MDS 9000 level.

Application servers can only access VDisks that have been made accessible to them. The SAN Volume Controller Storage Software for Cisco MDS 9000 detects the fibre-channel ports that are connected to the SAN. These correspond to the host bus adapter (HBA) worldwide port names (WWPNs) that are present in the application servers. The SAN Volume Controller Storage Software for Cisco MDS 9000 enables you to create logical hosts that group together WWPNs belonging to a single application server. VDisks can then be mapped to a host. The act of mapping a virtual disk to a host makes the virtual disk accessible to the WWPNs in that host, and hence the application server itself.

VDisks and host mappings:

The SAN concept known as LUN masking usually requires device driver software in each host. The device driver software masks the LUNs as instructed by the user. After the masking has been done, only some disks are visible to the operating system. The SAN Volume Controller Storage Software for Cisco MDS 9000 performs a similar function, but, by default, it presents to the host only those VDisks that are mapped to that host. You must therefore map the VDisks to the hosts that are to access those VDisks.

Each host mapping associates a virtual disk with a host object and allows all HBA ports in the host object to access the virtual disk. You can map a virtual disk to multiple host objects. When a mapping is created, multiple paths might exist across the SAN fabric from the hosts to the SAN Volume Controller Storage Software for Cisco MDS 9000s that are presenting the virtual disk. Most operating systems present each path to a virtual disk as a separate storage device. The SAN Volume Controller Storage Software for Cisco MDS 9000, therefore, needs the IBM Subsystem Device Driver (SDD) software to be running on the host. This software handles the many paths that are available to the virtual disk and presents a single storage device to the operating system.

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

Figure 9 and Figure 10 on page 27 show two VDisks, and the mappings that exist between the host objects and these VDisks.

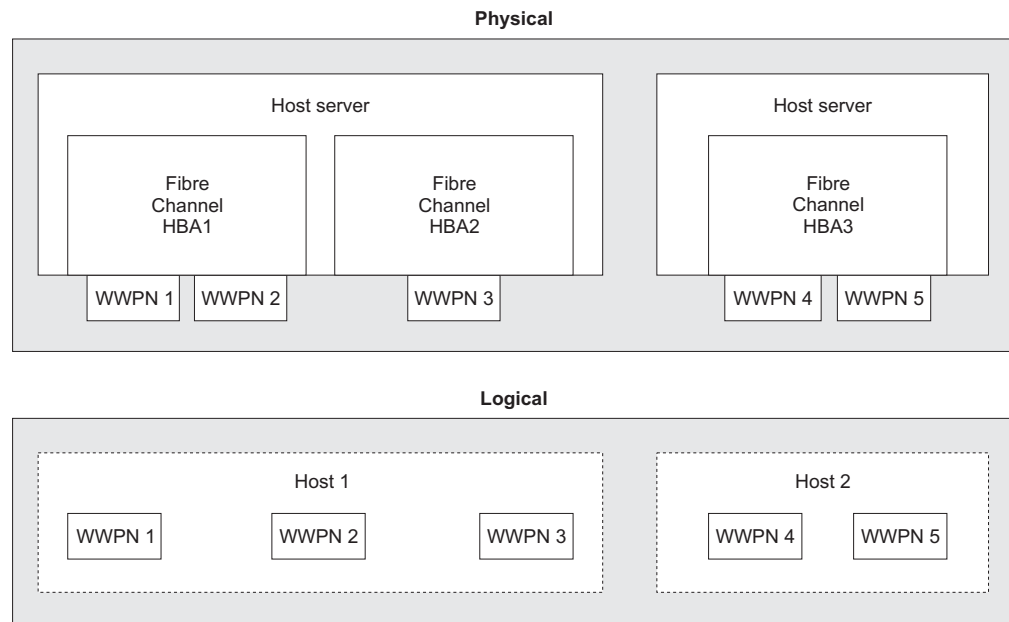


Figure 9. Hosts, WWPNs, and VDisks

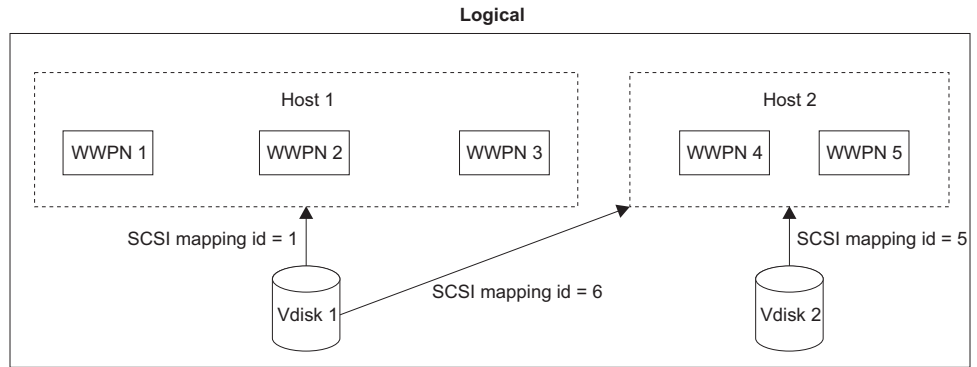


Figure 10. Hosts, WWPNs, VDisks and SCSI mappings

Related topics:

- “Managed disks” on page 18
- “Virtual disks” on page 22

Chapter 3. Copy Services

One is called FlashCopy and the other is synchronous Remote Copy (which is analogous to Peer-to-Peer Remote Copy or PPRC).

FlashCopy

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

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

Source VDIsks and target VDIsks must meet the following requirements:

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

Related topics:

- “FlashCopy consistency groups” on page 33
- “FlashCopy mappings”
- “Virtual disks” on page 22

FlashCopy mappings

Because FlashCopy copies one VDisk to another VDisk, the SAN Volume Controller Storage Software for Cisco MDS 9000 needs to be aware of that relationship. particular virtual disk can take part in only one mapping; that is, a virtual disk can be the source or target of only one mapping. You cannot, for example, make the target of one mapping the source of another mapping.

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

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

A FlashCopy mapping can be created between any two virtual disks in a cluster. It is not necessary for the virtual disks to be in the same I/O group or managed disk group. When a FlashCopy operation is started, a checkpoint is made of the source virtual disk. No data is actually copied at the time a start occurs. Instead, the

checkpoint creates a bitmap that indicates that no part of the source virtual disk has yet been copied. Each bit in the bitmap represents one region of the source virtual disk. Such a region is called a grain.

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

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

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

When a mapping is started and the background copy rate is greater than zero, the unchanged data is copied to the target, and the bitmap is updated to show that the copy has occurred. After a time, the length of which depends on the priority given and the size of the virtual disk, the whole virtual disk is copied to the target. The mapping returns to the idle/copied state. You can restart the mapping at any time to create a new copy at the target; the process copy starts again.

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

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

FlashCopy mapping states:

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

Idle or copied

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

Copying

The copy is in progress.

Prepared

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

Preparing

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

Stopped

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

Suspended

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

Before you start the mapping, you must prepare it. By preparing the mapping, you ensure that the data in the cache is destaged to disk and that a consistent copy of the source exists on disk. At this time the cache goes into write-through mode. That is, data that is written to the source is not cached in the SAN Volume Controller Storage Software for Cisco MDS 9000s; it passes straight through to the managed disks. The prepare operation for the mapping might take you a few minutes; the actual length of time depends on the size of the source virtual disk. You must coordinate the prepare operation with the operating system. Depending on the type of data that is on the source virtual disk, the operating system or application software might also cache data write operations. You must flush, or synchronize, the file system and application program before you prepare for, and finally start, the mapping.

For customers who do not need the complexity of consistency groups, the SAN Volume Controller Storage Software for Cisco MDS 9000 allows a FlashCopy mapping to be treated as an independent entity. In this case the FlashCopy mapping is known as a stand alone mapping. For FlashCopy mappings which have been configured in this way, the **Prepare** and **Start** commands are directed at the FlashCopy mapping name rather than the consistency group ID.

Veritas Volume Manager:

For FlashCopy target VDIs, the SAN Volume Controller Storage Software for Cisco MDS 9000 sets a bit in the inquiry data for those mapping states where the target VDisk could be an exact image of the source VDisk. Setting this bit enables the Veritas Volume Manager to distinguish between the source and target VDIs and thus provide independent access to both.

Related topics:

- “FlashCopy” on page 29
- “FlashCopy consistency groups” on page 33
- “Virtual disks” on page 22

FlashCopy mapping events

Create

A new FlashCopy mapping is created between the specified source virtual disk and the specified target virtual disk. The various supported parameters are also described there. The operation fails if either the source or target virtual disks are already a member of a FlashCopy mapping. The operation fails if the SAN Volume Controller Storage Software for Cisco MDS 9000 has insufficient bitmap memory. The operation also fails if the source and target virtual disks are different sizes.

Prepare

The prepare command is directed to either a consistency group for FlashCopy mappings which are members of a normal consistency group or

to the mapping name for FlashCopy mappings which are members of the special consistency group 0. The prepare command places the FlashCopy mapping in the preparing state.

It is important to note that the act of preparing for start may corrupt any data which previously resided on the target virtual disk since cached writes are discarded. Even if the FlashCopy mapping is never started, the data from the target may have been logically changed by the act of preparing for start.

Flush done

The FlashCopy relationship moves from the preparing state to the prepared state automatically once all cached data for the source has been flushed and all cached data for the target has been invalidated.

Start When all the FlashCopy mappings in a consistency group are in the prepared state, the FlashCopy relationships can be started. Some other FlashCopy products refer to this event as starting the FlashCopy.

The start of all of the FlashCopy mappings in the consistency group must be synchronized correctly with respect to I/Os directed at the virtual disks to preserve the cross volume consistency group. This is achieved as follows:

During the **start** command:

- New reads and writes to all source virtual disks in the consistency group are paused in the cache layer until all ongoing reads and writes below the cache layer have been completed.
- Once all FlashCopy mappings in the consistency group have been paused, internal cluster state is set to allow FlashCopy operations.
- Once all FlashCopy mappings in the consistency group have had their cluster state set, read and write operations are unpaused on the source virtual disks.
- The target virtual disks are brought online.

As part of the **start** command, read and write caching is enabled for both the source and target virtual disks.

Modify

A FlashCopy mapping has three properties that can be modified. These are the background copy rate, the consistency group, and the name of the FlashCopy mapping. The background copy rate can be modified in any state but attempting to modify the consistency group in any state other than idling, copied, or stopped will fail.

Stop There are two mechanisms by which a FlashCopy mapping can be stopped:

1. Either you have issued a command; or
2. An input/output (I/O) error has occurred.

Delete This command requests that the specified FlashCopy mapping be deleted. If the FlashCopy mapping is in the stopped state, the force flag must be used.

Deleting a FlashCopy mapping in the stopped state may allow unflushed write data from the cache to be destaged to what was the target virtual disk. This does not affect the data integrity of the system because following a forced delete, nothing can be certain about the contents of the target virtual disk. The data contained in the target virtual disk could be anything.

The destaging of old data to what was the target virtual disk does not affect the future use of the virtual disk because any new data will be written over this old data, in the cache or on disk.

Flush failed

If the flush of data from the cache cannot be completed then FlashCopy mapping will enter the stopped state.

Copy complete

Once every grain of the source and target has been copied, the source and target are independent and the state machine enters the copied state. The FlashCopy mapping is not automatically deleted at this time and can be re-activated by preparing and starting again.

Bitmap Online/Offline

The node has failed.

FlashCopy consistency groups

The consistency group is specified when the mapping is created. You can also change the consistency group later. When you use a consistency group, you prepare and trigger that group instead of the various mappings. This ensures that a consistent copy is made of all the source VDisks. Mappings that you want to control at an individual level instead of at a consistency group level, should not be put into a consistency group. These mappings are known as stand-alone mappings.

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

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

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

FlashCopy consistency-group states:

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

Idle or copied

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

Copying

The copy is in progress.

Prepared

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

Preparing

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

Stopped

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

Suspended

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

Related topics:

- “FlashCopy” on page 29
- “FlashCopy mappings” on page 29
- “Dependent writes”
- “Virtual disks” on page 22

Dependent writes

Think about the following typical sequence of write operations for a data base update transaction.

1. Run a write operation to update the data base log so that it indicates that a data base update is about to take place.
2. Run a second write operation to update the data base.
3. Run a third write operation to update the data base log so that it indicates that the data base update has completed successfully.

The database ensures correct ordering of these writes by waiting for each step to complete before starting the next. If, however, the database log (updates 1 and 3) and the database itself (update 2) are on different virtual disks and a FlashCopy mapping is started during this update, then the possibility that the database itself is copied slightly before the database log resulting in the target virtual disks seeing writes (1) and (3) but not (2) must be excluded. In this case, if the database were restarted from a backup made from the FlashCopy target disks, the data base log would indicate that the transaction had completed successfully when, in fact, that is not the case. The transaction would be lost and the integrity of the data base would be in question.

It may thus be the case that in order to create a consistent image of user data it is necessary to perform a FlashCopy operation on multiple virtual disks as an atomic operation. In order to meet this need, the SAN Volume Controller Storage Software for Cisco MDS 9000 supports the concept of a consistency group. A consistency group contains a number of FlashCopy mappings. A consistency group can contain an arbitrary number of FlashCopy mappings up to the maximum number of FlashCopy mappings supported by a SAN Volume Controller Storage Software for Cisco MDS 9000 cluster. The SAN Volume Controller Storage Software for Cisco MDS 9000 allows the **start** command which causes the point in time copy to occur, to be directed at a consistency group. In this case all of the FlashCopy mappings in the consistency group are started at the same time, resulting in a point in time copy which is consistent across all of the FlashCopy mappings which are contained in the consistency group. The SAN Volume Controller Storage Software for Cisco MDS 9000 supports 128 consistency groups per cluster.

Operations on consistency groups

You can create, rename, delete, prepare, start, and stop consistency groups by using the command line tool that is described in *IBM TotalStorage SAN Volume Controller: Command-Line Interface User's Guide* or you can use the Console for SAN Volume Controller for Cisco MDS 9000.

FlashCopy limits

FlashCopy applications

An important use of FlashCopy is to take consistent backups of changing data. In this application, a FlashCopy is created to capture the data at a particular time. The resulting image of the data can be backed up, for example, to a tape device. When the copied data is on tape, the data on the FlashCopy target disks becomes redundant and can now be discarded. Usually in this backup condition, the target data can be handled as read only.

Another use of FlashCopy data is in the application testing. It is often very important to test a new version of an application with real business data before the existing production version of the application is updated or replaced. This testing reduces the risk that the updated application fails because it is not compatible with the actual business data that is in use at the time of the update. Such an application test might require write access to the target data.

Other uses of FlashCopy in the business environment include creating copies for auditing purposes and for data mining.

In the scientific and technical arena one way in which FlashCopy is employed is to create restart points for long running batch jobs. This means that if a batch job fails many days into its run it may be possible to restart the job from a saved copy of its data rather than re-running the entire multi-day job.

FlashCopy indirection layer

The act of starting a FlashCopy mapping causes this indirection layer to become active in the I/O path. This occurs as an atomic command across all FlashCopy mappings in the consistency group.

The indirection layer makes a decision about each I/O. This decision is based upon the following:

- the virtual disk and LBA to which the I/O is addressed,
- its direction (read or write)
- the state of an internal data structure, the flash copy bitmap.

The indirection layer either allows the I/O through to the underlying storage, redirects the I/O from the target virtual disk to the source virtual disk, or stalls the I/O while it arranges for data to be copied from the source virtual disk to the target virtual disk.

Grains and the FlashCopy bitmap

In the SAN Volume Controller Storage Software for Cisco MDS 9000, the grain size is 256 KB. The FlashCopy bitmap contains one bit for each grain. The bit records whether the associated grain has yet been split by copying the grain from the source to the target.

Source and target reads

Source reads:

Reads of the source are always passed through to the underlying source disk.

Target reads:

In order for FlashCopy to process a read from the target disk it must consult its bitmap. If the data being read has already been copied to the target then the read is sent to the target disk. If it has not, then the read is sent to the source disk. Clearly this algorithm requires that, while this read is outstanding, no writes are allowed to execute which would change the data being read from the source. The SAN Volume Controller Storage Software for Cisco MDS 9000 satisfies this requirement by using a cluster wide locking scheme.

FlashCopy limits the number of concurrent reads to an unsplit target grain to one. If more than one concurrent read to an unsplit target grain is received by the FlashCopy mapping layer, they will be serialized.

Writes to the source or target

A specific optimization is performed where an entire grain is written to the target virtual disk. In this case the new grain contents are written to the target virtual disk and if this succeeds the grain is marked as split in the flash copy bitmap without a copy from the source to the target having been performed. If the write fails, the grain is not marked as split.

FlashCopy limits

Up to 128 FlashCopy mappings are supported in a single cluster. A maximum of 16 TB of VDisk space (both source and target) may be participating in FlashCopy mappings in any one I/O group on a single cluster.

Background copy

This is a value between 1 and 100. The background copy rate can be changed when the FlashCopy mapping is in any state.

If “NOCOPY” is specified, then background copy is disabled. One use of this is for short-lived FlashCopy mappings which are to be used for backup purposes only. Since the source data set is not expected to change much during the lifetime of the FlashCopy mapping, it is more efficient in terms of managed disk I/Os not to perform a background copy.

The relationship of the background copy rate value to the attempted number of grains to be split per second is given by the following table. (A grain is the unit of data represented by a single bit, which is 256K.)

Table 9. Background copy

User-specified value	KB/sec	Grains/sec
1 - 10	128	0.5
11 - 20	256	1
21 - 30	512	2
41 - 50	2048	8
91 - 100	64 MB	256

The grains/sec numbers represent goals that the code tries to achieve. The SAN Volume Controller Storage Software for Cisco MDS 9000 will be unable to achieve these goals if insufficient bandwidth is available from the SAN Volume Controller Storage Software for Cisco MDS 9000 nodes to the physical disks making up the managed disks after taking into account the requirements of foreground I/O. If this situation arises, then background copy I/O will contend for resources on an equal basis with I/O arriving from hosts. Both will tend to see an increase in latency and consequential reduction in throughput with respect to the situation had the bandwidth not been limited.

Degradation will be graceful. Both background copy and foreground I/O will continue to make forward progress, and will not stop, hang or cause the node to fail.

The background copy is performed by one of the nodes belonging to the I/O group in which the source virtual disk resides. This responsibility is failed over to the other node in the I/O group in the event of the failure of the node performing the background copy.

The background copy is performed backward; that is, it starts with the grain containing the highest logical block numbers (LBAs) and works backwards towards the grain containing LBA 0. This is done to avoid any unwanted interactions with sequential write streams from the using application.

Host considerations for FlashCopy integrity

The SAN Volume Controller Storage Software for Cisco MDS 9000 FlashCopy functionality transfers a point-in-time copy of one virtual disk onto a designated target virtual disk of the same size. Both virtual disks must already be created. All the data in the source virtual disk is copied to the destination virtual disk. This includes operating system control information as well as application data and meta-data. Because all the data is copied, some operating systems will not allow a source disk and a target disk to reside on the same host. In order to ensure the integrity of the copy made, it is necessary to completely flush the host cache of any outstanding reads or writes before proceeding with the FlashCopy. Host cache flushing is ensured by unmounting the source virtual disks from the source host before starting the FlashCopy.

Steps:

Perform the following steps to flush data from your host volumes and perform the FlashCopy:

1. If you are using UNIX or Linux operating systems, perform the following steps:
 - a. Quiesce all applications to the source volumes you wish to FlashCopy.
 - b. Use the **umount** command to unmount the designated drives.
 - c. Prepare and start the FlashCopy for those unmounted drives.
 - d. Mount back your volumes with the mount command and resume your applications.
2. If you are using Windows operating system using drive letter changes, perform the following steps:
 - a. Quiesce all applications to the source volumes you wish to FlashCopy.
 - b. Go into your disk management window and remove the drive letter on each drive to be copied (this unmounts the drive).
 - c. Prepare and start the FlashCopy for those unmounted drives.

- d. Mount back your volumes by restoring the drive letters, and resume your applications.

If you are using the **chkdsk** command, perform the following steps:

- a. Quiesce all applications to the source volumes you wish to FlashCopy.
- b. Issue the **chkdsk /x** command on each drive to be copied (the /x option will unmount, scan, and remount the volume).
- c. Ensure that all applications to the source volumes are still quiesced.
- d. Prepare and start the FlashCopy for those unmounted drives.

Note: If you can ensure that no reads and writes will be issued to the source volumes after unmounting, you can immediately remount and then perform the FlashCopy.

Because the target disks will be overwritten with a complete image of the source disks, it is important that any data held in the host operating system (or application) caches for the target disks is discarded before the FlashCopy mappings are started. The easiest way to ensure that no data is held in these caches is to unmount the target disks prior to starting FlashCopy.

Some operating systems and applications provide facilities to stop I/O operations and to ensure that all data is flushed from caches on the host. If these facilities are available then they can be used to prepare and start a FlashCopy in a less disruptive manner. Refer to your host and application documentation for details.

Some operating systems are unable to use a copy of a virtual disk without an additional step, which is called synthesis. Synthesis performs some transformation on the operating system meta-data on the target virtual disk to allow the operating system to use the disk. Refer to your host documentation on how to detect and mount the copied virtual disks.

Remote Copy

Although the application only writes to a single virtual disk, the SAN Volume Controller Storage Software for Cisco MDS 9000 maintains two copies of the data. If the copies are separated by a significant distance, then the remote copy can be used as a backup for disaster recovery. A prerequisite for the SAN Volume Controller Storage Software for Cisco MDS 9000 Remote Copy operations between two clusters is that the SAN fabric to which they are attached provides adequate bandwidth between the clusters.

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

Remote Copy supports the following features:

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

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

- Intercluster and intracluster Remote Copy can be used concurrently within a cluster.

Related topics:

- “Remote Copy relationships”
- “Remote Copy partnerships”

Synchronous Remote Copy

The host application writes data to the primary VDisk but does not receive the final status on the write operation until the data is written to the secondary VDisk. For disaster recovery, this mode is the only practical mode of operation because a consistent copy of the data is maintained.

Related topics:

- “Remote Copy” on page 38

Remote Copy partnerships

The SAN Volume Controller Storage Software for Cisco MDS 9000 needs to know not only about the relationship between the two VDisks but also about the relationship between the two clusters. A Remote Copy partnership defines the relationship between the two clusters.

To establish a cluster partnership between two clusters it is necessary to issue the **svctask mkpartnership** command from both clusters. For example, to establish a partnership between clusterA and clusterB, you would first issue the **svctask mkpartnership** command from clusterA, specifying clusterB as the remote cluster. At this point the partnership is partially configured, and sometimes described as one-way. Next, you would issue the **svctask mkpartnership** command from clusterB, specifying clusterA as the remote cluster. When this completes, the partnership is fully configured for two-way communication between the clusters. See *IBM TotalStorage SAN Volume Controller: Command-Line Interface User's Guide*.

Background copy management:

You can specify the rate at which the initial background copy from the local cluster to the remote cluster is performed. The bandwidth parameter controls this rate.

Related topics:

- “Remote Copy” on page 38
- “Remote Copy relationships”

Remote Copy relationships

In most cases, the master VDisk contains the production copy of the data and is the VDisk that the application normally accesses. The auxiliary VDisk typically contains a backup copy of the data and is used in disaster recovery scenarios.

The master and auxiliary VDisks are defined when the relationship is created, and these attributes never change. However, either VDisk may operate in the primary or secondary role according to the circumstances. The primary VDisk is the one currently receiving updates from the application, analogous to source VDisk . The secondary VDisk receives a copy of any updates to the primary VDisk, because

these updates are all transmitted across the Remote Copy link. Therefore, the secondary VDisk is analogous to a continuously updated target VDisk.

Primary

Contains a valid copy of the application data, and it is accessible for application write operations.

Secondary

Might contain a valid copy of the application data, but it is not available for application write operations.

When a relationship is created, the master VDisk is assigned the role of primary VDisk and the auxiliary VDisk is assigned the role of secondary VDisk. Therefore, the initial copying direction is from master to auxiliary. When the relationship is in a consistent state, the copy direction can be reversed by issuing the **svctask switchrcrelationship** command, and specifying the auxiliary disk as the primary.

The two VDIs in a relationship must be the same size. When the two VDIs are in the same cluster, they must be in the same input/output (I/O) group.

A relationship can be added to a Remote Copy consistency group, for ease of application management (see consistency groups below).

Note: Membership of a consistency group is an attribute of the relationship, not the consistency group. Therefore, the **svctask chrcrelationship** command is used to add or remove a relationship to or from a consistency group. See *IBM TotalStorage SAN Volume Controller: Command-Line Interface User's Guide*.

Remote Copy states

When a Remote Copy relationship is created with two virtual disks in different clusters, the distinction between the connected and disconnected states is important. These states apply to both the cluster, relationships, and consistency groups.

Inconsistent (Stopped)

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

Inconsistent (Copying)

The primary VDisk is accessible for read and write I/O operations but the secondary VDisk is not accessible for either. This state is entered after a **Start** command is issued to an consistency group in the InconsistentStopped state. This state is also entered when a **Start** command is issued, with the force option, to a consistency group in the Idling or ConsistentStopped state.

Consistent (Stopped)

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

Consistent (Synchronized)

The primary VDisk is accessible for read and write I/O operations. The secondary VDisk is accessible for read-only I/O operations.

Idling A master VDisk and a auxiliary VDisk operates in the primary role. Consequently the VDisk is accessible for write I/O operations.

Idling (Disconnected)

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

Inconsistent (Disconnected)

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

Consistent (Disconnected)

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

Related topics:

- “Remote Copy” on page 38
- “Remote Copy partnerships” on page 39

Remote Copy consistency groups

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

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

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

Remote Copy consistency group states:**Inconsistent (Stopped)**

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

Inconsistent (Copying)

The primary VDIsks are accessible for read and write I/O operations but the secondary VDisk are not accessible for either. This state is entered after a

Start command is issued to an consistency group in the InconsistentStopped state. This state is also entered when a **Start** command is issued, with the force option, to a consistency group in the Idling or ConsistentStopped state.

Consistent (Stopped)

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

Consistent (Synchronized)

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

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

Idling (Disconnected)

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

Inconsistent (Disconnected)

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

Consistent (Disconnected)

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

Empty

The consistency group contains no relationships.

Related topics:

- “Remote Copy relationships” on page 39
- “Remote Copy” on page 38
- “Virtual disks” on page 22

Chapter 4. Configuration rules and requirements

Configuration rules

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

A SAN configuration that contains SAN Volume Controller Storage Software for Cisco MDS 9000s is valid if it observes *all* of the following rules. These rules are discussed in the following section.

Storage subsystems

All SAN Volume Controller Storage Software for Cisco MDS 9000 nodes of a cluster must be able to see the same set of storage subsystem ports on each device. Any operation that is in this mode in which two nodes do not see the same set of ports on the same device is degraded, and the system logs errors that request a repair action. This rule can have important effects on storage subsystem such as FASTT, which has exclusion rules that determine to which host bus adapter (HBA) WWNNs a storage array can be mapped.

A configuration in which a SAN Volume Controller Storage Software for Cisco MDS 9000 bridges a separate host device and a RAID array is supported. Typical compatibility matrixes are shown in a document titled *Supported Hardware List* on the following Web page:

<http://www.ibm.com/storage/support/2062-2300>

The SAN Volume Controller Storage Software for Cisco MDS 9000 cannot share its storage subsystem with hosts because exclusive access is required to virtualize the storage. However, device sharing is supported. A device can be shared with a host under certain conditions as described in this topic.

Two SAN Volume Controller Storage Software for Cisco MDS 9000 clusters must not share the storage subsystem. That is, one device cannot present LUs to two different SAN Volume Controller Storage Software for Cisco MDS 9000 clusters. Although such a configuration is more unsupported than invalid, any operation that runs in this mode causes serious problems because the same managed disk (MDisk) can appear in two different SAN Volume Controller Storage Software for Cisco MDS 9000 clusters and can be concurrently mapped to different virtual disk (VDisks). This condition can cause data corruption.

The SAN Volume Controller Storage Software for Cisco MDS 9000 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):

When a storage subsystem is detected on the SAN, the SAN Volume Controller Storage Software for Cisco MDS 9000 attempts to recognize it using its Inquiry data. If the device is recognized as one of the explicitly supported storage models, then the SAN Volume Controller Storage Software for Cisco MDS 9000 uses error recovery programs that are potentially tailored to the known needs of the storage subsystem. If the device is not recognized, then the SAN Volume Controller Storage Software for Cisco MDS 9000 configures the device as a generic device. A generic

device may or may not function correctly when addressed by a SAN Volume Controller Storage Software for Cisco MDS 9000. In any event, the SAN Volume Controller Storage Software for Cisco MDS 9000 does not regard accessing a generic device as an error condition and, consequently, does not log an error. MDisks presented by generic devices are not eligible to be used as quorum disks.

Split device configurations:

In a split device configuration, a RAID array presents LUs to both a SAN Volume Controller Storage Software for Cisco MDS 9000 (which treats the LU as an MDisk) and to another host. The SAN Volume Controller Storage Software for Cisco MDS 9000 presents VDisks created from the MDisk to another host. There is no requirement for the pathing driver in the two hosts to be the same (although, if the RAID controller were an ESS, both hosts would use SDD). In Figure 11, the RAID controller is a FASTT, with RDAC used for pathing on the directly attached host, and SDD used on the host that is attached with the SAN Volume Controller Storage Software for Cisco MDS 9000. Hosts can simultaneously access LUs, which are provided by the SAN Volume Controller and directly by the device.




Figure 11. Disk controller system shared between SAN Volume Controller Storage Software for Cisco MDS 9000 and a host

In the case where the RAID controller is an ESS, the pathing driver in the host would be IBM Subsystem Device Driver (SDD) for the ESS and SDD for the SAN Volume Controller Storage Software for Cisco MDS 9000 LUs. Figure 12 on page 45

45 a supported configuration since the same pathing driver is used for both direct and virtual disks.




Figure 12. ESS LUs accessed directly and via SAN Volume Controller Storage Software for Cisco MDS 9000

Host bus adapters

The SAN Volume Controller Storage Software for Cisco MDS 9000 must be configured to export virtual disks 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/2062-2300>

Operation with other HBAs is not supported.

The number of paths from the SAN Volume Controller Storage Software for Cisco MDS 9000 nodes to a host must not exceed two. SAN Volume Controller Storage Software for Cisco MDS 9000 nodes present a minimum of three external ports to the SAN. Additional worldwide port names are created when a port is put into multiple VSANs. To restrict the number of paths to two, you should place the target port in a zone or VSAN with the hosts, the initiator port in a zone or VSAN with the disk controller systems, and the management ports together in a zone or VSAN that contains only the nodes. In this way, the hosts, storage, and inter node traffic are isolated from one another.

Nodes

The SAN Volume Controller Storage Software for Cisco MDS 9000 nodes must always be deployed in pairs. If a node fails or is removed from the configuration, the remaining node operates in a degraded mode, but the configuration is still valid.

A cluster contains 2 nodes. Each pair of nodes is known as an I/O group. The nodes in an I/O group must be on different Cisco MDS 9000 Caching Services Modules. Each node may be in *only* one I/O group. A cluster can be composed of nodes that reside on different switches.

Fibre-channel switches

The SAN must contain only supported switches. The SAN Volume Controller Storage Software for Cisco MDS 9000 supports the Cisco MDS 9000 switch only.

See the following Web site for specific firmware levels and the latest supported hardware: <http://www.ibm.com/storage/support/2145/>

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

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

Operation with fibre-channel switch extenders is not supported.

The switch configuration of a SAN Volume Controller Storage Software for Cisco MDS 9000 SAN must observe the switch manufacturer's configuration rules. These rules might put restrictions on the switch configuration; for example, the switch manufacturer might not permit other manufacturer's switches to be in the SAN. Any operations that run outside the manufacturer's rules is not supported.

The switch must be configured so that the SAN Volume Controller Storage Software for Cisco MDS 9000 nodes can see the back-end storage and the front-end HBAs. However, the front-end HBAs and the back-end storage must not be in the same zone or VSAN. Any operation that runs outside these zoning rules is not supported.

Each SAN Volume Controller for Cisco MDS 9000 node has three ports, one port for communication with hosts (the target port), one port for communication with back-end storage (the initiator port) and one port for communication between nodes (the management port). The switch must be configured so that the proper port can communicate with the hosts, back-end storage, and other nodes.

With Remote Copy, 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.

Fibre-channel switches and Inter-Switch Links:

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

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

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

Configuration requirements

A valid SAN configuration meets the following criteria:

- The two nodes on the same Cisco MDS 9000 Caching Services Module cannot be placed in the same I/O group. This means that the minimum SAN Volume Controller Storage Software for Cisco MDS 9000 configuration requires two Cisco MDS 9000 Caching Services Modules, either in the same switch or different switches.
- The nodes on a Cisco MDS 9000 Caching Services Module can belong to different clusters.
- A switch can contain nodes from more than one cluster.
- A cluster can span switches. A cluster spanning more than two switches is not supported.
- The two nodes in an I/O group can be in different switches.
- If a cluster spans switches, the switches must all be in the same subnet. To be in the same subnet, the first three numbers separated by periods of the IP addresses must be the same, and all four numbers separated by periods must be the same for the Subnet Mask. For example, if your subnet mask is 255.255.255.0, the first three numbers of your IP addresses must be the same.
- If a cluster spans switches, you can synchronize the time across the switches so that the log entry timestamps have the same basis, although this is not a requirement.
- If a cluster spans switches, you can define the same users and passwords on all the switches that the cluster has nodes in or use RADIUS to maintain consistency across all switches.
- The maximum number of virtual storage area networks (VSANs) in which a management port can be placed is four.
- The maximum number of VSANs in which the three ports of a SAN Volume Controller Storage Software for Cisco MDS 9000 node can be placed is 64. (That is, the number of VSANs the target port is in, plus the number of VSANs that the initiator port is in, plus the number of VSANs that the management port is in cannot exceed 64.)

Steps:

Perform the following steps:

1. Ensure that the Cisco MDS 9000 switch is installed.
2. Your IBM service representative must have installed the SAN Volume Controller Storage Software for Cisco MDS 9000.
3. Install and configure your disk controller systems and create the RAID resources that you intend to virtualize. To prevent loss of data, virtualize only those RAID that provide some kind of redundancy, that is, RAID 1, RAID 10, RAID 0+1, or RAID 5. Do *not* use RAID 0 because a single physical disk failure might cause the failure of many virtual disks. RAID 0, like other types of RAID offers cost-effective performance by using available capacity through data striping. However, RAID 0 does not provide a parity disk drive for redundancy (RAID 5) or mirroring (RAID 10).

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

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

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

Most back-end disk controller systems enable RAID to be divided up into more than one SCSI logical unit (LU). When configuring new storage for use with the SAN Volume Controller Storage Software for Cisco MDS 9000, you do not need to divide up the array. New storage should be presented as one SCSI LU. This will give a one-to-one relationship between MDisks and RAID.

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

4. Configure your Cisco MDS 9000s to create the zones and VSANs that the SAN Volume Controller Storage Software for Cisco MDS 9000 needs. One zone must contain all the disk controller systems and the SAN Volume Controller Storage Software for Cisco MDS 9000 nodes. For hosts, use switch zoning to ensure that each host fibre-channel port is zoned to exactly one fibre-channel port of each SAN Volume Controller Storage Software for Cisco MDS 9000 node in the cluster.
5. If you want the SAN Volume Controller Storage Software for Cisco MDS 9000 to export redundant paths to VDIs, you must install the Subsystem Device Driver (SDD) on all of the hosts that are connected to the SAN Volume Controller Storage Software for Cisco MDS 9000. Otherwise, you will not be able to use the redundancy inherent in the configuration. Install the SDD from the following Web site:

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

Be sure to install version 1.5.x.x or higher.

6. Install and configure the SAN Volume Controller Storage Software for Cisco MDS 9000 console workstation. The communication between the console workstation and the SAN Volume Controller Storage Software for Cisco MDS

9000 runs under a client-server network application called Secure Shell (SSH). SSH Server software and the SSH Client software called PuTTY are shipped with SAN Volume Controller Storage Software for Cisco MDS 9000. You will need to configure SSH and PuTTY on the console workstation. Once you have installed your console workstation, you can configure and administer the SAN Volume Controller Storage Software for Cisco MDS 9000 using a graphical interface or a command-line interface.

- You can configure the SAN Volume Controller Storage Software for Cisco MDS 9000 using the Console for SAN Volume Controller for Cisco MDS 9000, Web-based application on the console workstation.

Note: You must enable SSH on the switch if you use the console workstation.

- You can configure the SAN Volume Controller Storage Software for Cisco MDS 9000 using the command-line interface (CLI) commands.
- You can install an SSH client if you only want to use the CLI commands. If you want to use the CLI from a host other than the console workstation, ensure that the host has an SSH client installed on it.

Notes:

- a. AIX[®] comes with an installed SSH client.
- b. Linux comes with an installed SSH client.
- c. PuTTY is recommended for Windows[®].

Result:

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

1. Create the virtual storage area networks (VSANs) and zones that you will use in the switch.
2. Create the host, management, and backend storage interfaces for each node.
3. Create the cluster.
4. Add nodes to the cluster and set up the cluster properties.
5. Create managed disk groups from the managed disks to make pools of storage from which you can create virtual disks.
6. Create host objects from the HBA fibre-channel ports to which you can map virtual disks.
7. Create virtual disks from the capacity that is available in your managed disk groups.
8. Map the virtual disks to the host objects to make the disks available to the hosts, as required.
9. Optionally, create Copy Services (FlashCopy and Remote Copy) objects as required.

Related topics:

- “Managed disk groups” on page 20
- “Creating virtual disks” on page 81
- “Fibre-channel switches” on page 46

Part 2. Preparing to configure the SAN Volume Controller Storage Software for Cisco MDS 9000

Related topics:

- Chapter 5, “Centralized management functions,” on page 53

Chapter 5. Centralized management functions

The centralized management functions available are:

- Browser support for:
 - Console for SAN Volume Controller for Cisco MDS 9000
 - Fibre channel switch
- CLI configuration support using Secure Shell (SSH)
- Remote Service capability through VPN

Secure Shell (SSH)

Overview:

SSH is a client-server network application intended to ensure secure communications. The SAN Volume Controller Storage Software for Cisco MDS 9000 uses the SSH facilities within the MDS 9000 to provide these services.

When using the console workstation, you must enable SSH. To enable SSH on the SAN Volume Controller Storage Software for Cisco MDS 9000, at the switch, use the **switch#config-t** and **switch (config)#ssh server enable** commands.

Secure Shell (SSH) is an optional communication vehicle between the host system you are using and either:

- the SAN Volume Controller Storage Software for Cisco MDS 9000 command-line interface (CLI)
- or the system on which the Console for SAN Volume Controller for Cisco MDS 9000 is installed.

Authenticating SSH logins:

When you are using AIX hosts, SSH logins are authenticated on the cluster using the RSA-based authentication that is supported in the OpenSSH client available for AIX. This scheme is based on public-key cryptography, using a scheme known commonly as RSA.

Note: The authentication process for non-AIX hosts systems is similar.

With this scheme (as in similar OpenSSH systems on other host types) the encryption and decryption is done using separate keys. This means it is not possible to derive the decryption key from the encryption key.

Because physical possession of the private key allows access to the cluster, the private key must be kept in a protected place, such as the .ssh directory on the AIX host, with restricted access permissions.

Running the command-line interface (CLI):

When establishing a session with the SAN Volume Controller Storage Software for Cisco MDS 9000 using the CLI, any of the SSH facilities provided by the SAN Volume Controller Storage Software for Cisco MDS 9000 may be used, provided that you have a compatible SSH client.

We recommend that you install the SSH Client in order to use the command-line interface (CLI). However, it is not required.

The console workstation requires the SSH client software called PuTTY. This software provides the Secure Shell (SSH) client function for users logged into the console workstation who wish to invoke the SAN Volume Controller Storage Software for Cisco MDS 9000 command-line interface (CLI). You must install the PuTTY client software.

If you wish to run the SAN Volume Controller Storage Software for Cisco MDS 9000 command-line interface (CLI) from a different system than the console workstation, you must install an SSH client. For your convenience, the installation program to install the PuTTY software on Windows can be found in the SSH client directory of the Console for SAN Volume Controller for Cisco MDS 9000 CD-ROM.

The console workstation also has the Console for SAN Volume Controller for Cisco MDS 9000 Web server and Common Information Model (CIM) Object Man

Configuring the Secure Shell (SSH) client system

IBM has provided the PuTTY Secure Shell client software on the console workstation. You can also install PuTTY on any Windows 2000 server where you will run the command-line interface (CLI) or where you install the Console for SAN Volume Controller for Cisco MDS 9000. If you have some other Secure Shell client software to run on another host, follow that software documentation to perform the tasks equivalent to the following steps.

1. Install SSH client software.
2. Configure the PuTTY session, if required, on the SSH client system.
3. Copy the SSH public key to the console workstation.

Related topics:

- “Configuring the PuTTY session for the command-line interface”

Configuring the PuTTY session for the command-line interface

This step is only required if you are preparing to run the CLI from the console workstation.

Steps:

Perform these steps to configure the PuTTY session on the SSH client system:

1. Click **Start -> Programs -> PuTTY -> PuTTY** to open the PuTTY Configuration interface window. The items you select in the Category pane on the left side of the window affect the content in the right pane of the window.
2. In the Category pane, click **Session**.
3. Click **SSH**.
4. In the Connection tree, click **Connection -> SSH**. This will bring up a different view in the right pane.
5. Ensure that the button labeled **2** is selected.
6. In the SSH tree, click **Auth**. A different view opens in the right pane.
7. In the Category pane, click **Session**.
8. In the Load, save or delete a stored session section in the right pane, click **Default Settings -> Save** in the **save or delete a stored session** field.

Upgrading software on the console workstation

Vendor software:

The following table provides information about upgrading your vendor software.

Table 10. Upgrading vendor software

Software	Reasons to upgrade
Microsoft Windows 2000 Server Edition and Service Pack 3 or Service Pack 4	Requires an upgrade only if new functions are required.
Windows 2000 Security Patches	Critical updates should be installed as they become available in order to avoid security exposures. Visit the Web site: http://windowsupdate.microsoft.com
PuTTY	Required an upgrade only if a problem is found or a new function is required.
Adobe Acrobat Reader	Requires an upgrade only if a problem is found or a new function is required.

IBM software:

The following table provides information about upgrading your IBM software.

Table 11. Reasons to upgrade IBM software

Software	Reasons to upgrade
Console for SAN Volume Controller for Cisco MDS 9000	Requires an upgrade only if a problem is found or a new function is required.

The console workstation is provided on CDs. The console workstation installation instructions for the software packages are located in the individual software installation guides.

It is the user's responsibility to download and install software upgrades for the console workstation. For recommended IBM upgrades, see the following Web site:

www.ibm.com/storage/support/2062-2300

Remote Service

Remote service is handled through remote support facilities of the Cisco MDS 9000. Refer to Cisco MDS 9000 publications for information about Cisco remote support facilities.

Related topics:

- "Related publications" on page xi

Troubleshooting console workstation problems

- Recovering from SAN Volume Controller Console signing off
- Resolving Windows 2000 boot problem

Related topics:

- “Recovering from Console for SAN Volume Controller for Cisco MDS 9000 signing off”

Recovering from Console for SAN Volume Controller for Cisco MDS 9000 signing off

Problem:

You are presented with a dialogue box containing the words: You have signed off. This window will be closed. Before checking for hardware errors, open a new browser window and try to reconnect to the Console for SAN Volume Controller for Cisco MDS 9000. The sign off message is generally caused by the open browser session timing out. If you have left the browser window open from a previous session, this is the likely cause. You should be able to reconnect. If you cannot reconnect, perform the following steps.

Investigation steps:

Try the following actions to resolve the problem:

The problem could be due to:

- **A memory failure in the console workstation and it is running with less than the required one gigabyte of memory.**
Check and correct the memory problem.
- **The IP address of the console workstation changing since the last reboot.**

Restart the console workstation to correct this problem.

Installing antivirus software

In order to protect your workstations and your enterprise, install the latest patches for your Windows operating system and the latest levels of your antivirus software on the console workstation.

If the console workstation is either connected to your local area network or you have enabled remote support, you must install the latest levels of the antivirus software that your enterprise uses.

Related topics:

- Chapter 5, “Centralized management functions,” on page 53

Part 3. SAN Volume Controller Console

Chapter 6. Console for SAN Volume Controller for Cisco MDS 9000

Overview:

The SAN Volume Controller Storage Software for Cisco MDS 9000 is provided with a console, which is Web-browser based. It can be used to create and maintain the configuration of storage associated with the SAN Volume Controller Storage Software for Cisco MDS 9000. It also provides user management and access to multiple clusters.

The functions that can be performed with the Console for SAN Volume Controller for Cisco MDS 9000:

- Initial setup of the cluster, its nodes, and the I/O groups (or node pairs). This function includes diagnostics and error log analysis of the cluster.
- Setup and maintenance of managed disks and managed disk groups.
- Enabling SSH keys.
- Setup and maintenance of virtual disks.
- Setup of logical host objects.
- Mapping of virtual disks to hosts.
- Navigation from managed hosts to virtual disk and to managed disk groups, and the reverse direction up the chain.
- Set up and start of Copy Services:
 - FlashCopy and FlashCopy Consistency groups
 - Synchronous Remote Copy and Remote Copy Consistency groups

The Console for SAN Volume Controller for Cisco MDS 9000 is Storage Management Initiative Specification (SMI-S) compliant.

Accessing the Console for SAN Volume Controller for Cisco MDS 9000

The Console for SAN Volume Controller for Cisco MDS 9000 is a Web-based application that you can use to manage multiple clusters. Because the application is Web-based, do not set the browser to disable popup windows. This can prevent the windows in the Console for SAN Volume Controller for Cisco MDS 9000 from opening.

You access the Console for SAN Volume Controller for Cisco MDS 9000 by pointing a Web browser at the following URL on your console workstation:

```
http://<svccconsoleip>:9080/ica
```

where *<svccconsoleip>* is the IP address of your console workstation.

Note: If you are accessing the Console for SAN Volume Controller for Cisco MDS 9000 from a workstation that is locally installed, you can point your browser to the following URL:

```
http://localhost:9080/ica
```

Log on to the Console for SAN Volume Controller for Cisco MDS 9000 using the superuser user name, which is superuser, and the superuser password, which is passw0rd. (Upon first access, you will be required to change the superuser password.)

Use the Console for SAN Volume Controller for Cisco MDS 9000 panels to configure SAN Volume Controller Storage Software for Cisco MDS 9000 clusters in your environment. Once the cluster has been configured, you can use the View Clusters panel to launch another browser window with specific information for a specific cluster.

Console for SAN Volume Controller for Cisco MDS 9000 layout

The basic frame layout consists of a banner, task bar, portfolio and a work area. An optional frame can be added for embedded task assistance or help.

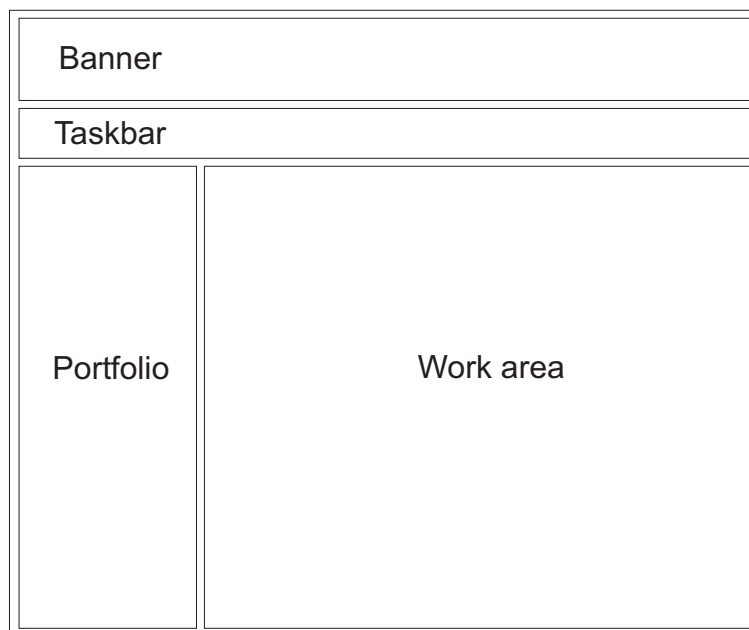


Figure 13. Basic frame layout

Console for SAN Volume Controller for Cisco MDS 9000 banner area

This area is used for product or customer identification.

Console for SAN Volume Controller for Cisco MDS 9000 task bar

The task bar keeps track of all opened primary tasks and allows the user to quickly go back to the previous task or move forward to the next task. Clicking the **question mark (?)** icon on the right side brings up the information center in a separate browser window; clicking the (I) icon brings up a help topic for the panel that is currently displayed in the work area.

Console for SAN Volume Controller for Cisco MDS 9000 portfolio

The portfolio area contains task based links that open panels in the work area. Common tasks are grouped under task headings and are expandable and collapsible.

Console for SAN Volume Controller for Cisco MDS 9000 work area

The work area is where you work with your cluster and the objects it contains. The work area is the main area of the application.

Chapter 7. Overview of creating a cluster using the Console for SAN Volume Controller for Cisco MDS 9000

Overview:

The create cluster wizard of the Console for SAN Volume Controller for Cisco MDS 9000 enables you to create a cluster through its console.

Browser requirements for the Console for SAN Volume Controller for Cisco MDS 9000

The following Web browser versions are required.

- Windows and UNIX operating systems
 - Netscape version 7.1
 - Netscape is available from the following Web site:

<http://wp.netscape.com/download/archive.html>
 - Internet Explorer Version 6 or later
 - You can get version 6 or later from the following Web site:

<http://www.microsoft.com/windows/ie/downloads/ie6/default.asp>
- AIX operating system
 - You can get AIX Netscape version 7.0 from the following Web site:

<http://devedge.netscape.com/central/gecko/2002/download/>

You must ensure that the proxy setting is disabled.

- For Netscape, perform the following steps:
 1. Open your Netscape browser and click **Edit -> Preferences**. The Preferences window is displayed.
 2. From the left side category, click **Advanced** to expand the sub options. The sub option Proxies is displayed.
 3. Click **Proxies**. The Proxies window is displayed.
 4. Select **Direct connection to Internet**.
- For Internet Explorer, perform the following steps:
 1. Click **Tools -> Internet Options -> Connections -> LAN Settings**.
 2. Click to clear the **Use a proxy server** box.

Related topics:

- “Configuring a cluster using the Console for SAN Volume Controller for Cisco MDS 9000” on page 67

Preparing the Cisco MDS switch

To prepare for both hardware and software setup, be sure you have the following in place before getting started:

- Installed and functioning Cisco MDS 9000 Family switch fabric. Each Cisco MDS 9000 Family switch needs a switch IP (mgmt0, one for each switch).

- One subnet for fabric-wide communications: This subnet may be private (only seen by the MDS switches and the IBM TotalStorage server) and should have a range large enough for all supervisor modules in the IP fabric.
- One IP address for each cluster.

Refer to the *Cisco MDS 9216 Switch Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide*, and the *Cisco MDS 9000 Family Configuration Guide*.

Setting up the Cisco MDS switch

This topic explains the process to set up a new, or update an existing MDS 9000 Family switch in preparation for the SAN Volume Controller Storage Software for Cisco MDS 9000. Before setting up your hardware, ensure you have correctly installed the Cisco MDS chassis and its components as specified in the *Cisco MDS 9216 Switch Hardware Installation Guide, Chapter 2: Installing the Cisco MDS 9000 Family Switch* or the *Cisco MDS 9500 Series Hardware Installation Guide, Chapter 2: Installing the Cisco MDS 9000 Family Switch*.

Be sure to save your work frequently using the Cisco MDS 9000 command:

```
copy running-config startup-config
```

Setting up a new Cisco MDS switch:

To set up a new Cisco MDS 9000 Family switch, follow these steps.

1. Follow instructions for preinstallation, installing the chassis in the rack, grounding the chassis, installing modules, installing CompactFlash cards, installing Power Supplies, and installing the Fan Assembly as specified in the *Cisco MDS 9216 Switch Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide*.
2. Connect to the supervisor module as specified in the *Cisco MDS 9216 Switch Hardware Installation Guide, Chapter 3: Connecting the Cisco MDS 9000 Family Switch* or the *Cisco MDS 9500 Series Hardware Installation Guide, Chapter 3: Connecting the Cisco MDS 9000 Family Switch*.
3. Login to the Cisco MDS 9000 Family switch using the Cisco MDS 9000 Family Command-Line Interface (CLI).
4. Configure the switch as specified in the *Cisco MDS 9000 Family Configuration Guide, Chapter 3: Initial Configuration*.
 - a. Perform the initial setup routine.
 - b. Assign a switch name.
 - c. Access the switch.
 - d. Create at least one user that has the svc-admin role. Users with the svc-admin role have all privileges to administer the cluster and storage. They do not have any access to configure the switch or the SAN. The default Cisco "admin" user ID includes svc-admin role privileges.
 - e. Prior to using the Console for SAN Volume Controller for Cisco MDS 9000 with the switch, enable SSH on the switch.

After reviewing the default configuration, you can change it or perform other configuration or management tasks. The initial setup can only be performed at the Cisco MDS 9000 command line interface. You can continue to configure other software features, or access the switch after initial configuration by using either the Cisco MDS 9000 command line interface, the IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000 command line

interface, or the Console for SAN Volume Controller Storage Software for Cisco MDS 9000 graphical user interface (GUI). The rest of this procedure uses the Cisco CLI to configure the switch.

5. Verify the module status as specified in the *Cisco MDS 9000 Family Configuration Guide, Chapter 3: Initial Configuration*.
6. Configure the management port as specified in the *Cisco MDS 9000 Family Configuration Guide, Chapter 3: Initial Configuration*.

Updating an existing Cisco MDS switch:

The following steps assume you have a remote FTP, TFTP, SFTP, or SCP server that contains the switch image. Be sure to save your work frequently using the Cisco CLI command:

```
copy running-config startup-config
```

To update an existing Cisco MDS switch, follow these steps.

1. Connect to the supervisor module as specified in the *Cisco MDS 9216 Switch Hardware Installation Guide, Chapter 3: Connecting the Cisco MDS 9000 Family Switch* or the *Cisco MDS 9500 Series Hardware Installation Guide, Chapter 3: Connecting the Cisco MDS 9000 Family Switch*.

2. Login to the Cisco MDS switch using the Cisco MDS 9000 Family CLI.

3. Copy the SAN Volume Controller Storage Software for Cisco MDS 9000 software image file from the remote directory to the bootflash directory in the active supervisor module, as shown in the following example.

```
switch# copy ftp://10.1.7.2/m9000-ckg-csm-svc-mz.1.3.1.bin  
bootflash:m9000-ckg-csm-svc-mz.1.3.1.bin
```

In this example, *m9000-ckg-csm-svc-mz.1.3.1.bin* represents the file name of the software image. Obtain the current file name and FTP site from your IBM or Cisco representative.

4. The image file can be locally accessible (bootflash/slot0) or on an external machine on the network (scp, sftp, ftp, tftp). This example copies the file from the remote directory using FTP and then installs it from the bootflash directory.

5. Verify that the file was copied in the required directory, as shown in the following example.

```
switch# dir bootflash: admin 524288 Sep 22 04:27:42 2003  
m9000-ckg-csm-svc-mz.1.3.1.bin
```

In this example, *m9000-ckg-csm-svc-mz.1.3.1.bin* represents the file name of the software image. Obtain the current file name and FTP site from your IBM or Cisco representative.

6. Use the show version image command to verify that you have successfully downloaded the correct version of the SAN Volume Controller Storage Software for Cisco MDS 9000 image, as shown in the following example.

```
switch# show version image bootflash:m9000-ckg-csm-svc-mz.1.3.1.bin<  
image name: m9500-sflek9-svc-mzg.1.0.3.bin  
diagnostic: version 1.3(1) [gdb]  
compiled: 9/22/2003 12:00:00
```

In this example, *m9000-ckg-csm-svc-mz.1.3.1.bin* represents the file name of the software image. Obtain the current file name and FTP site from your IBM or Cisco representative.

7. Issue the show install command to verify if the downloaded SAN Volume Controller Storage Software for Cisco MDS 9000 image is compatible with the existing software images (system image, kickstart image), as shown in the following example. In this example, *m9000-ckg-csm-svc-mz.1.3.1.bin*

represents the file name of the software image. Obtain the current file name and FTP site from your IBM or Cisco representative.

```
switch# show install bootflash:m9000-ckg-csm-svc-mz.1.3.1.bin
```

In this example, *m9000-ckg-csm-svc-mz.1.3.1.bin* represents the file name of the software image. Obtain the current file name and FTP site from your IBM or Cisco representative. Refer to the www.cisco.com web site to verify compatibility issues or the appropriate Cisco MDS 9000 Family release notes to ensure that your system and setup meets the minimum requirements, or images will not install properly.

8. Change to configuration mode.

```
switch# config terminal
switch(config)#
```

9. Insert the Cisco MDS 9000 Caching Services Module (CSM) as specified in *Cisco MDS 9216 Switch Hardware Installation Guide, Chapter 2: Installing the Cisco MDS 9000 Family Switch* or the *Cisco MDS 9500 Series Hardware Installation Guide, Chapter 2: Installing the Cisco MDS 9000 Family Switch*. The updated SAN Volume Controller Storage Software for Cisco MDS 9000 software automatically detects the presence of the newly-inserted CSM.

10. Install the new image on each SAN Volume Controller Storage Software for Cisco MDS 9000 node in each module in the fabric, as shown in the following example. Each module has two nodes.

```
switch# install module 2 node 1 image svc-system
bootflash:m9000-ckg-csm-svc-mz.1.3.1.bin
SVC reimage going on. Please wait
m9000-ckg-csm-svc-mz.1.3.1.bin 100% |*****|
45408 KB 00:53
svc 2/1 software reimage succeeded
```

```
switch# install module 2 node 2 image svc-system
bootflash:m9000-ckg-csm-svc-mz.1.3.1.bin
SVC reimage going on. Please wait
m9000-ckg-csm-svc-mz.1.3.1.bin 100% |*****|
45408 KB 00:55 svc 2/2 software reimage succeeded
```

In this example, *m9000-ckg-csm-svc-mz.1.3.1.bin* represents the file name of the software image. Obtain the current file name and FTP site from your IBM or Cisco representative.

All prior information on the upgraded node is lost with each upgrade.

11. Verify that the CSM booted successfully, by issuing the command:

```
show module
```

The status will progress from testing to ok when it has fully booted.

Mod	Ports	Module-Type	Model	Status
2	0	Caching Services Module	DS-X9560-SMAP ok	
4	0	Caching Services Module		powered-dn
5	0	Supervisor/Fabric-1	DS-X9530-SF1-K9	active *
7	0	Caching Services Module	DS-X9560-SMAP	ok
8	32	1/2 Gbps FC Module	DS-X9032	ok

Mod	Sw	Hw	World-Wide-Name(s) (WWN)
2	1.3(0.58m)	0.702	--
5	1.3(0.58m)	0.602	--
7	1.3(0.58m)	0.602	--
8	1.3(0.58m)	0.3	21:c1:00:05:30:00:00:00 to 21:e0:00:05:30:00:00:00

Mod	MAC-Address(es)	Serial-Num
2	00-05-30-01-37-7a to 00-05-30-01-37-fe	JAB072705JA

```
5 00-05-30-00-64-be to 00-05-30-00-64-c2
7 00-05-30-00-b1-2e to 00-05-30-00-b1-b2 JAB071007H9
8 00-05-30-00-2d-e2 to 00-05-30-00-2d-e6 JAB06280ae9
```

* this terminal session

You might also need to update the Supervisor code to support the Cisco MDS 9000 Caching Services Module (CSM). See the *Cisco MDS 9000 Family Configuration Guide* for more information.

Creating SAN Volume Controller Storage Software for Cisco MDS 9000 interfaces

See the *Cisco MDS 9000 Family Configuration Guide* for more information about VSANs.

Steps:

Each node in a CSM is represented by an interface (interface svc 2/1 and interface svc 2/2). To create a SAN Volume Controller Storage Software for Cisco MDS 9000 interface and N-port VSANs in a Cisco MDS switch using the Cisco MDS 9000 commands, follow these steps:

1. Enter configuration mode by issuing the command:

```
switch# config t
switch(config)#
```

2. Create the SAN Volume Controller Storage Software for Cisco MDS 9000 interface for Node in the CSM module that resides in slot 2 by issuing the command:

```
switch(config)# interface svc 2/1
switch(config-if)#
```

3. Enable the SAN Volume Controller Storage Software for Cisco MDS 9000 interface by issuing the command:

```
switch(config-if)# no shutdown
```

Configuring a cluster using the Console for SAN Volume Controller for Cisco MDS 9000

Steps:

Perform the following steps to create a cluster using the Create a Cluster wizard:

1. Start the Console for SAN Volume Controller for Cisco MDS 9000 by clicking on the desktop icon or by pointing your Web browser to <http://localhost:9080/ica> . The Sign on panel is displayed.
2. Type superuser for the user ID and password for the password. The first time you sign on as the superuser, you must change the password for the superuser. After you change the password, the Welcome panel is displayed.
3. Click **Clusters** from the navigation panel on the left.
4. If this is the first time that you have accessed the Console for SAN Volume Controller for Cisco MDS 9000 after installing it, you will see two buttons in the work area off the Welcome panel: Add SAN Volume Controller Cluster, and Create New SAN Volume Controller Cluster.

If the cluster has not been created with the CLI, click the **Create New SAN Volume Controller Cluster** button and follow the instructions in the Creating a Cluster wizard.

If the cluster has already been created with the CLI, click the **Add SAN Volume Controller Cluster** button and follow the instructions in the Adding a Cluster wizard.

Note: If one or more clusters have been added using the Console for SAN Volume Controller for Cisco MDS 9000 GUI, the Add SAN Volume Controller Cluster, and Create New SAN Volume Controller Cluster buttons will not be displayed in the work area of the Welcome panel. You can create or add more clusters by clicking on **Clusters** from the Portfolio. A Viewing Panels screen then displays in the work area, where you can select the option you want from the drop-down list, and you will be directed to the appropriate wizard to add or create clusters.

5. The Create a Cluster wizard is displayed.

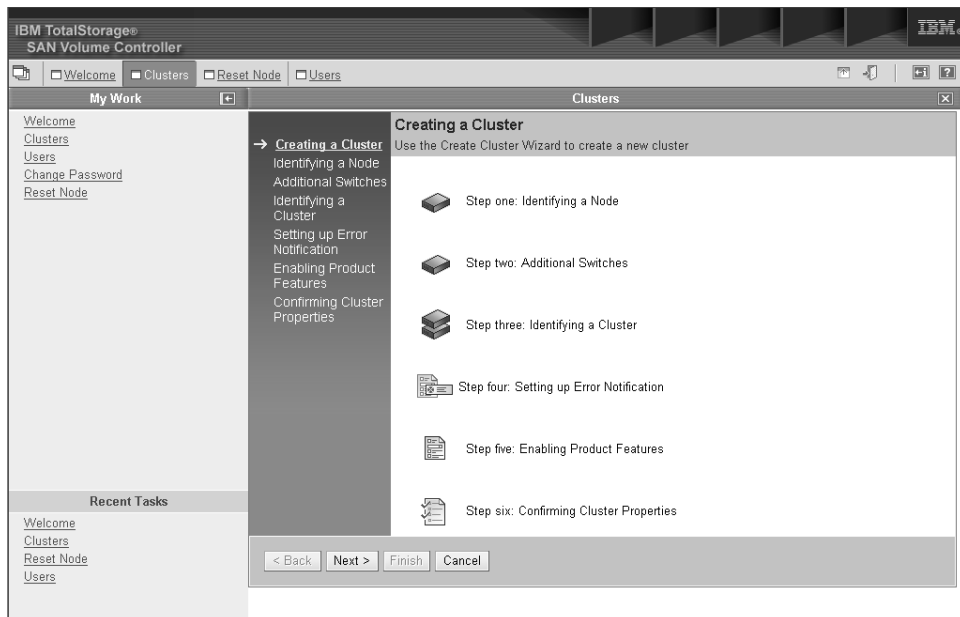


Figure 14. Create a Cluster Wizard: Introduction panel

6. Click **Next**. The Identifying a Node panel of the wizard is displayed.

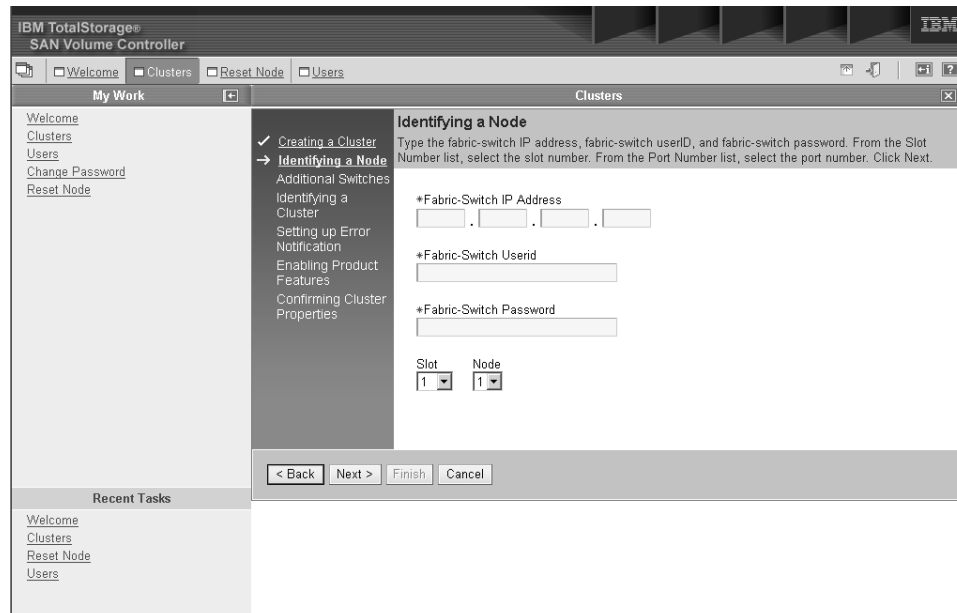


Figure 15. Create a Cluster Wizard: Identifying a Node panel

7. Type the fabric-switch IP address, fabric-switch user ID, and fabric-switch password. Select the slot number and node number. Click **Next**. The Identifying a Cluster panel of the wizard is displayed.
8. Click **Additional Switches** in the portfolio. The Additional Switches panel of the wizard is displayed. Enter the IP address of an additional switch, if desired, or click **Next** to go to the next panel.

Note: This panel is used to add nodes on another switch, if desired. If you add nodes on another switch, you need to enter the IP address of the other fabric switch.

9. Type the cluster IP address and cluster name. Click **Next** to continue. The Setting Up Error Notification panel of the wizard is displayed.
10. Select an SNMP setting and an E-mail setting. You can specify the following kinds of errors for notification:

All Notify for all occurrences, including state changes for objects.

Hardware only

Notify for all occurrences except state changes.

None Do not notify for any occurrences.

Click **Next** to continue. The Enabling Product Features panel of the wizard is displayed.

11. Select the features that you want to enable for this cluster. Set the virtualization limit to the number of gigabytes that you can virtualize as indicated in your license. Click **Next** to continue. The Confirming Cluster Properties panel of the wizard is displayed.
12. Verify that the information for the cluster is correct. Click **Finish** to create the cluster.

Result:

The cluster is created and it should be listed on the Viewing Clusters panel.

Note: You may have to press **Refresh** on the Viewing Clusters panel to see the new cluster.

Related topics:

- “Browser requirements for the Console for SAN Volume Controller for Cisco MDS 9000” on page 63
- Chapter 8, “Scenario: typical usage for the Console for SAN Volume Controller for Cisco MDS 9000,” on page 73

Launching the Console for SAN Volume Controller for Cisco MDS 9000

You can launch the SAN Volume Controller Storage Software for Cisco MDS 9000 from the Viewing Clusters panel. The SAN Volume Controller Storage Software for Cisco MDS 9000 is the centralized Web application that is used to manage your clusters.

Context:

This procedure assumes that you are at the Welcome panel for the SAN Volume Controller Storage Software for Cisco MDS 9000.

Steps:

Perform the following steps to launch the SAN Volume Controller Storage Software for Cisco MDS 9000 application:

1. Click **Clusters** from the portfolio. The Viewing Clusters panel is displayed.
2. Select the cluster that you want to manage with the application.
3. Select **Launch the SAN Volume Controller application** from the drop-down list and click **Go**. A secondary browser window opens.

Related topics:

- “Clusters” on page 12

Displaying cluster properties using the Console for SAN Volume Controller for Cisco MDS 9000

Steps:

Perform the following steps to display the cluster properties:

1. From the secondary browser window, click **Manage Cluster** from the portfolio.
2. Click **View Cluster properties** to view the properties for the cluster. The Cluster Properties notebook is displayed.
3. Click:
 - **General** tab in the notebook to display the general properties.
 - **IP Addresses** to view the IP address of the cluster .
 - **Space** to view the space and capacity within the virtual disks (VDisks) and managed disk (MDisk) groups.
 - **SNMP** to view the SNMP details.
 - **Statistics** to view the cluster statistics details.
 - **Remote Copy** to view the remote copy properties of the cluster.

Changing the password for the Console for SAN Volume Controller for Cisco MDS 9000

Prerequisites:

This procedure assumes that you are at the Welcome panel for the Console for SAN Volume Controller for Cisco MDS 9000.

Steps:

Perform the following steps to change your password:

1. Click on Change Password in the portfolio. The Change Password Screen is displayed.
2. Type the new password.
3. Type the new password again.
4. Click OK.

Related topics:

- “Displaying cluster properties using the Console for SAN Volume Controller for Cisco MDS 9000” on page 70
- “Launching the Console for SAN Volume Controller for Cisco MDS 9000” on page 70

Managing SAN Volume Controller Users

Prerequisites:

This procedure assumes that you are at the Welcome panel for the Console for SAN Volume Controller for Cisco MDS 9000.

Note: You must have superuser authority to perform this task.

Steps:

Perform the following steps to manage SAN Volume Controller Users:

1. Click on Users in the portfolio. The Users panel is displayed.
2. Select one of the following actions from the drop down list:
 - Add a superuser
 - Add a user
 - Modify a user
 - Delete a user
3. Select a user and click Go.

Related topics:

- “Changing the password for the Console for SAN Volume Controller for Cisco MDS 9000”

Chapter 8. Scenario: typical usage for the Console for SAN Volume Controller for Cisco MDS 9000

The main focus of the following example is to provide storage to your host system. Our hypothetical example is the following:

For example, you wish to provide a host system with two disks and create a FlashCopy of these two disks. The copy is to be made available to a second host. These two hosts require that the host objects that are created, correspond with the group of WWPNs presented by their fibre-channel HBAs to the SAN. You also need to create four virtual disks, one for each of the disks that are to be presented to the hosts. Once the VDisks are created, you can map two of them to each host. In order to create the VDisks you need to have a managed disk group to be able to create them from. You wish to spread the 8 managed disks across two groups and create the source VDisks from one and the target VDisks from the other. In order to create any of these objects you need to create a cluster and at least one more node to the cluster.

The following steps illustrates how this can be done:

1. Create a cluster.
2. Configure the cluster with an IP address of 9.20.123.456. Name the cluster `examplecluster`.
3. Launch the SAN Volume Controller Storage Software for Cisco MDS 9000 application for the cluster. A secondary browser window opens to the SAN Volume Controller Storage Software for Cisco MDS 9000 Web application. Now you can work with the specific SAN Volume Controller Storage Software for Cisco MDS 9000 cluster which you selected.
4. Add nodes
 - `knode` and `lnode` to the I/O group called `io_group0` in the `examplecluster` cluster
 - `mnode` and `nnode` to the I/O group called `io_group1` in the `examplecluster` cluster
5. Create the managed disk (MDisk) groups `maindiskgroup` and `bkpdiskgroup`
6. Create four virtual disks (VDisks)
 - 2 VDisks from `maindiskgroup`
 - 2 VDisks from `bkpdiskgroup`
7. Create two host objects
 - A host object called `demohost1` with HBAs that have WWPNs of `210100e08b251dd4`, and `210100e08b251dd5`
 - A host object called `demohost2` with HBAs that have WWPNs of `210100e08b251dd6`, and `210100e08b251dd7`
8. Create the VDisk-to-host mappings
 - Create a VDisk-to-host mapping for `demohost1`
 - Create a VDisk-to-host mapping for `demohost2`

Once this step is complete, you have successfully created storage on your host system.
9. Create a FlashCopy consistency group called `maintobkpcopy` and add the two FlashCopy mappings to it.

Note: You must first create FlashCopy mappings to define the relationships.

Related topics:

- “Configuring a cluster using the Console for SAN Volume Controller for Cisco MDS 9000” on page 67
- “Creating virtual disks” on page 81

Adding nodes to a cluster

Before you begin: You must create and launch a cluster. For more information, see the related topics below.

Prerequisites:

Attention: Before adding a node to a cluster, make sure that the switch zoning is configured such that the node being added is in the same zone as all other nodes in the cluster. In particular, if you are replacing a node and the switch is zoned by worldwide port name (WWPN) rather than by switch port, make sure the switch is configured such that the node being added is in the same VSAN/zone as you must update the switch configuration.

Before adding a node to the cluster check to see if any of the following conditions are true:

- The cluster has more than one I/O group.
- The node being added to the cluster uses physical node hardware or a slot which has previously been used as for a node in the cluster.
- The node being added to the cluster uses physical node hardware or a slot which has previously been used for a node in another cluster and both clusters have visibility to the same hosts and backend storage.

Attention: If any of these conditions are true, then you must perform the following special procedures. Failure to perform the special procedure is likely to result in the corruption of all data managed by the cluster.

Special procedures when adding a node to a cluster:

If any of the previous conditions are true, then the following special procedures apply. These special procedures apply when you use either the **svctask addnode** command or the Console for SAN Volume Controller for Cisco MDS 9000. When a node is added to a cluster then either:

- The node must be added back to the same I/O group that it was previously in.

Note: The WWNN of the nodes in the cluster can be determined using the command:

```
svcinfolnode
```

or, if this information is not available, then

- *Before* the node is added back into the cluster all the hosts using the cluster must be shut down. The node must then be added before the hosts are rebooted. Or, if the I/O group information is not available and it is inconvenient to shutdown and reboot all of the hosts using the cluster, then
- On all the hosts connected to the cluster, unconfigure the Fibre Channel adapter device driver, the disk device driver, and the SDD device driver, before you add the node to the cluster.

Reconfigure the Fibre Channel adapter device driver, the disk device driver, and the SDD device driver, after adding the node into the cluster.

Note: This may not be possible on all operating systems in all circumstances.

Steps:

Perform the following steps to add a node to a cluster:

1. Click **Work with Nodes** from the portfolio.
2. Click **Nodes** from the portfolio. The Viewing Nodes panel is displayed.
3. Select **Add a Node** from the list and click **Go**. The Adding nodes to a cluster panel is displayed.

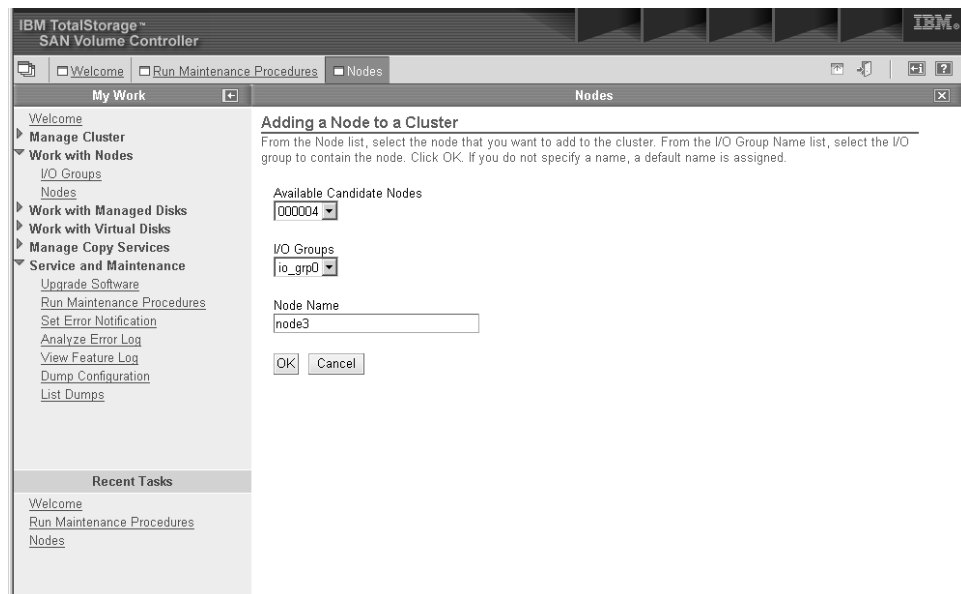


Figure 16. SAN Volume Controller Console Nodes panel

4. From the list of candidate nodes, select the number of the node that you want to add. If you get a “Candidate nodes not available” error, make sure your zoning is correct.
5. Select the I/O group for the node.

Hypothetical scenarios where the special procedures may apply.:

The following are two hypothetical scenarios where the special procedures may apply:

- Two nodes of a four-node cluster have been lost because of a power failure. In this case the two lost nodes must be added back into the cluster using the **svctask addnode** command or the Console for SAN Volume Controller for Cisco MDS 9000.
- A user decides to delete two nodes from the cluster and add them back into the cluster using the **svctask addnode** command or the Console for SAN Volume Controller for Cisco MDS 9000.

Background:

Applications on host systems direct I/O operations to filesystems or logical volumes which are mapped by the operating system to vpaths which are pseudo disk objects supported by the SDD driver. See the *IBM TotalStorage Subsystem Device Driver: User's Guide*.

The SDD driver maintains an association between a vpath and a SAN Volume Controller Storage Software for Cisco MDS 9000 VDisk. This association uses an identifier (UID) which is unique to the VDisk and is never re-used. This allows the SDD driver to unambiguously associate vpaths with VDIs.

The SDD device driver operates within a protocol stack which also contains Disk and Fibre Channel device drivers which allow it to communicate with the SAN Volume Controller Storage Software for Cisco MDS 9000 using the SCSI protocol over Fibre Channel as defined by the ANSI FCS standard. The addressing scheme provided by these SCSI and Fibre channel device drivers uses a combination of a SCSI Logical unit number (LUN) and the World Wide Name for the Fibre Channel Node and Ports.

In the event of errors occurring, error recovery procedures (ERPs) operate at various tiers in the protocol stack. Some of these ERPs cause I/O to be redriven using the same WWN and LUN numbers which were previously used.

The SDD device driver does not check the association of the VDisk with the VPath on every I/O that it performs.

Data Corruption Scenario:

Consider a four-node SAN Volume Controller Storage Software for Cisco MDS 9000 configuration.

The nodes, Node1 and Node2, are in I/O group 0 which supports the VDisk, VDisk0.

The nodes, Node3 and Node4, are in I/O group 1 which supports the VDisk, VDisk1.

Assume that VDisk 0 is mapped to a host as LUN 0. This will be LUN 0 associated with the ports in Node1 and Node2. We might represent this as N1/0 and N2/0 respectively. Assume also that VDisk1 is also mapped to the host as LUN 0. Thus N3/0 and N4/0 are mapped to VDisk1.

Now assume that nodes, Node2 and Node4, are removed from the cluster.

If Node2 is added back into the cluster into I/O Group 1 a data corruption could occur because:

- N2/0 now maps to VDisk1 whereas previously it mapped to VDisk0.
- There are scenarios where I/O intended for VDisk0 could be sent to the old address, N2/0, which now is mapped to VDisk1.

Context:

Assume that the cluster has been created.

Steps:

Perform the following steps to add nodes to the cluster:

1. From the Welcome panel, click **Work with Nodes** in the portfolio.
2. Click **Nodes** in the portfolio. The Nodes panel is displayed.

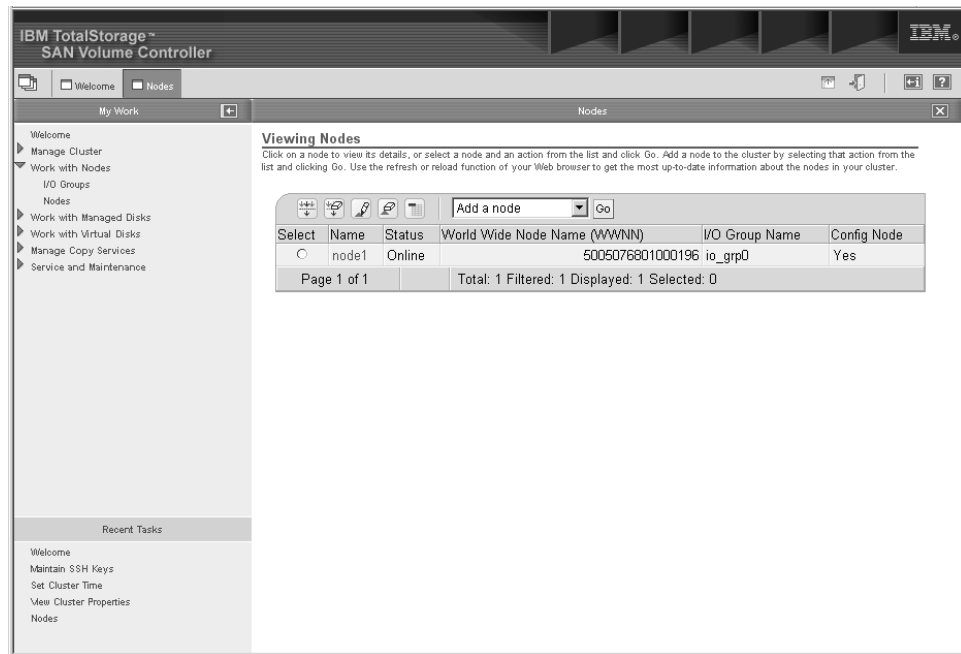


Figure 17. Nodes panel

3. Select **Add Node** from the drop down list and click **Go**.



Figure 18. Add Node drop down list

4. **Attention:** If you are re-adding a node to the SAN, ensure that you are adding the node to the same I/O group that it was removed from. Failure to do this can result in data corruption. Use the information that was recorded when the node was originally added to the cluster. If you do not have access to this information, call IBM Service to add the node back into the cluster without corrupting the data.

If you are adding the node into the cluster for the first time, record the following information:

- Card serial number
- Slot the Cisco MDS 9000 Caching Services Module is in
- Node (1 or 2) on the Cisco MDS 9000 Caching Services Module
- WWNN
- All WWPNs
- I/O group that contains the node

This can avoid a possible data corruption exposure if the node must be removed from and re-added to the cluster.

Note: This warning also is displayed on the Console for SAN Volume Controller for Cisco MDS 9000 panel when adding the node.

Select the node that you want to add from the list and the I/O group name that you want to add the node to and click **OK**. This will add the node to the I/O group.

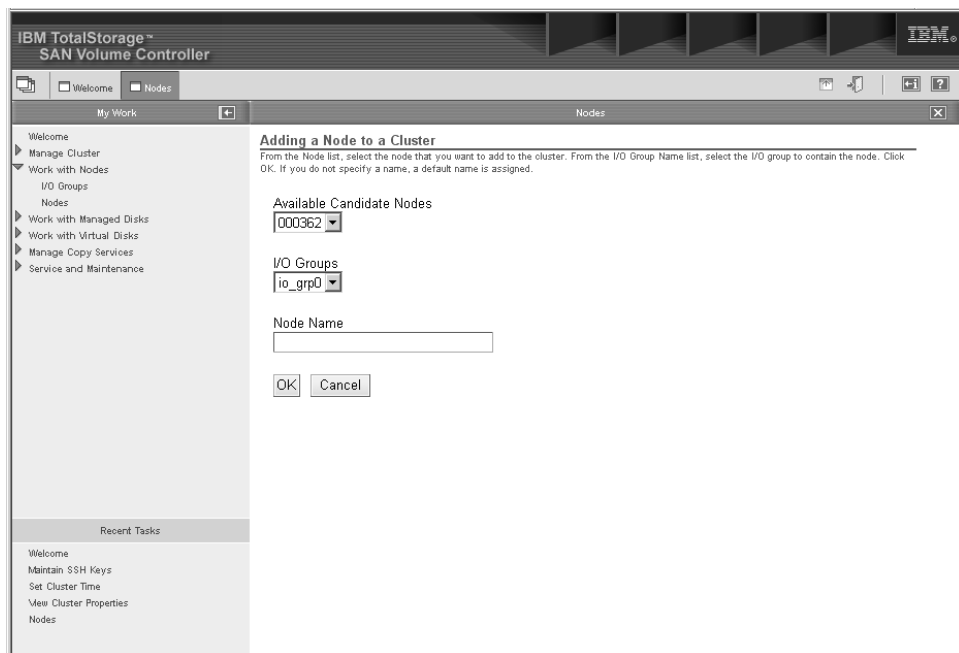


Figure 19. Add Node to Cluster panel

Notes:

- a. Each node in an I/O group must be on a different Cisco MDS 9000 Caching Services Module.
- b. If you do not supply a name, the cluster assigns a default name to the object. Whenever possible you should provide a meaningful name for objects to make identifying that object easier in the future.

Example:

In our hypothetical scenario, the nodes are called:
knode and lnode

In our hypothetical scenario, the I/O group is called:
io_group0

In our hypothetical scenario, the nodes are called:
mnode and nnode

In our hypothetical scenario, the I/O group is called:
io_group1

5. Repeat step 4 on page 78 for each of the nodes that you want to add to the cluster.

Displaying node properties using the Console for SAN Volume Controller for Cisco MDS 9000

Steps:

Perform the following steps to display the node properties:

1. Click **Work with Nodes** from the portfolio.
2. Click **Nodes** from the portfolio. The Nodes panel is displayed.
3. Select the name of the node that you want to view the details for. The Viewing General Details panel is displayed.

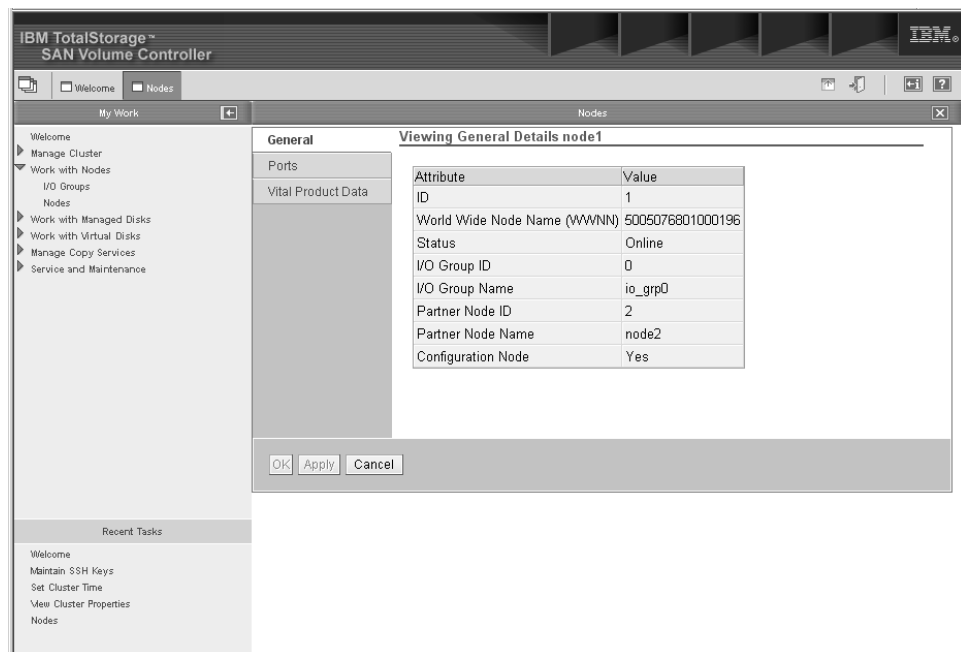


Figure 20. Viewing general details panel

4. Click **Ports** to view the WWPN port details. The Viewing Port Details panel for node1 is displayed.

5. Click **Vital Product Data** to view the node hardware details. The Viewing Vital Product Data panel is displayed.

Creating managed disk groups

Prerequisites:

If you intend to keep the virtual disk allocation within one disk controller system, you should ensure that the MDisk group that corresponds with a single disk controller system is presented by that disk controller system. This also enables non-disruptive migration of data from one disk controller system to another disk controller system and simplifies the decommissioning process should you wish to decommission a disk controller system at a later time.

You should also ensure that all MDisks allocated to a single MDisk group are of the same RAID type. This ensures that a single failure of a physical disk in the disk controller system does not take the entire group offline. For example, if you had three RAID-5 arrays in one group and added a non-RAID disk to this group, if the non-RAID disk fails then you will lose access to all the data striped across the group. Similarly, for performance reasons you should not mix RAID types.

Steps:

Perform the following steps to create a new MDisk group:

1. Click **Work with Managed Disks** from the portfolio.
2. Click **Managed Disk Groups** from the portfolio. The Filtering Managed Disk Groups panel is displayed.

Note: The filter panels can be used to prefilter the list of objects that are displayed. This reduces the number of objects returned to the Console for SAN Volume Controller for Cisco MDS 9000. This can be useful when you have a large number of objects (for example 4096 MDisks or 1024 VDIs) and you do not want to display them all. You can bypass the filtering and display all objects by clicking **Bypass Filter**.

3. Specify the filter criteria that you want to use. Click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Managed Disk Groups panel is displayed.
4. Select **Create MDisk Group** from the list. Click **Go**. The Create a Managed Disk Group wizard is displayed.
5. Type the name of the MDisk group and add the MDisks from the **Managed Disk Candidates** list.

Example:

In our hypothetical scenario, type

```
maindiskgroup
```

```
add MDisks
```

```
m1sk0, m1sk1, m1sk2, m1sk3
```

from the **Managed Disk Candidates** list.

6. Select the extent size from the list.

Example:

In our hypothetical scenario, select

for the extent size used within this MDisk group and click **OK**.

7. Repeat steps 4 on page 80 through 6 on page 80 for all of the MDisk groups that you want to create.

Example:

In our hypothetical scenario, repeat steps 4 on page 80 through 6 on page 80, in which the second MDisk group is named

bkpdiskgroup

with the following MDisks attached,

mnsk4, mnsk5, mnsk6, mnsk7

The extent size will be

16

MB.

Related topics:

- “Managed disks” on page 18
- “Managed disk groups” on page 20

Creating virtual disks

Steps:

Perform the following steps to create virtual disks:

1. Click **Work with Virtual Disks** from the portfolio.
2. Click **Virtual Disks** from the portfolio. The Filtering Virtual Disks panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The Viewing Virtual Disks panel is displayed.
4. Select **Create virtual disks** from the list. Click **Go**. The Create Virtual Disks wizard is displayed.
5. Perform the following steps to complete the wizard:
 - Select an I/O group, preferred node, and managed disk group.

Note: If there are no MDisk groups to select, you will be prompted to create one.

- Select the type and quantity of the virtual disks that you want to create.
- Type a name for the virtual disks.
- Set the attributes, such as MDisk candidates, capacity of the VDisks, and type of VDisks, for the virtual disks.
- Verify the attributes.

Related topics:

- Chapter 8, “Scenario: typical usage for the Console for SAN Volume Controller for Cisco MDS 9000,” on page 73

Creating hosts

Steps:

Perform the following steps to create a new host object:

1. Click **Work with Virtual Disks** in the portfolio.
2. Click **Hosts** in the portfolio. The Filtering Hosts panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The Hosts panel is displayed.
4. Select **Create Host** from the list. Click **Go**. The Creating Hosts panel is displayed.
5. Type the name of the logical host object.

Note: If you do not specify a name, a default name is assigned (for example, host0). Then assign a WWPN. A WWPN consists of 16 hexadecimal digits (for example, 210100e08b251dd4). You can select a WWPN from the list of candidates or you can manually enter a WWPN that is not on the list of candidates. Multiple WWPN's can be assigned to a single logical host object. Click **OK**.

Example:

In our hypothetical scenario, because a host name was not specified, the default name is:

host0

The World Wide Port Names (WWPNs) assigned to the host are:

210100e08b251dd4, 210100e08b251dd5

These WWPNs can be found by using your specific switches management application.

6. Repeat steps 4 through 5 for each host object that you want to create.

Example:

In our hypothetical scenario, repeat steps 4 through 5 and name the host:

demohost2

The World Wide Port Names (WWPNs) assigned to the host are:

210100e08b251dd6, 210100e08b251dd7

Related topics:

- "Host objects" on page 24

Creating virtual disk-to-host mappings

Steps:

Perform the following steps to create a new mapping:

1. Click **Work with Virtual Disks** from the portfolio.
2. Click **Virtual Disks** from the portfolio. The Filtering Virtual Disks panel is displayed.

3. Specify the filter criteria that you want to use. Click **OK**. The Viewing VDisks panel is displayed.
4. Select the virtual disk that you want to map to your host.
5. Select **Map VDisk to a Host** from the list. Click **Go**. The Map VDisk to Host panel is displayed.
6. Select the host to which you want to map the virtual disk or disks, and click **OK**.

Related topics:

- “Virtual disks” on page 22
- “Virtual disk-to-host mapping” on page 25

Showing VDisks mapped to a host

If a number of new VDisks are mapped to a host, and a number of devices are already running I/O operations, then a lot of errors may be logged. At the time the new VDisk is mapped, multiple recoverable errors can be logged in the event log. Decoding of the event log shows the errors to be caused by a check condition. The error states that there has been a change to the device information since the last LUN operation.

Steps:

Perform the following steps to show the VDisks that are mapped to a host:

1. Click **Work with Virtual Disks** in the portfolio.
2. Click **Hosts** in the portfolio. The Filtering Hosts panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The Hosts panel is displayed.
4. Select the host and select **Show the VDisks Mapped to this Host** from the list. Click **Go**.

Result:

The virtual disks mapped to this host are displayed in the Viewing Virtual Disks panel

Creating consistency groups

Steps:

Perform the following steps to create a FlashCopy consistency group:

1. Click **Manage Copy Services** in the portfolio.
2. Click **FlashCopy consistency groups** in the portfolio. The Filtering FlashCopy consistency groups panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The FlashCopy consistency groups panel is displayed.
4. Click **Create FlashCopy consistency groups**. The Create FlashCopy consistency groups panel is displayed.
5. Type the name of the consistency group in the **FCCGroup name** field. From the **FlashCopy Mappings** list, select the mappings you want in the consistency group and click **OK**. If you do not specify a name, a default name is assigned.

Example:

In our hypothetical scenario, the name of the consistency group is:
maintobkpfcopy

The mappings that should be added are:
main1copy, main2copy

Note: You could have created the FlashCopy consistency group before you created the mappings and then added the FlashCopy mappings to the consistency group. To add FlashCopy mappings in this way, you must use the Modifying FlashCopy Mapping panel or the Creating FlashCopy Mappings panel.

Creating FlashCopy mappings

1. Click **Manage Copy Services** in the portfolio.
2. Click **FlashCopy mappings** in the portfolio. The Filtering FlashCopy Mappings panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The Viewing FlashCopy Mappings panel is displayed.
4. Select **Create Mapping** from the list. Click **Go**. The Creating FlashCopy Mappings panel is displayed.
5. Type the name of the new FlashCopy mapping.

Example:

In our hypothetical scenario, the name of the FlashCopy mapping is:
main1copy

6. Select the source VDisk from the list.

Example:

In our hypothetical scenario, the name of the source VDisk is:
maindisk1

7. Select the target VDisk from the list.

Example:

In our hypothetical scenario, the name of the target VDisk is:
bkpdisk1

8. Select the priority for the background copy. Click **OK**.
9. Repeat steps 4 through 8 for each FlashCopy mapping that you want to create.

Chapter 9. Advanced function FlashCopy overview

Overview:

The following sections details the advanced FlashCopy functions that you can perform using the Console for SAN Volume Controller for Cisco MDS 9000.

Starting FlashCopy mappings

You can start or trigger a FlashCopy mapping from the Starting FlashCopy Mappings panel.

Perform the following steps to start a FlashCopy mapping:

1. Click **Manage Copy Services** in the portfolio.
2. Click **FlashCopy Mappings** in the portfolio. The Filtering FlashCopy mappings panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The Viewing FlashCopy Mappings panel is displayed.
4. Select the appropriate mapping's row from the table.
5. Select **Start a Mapping** from the drop down list. Click **Go**. The Starting FlashCopy mappings panel is displayed.

Stopping FlashCopy mappings

You can stop a FlashCopy mapping from the Stopping FlashCopy Mappings panel.

Steps:

Perform the following steps to stop a FlashCopy mapping:

1. Click **Manage Copy Services** in the portfolio.
2. Click **FlashCopy Mappings** in the portfolio. The Filtering FlashCopy Mappings panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The Viewing Mappings panel is displayed.
4. Select the appropriate mapping's row from the table.
5. Select **Stop a mapping** from the drop down list. Click **Go**. The Stopping FlashCopy mappings panel is displayed.

Deleting FlashCopy mappings

You can delete a FlashCopy mapping from the Deleting FlashCopy Mappings panel.

Perform the following steps to delete a FlashCopy mapping:

1. Click **Manage Copy Services** in the portfolio.
2. Click **FlashCopy mappings** in the portfolio. The Filtering FlashCopy mappings panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The Viewing FlashCopy mappings panel is displayed.
4. Select the appropriate mapping's row from the table.
5. Click **Delete a mapping** and click **Go**. The Deleting FlashCopy mapping panel is displayed.

Preparing FlashCopy consistency groups

Steps:

Perform the following steps to prepare a FlashCopy consistency group:

1. Click **Manage Copy Services** in the portfolio.
2. Click **FlashCopy Consistency Groups** in the portfolio.
The Filtering FlashCopy Consistency Groups panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The FlashCopy Consistency Groups panel is displayed.
4. Select the appropriate mapping's row from the table.
5. Click **Prepare a Consistency Group** from the drop-down menu and click **Go**.
6. When the status says "Prepared", you can start the FlashCopy consistency group.

Starting FlashCopy consistency groups

You can start or trigger a FlashCopy consistency group from the Starting FlashCopy Consistency Group panel.

Steps:

Perform the following steps to start or trigger a FlashCopy consistency group:

1. Click **Manage Copy Services** in the portfolio.
2. Click **FlashCopy Consistency Groups** in the portfolio. The Filtering FlashCopy Consistency Groups panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The FlashCopy Consistency Groups panel is displayed.
4. Select the appropriate mapping's row from the table.
5. Click **Start Consistency Group**. The Starting FlashCopy Consistency Groups panel is displayed.

Related topics:

- "FlashCopy consistency groups" on page 33

Stopping FlashCopy consistency groups

You can stop a FlashCopy consistency group from the Stopping FlashCopy Consistency Groups panel.

Steps:

Perform the following steps to stop a FlashCopy consistency group:

1. Click **Manage Copy Services** in the portfolio.
2. Click **FlashCopy Consistency Groups** in the portfolio. The Filtering FlashCopy Consistency Groups panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The FlashCopy Consistency Groups panel is displayed.
4. Select the appropriate mapping's row from the table.
5. Click **Stop Consistency Group**. The Stopping Consistency Groups panel is displayed.

Related topics:

- “FlashCopy consistency groups” on page 33

Deleting FlashCopy consistency groups

You can delete a FlashCopy consistency group from the Deleting FlashCopy consistency groups panel.

Steps:

Perform the following steps to delete a FlashCopy consistency groups:

1. Click **Manage Copy Services** in the portfolio.
2. Click **FlashCopy consistency groups** in the portfolio. The Filtering FlashCopy consistency groups panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The FlashCopy consistency groups panel is displayed.
4. Select the appropriate mapping’s row from the table.
5. Click **Delete Consistency Groups**. The Delete Consistency Groups panel is displayed.

Related topics:

- “FlashCopy” on page 29
- “FlashCopy consistency groups” on page 33

Chapter 10. Advanced functions overview for the SAN Volume Controller Console

Determining the WWPNs for a node using the Console for SAN Volume Controller for Cisco MDS 9000

Before you begin: In order to perform these steps, you must open a browser window and type `http://localhost:9080/ica` in the URL field. Then, select and launch the desired cluster. This will open up a secondary browser window where you can perform the steps below.

Steps:

Perform the following steps to determine a node's WWPNs:

1. List the nodes in the cluster by clicking **Work with Nodes** -> **Nodes** from the Portfolio.
2. For the node or nodes in question, select the node name link to view the node details.
3. Select the ports tab and note each WWPN.

Determining a storage subsystem name from its SAN Volume Controller Storage Software for Cisco MDS 9000 name using the Console for SAN Volume Controller for Cisco MDS 9000

Steps:

Perform the following steps to determine the storage subsystem name:

1. Click **Work with Managed Disks** --> **Disk Controllers** from the Portfolio.
2. Select the ID link for the storage subsystem in question. Write down the WWNN. This can be used to determine the actual storage subsystem by launching the native user interface or using the command line tools it provides to verify the actual storage subsystem that has this WWNN.

Determining the relationship between VDIs and MDIs using the Console for SAN Volume Controller for Cisco MDS 9000

Steps:

Perform the following steps to determine the relationship between VDIs and MDIs:

1. Click **Work with Virtual Disks** -> **Virtual Disks** from the portfolio.
2. Click **OK** on the Filtering Virtual Disks panel.
3. From the Viewing Virtual Disks panel, select the VDisk that you want to view the relationship between this VDisk and its MDIs.
4. Select the **Show MDIs** task. The Work with MDIs panel is displayed. This panel lists the MDIs and make up the selected VDisk. From the drop-down list, select **Show the MDIs** and then click **Go**. The View the MDIs for the VDisk panel is displayed. This panel lists the Managed Disks that make up the selected VDisk.

Steps:

Perform the following steps to determine the relationship between MDisks and VDIs:

1. Click **Work with Managed Disks** -> **Managed Disks** from the Portfolio. Click **OK** on the Filtering Managed Disks panel.
2. Select the MDisk that you want to view the relationship between this MDisk and its VDIs.
3. From the drop-down list, select **Show the VDisk for an MDisk**, and then click **Go**. The View the Virtual Disk for an MDisk panel is displayed. This panel lists the VDIs that make up the selected MDisk.

Determining the relationship between managed disks and RAID arrays or LUNs using the Console for SAN Volume Controller for Cisco MDS 9000

Each MDisk corresponds with a single RAID array, or a single partition on a given RAID array. Each RAID controller will define a LUN number for this disk. The LUN number and controller name or ID are needed to be able to determine the relationship between MDIs and RAID arrays or partitions.

Steps:

Perform the following steps to determine the relationship between MDIs and RAID arrays:

1. Click **Work with Managed Disks** -> **Managed Disks** from the Portfolio.
2. Click **OK** on the Filtering Managed Disks panel.
3. Click the name of the MDisk that you wish to view details for. Write down the controller name and controller LUN number.
4. Click **Work with Managed Disks** -> **Disk Controllers** from the Portfolio.
5. From the Viewing Disk Controller Systems panel, click the **ID** link to show the detailed view of the selected controller. Write down the vendor ID and the product ID and WWNN and use these to determine the controller that is presented to the MDisk.
6. From the native user interface for the given controller, list the LUNs it is presenting and match the LUN number with that noted in 3. This will tell you the exact RAID array and partition that corresponds with the MDisk.

Virtual disk-to-host mappings

You can view the virtual disks that are mapped to a host by using the Viewing Virtual Disks panel.

Steps:

Perform the following steps to view your virtual disk-to-host mappings:

1. Click **Work with Virtual Disks** -> **Hosts** from the portfolio.
2. The Filtering Virtual Disks panel is displayed. Specify the filter criteria that you want to use and click **OK**.
3. The Hosts panel is displayed. Select the host and from the drop-down list, select **Show the VDIs Mapped to this host**. Click **Go**. The View VDIs Mapped to Host panel is displayed.

Increasing the size of your cluster using the Console for SAN Volume Controller for Cisco MDS 9000

To increase the size of your cluster you need to add nodes in pairs to a new I/O group. Your existing cluster may have become a bottleneck and so you wish to increase throughput by adding more nodes to the cluster.

Steps:

Perform the following steps to increase the size of your cluster:

1. Add a node to increase the size of your cluster and repeat this procedure for the second node.
2. If you wish to balance the load between the existing I/O groups and the new I/O groups, migrate your VDisks to new I/O groups. Repeat this procedure for all VDisks you want to assign to the new I/O group.

Related topics:

- “Adding a node to increase the size of your cluster”
- “Migrating a VDisk to a new I/O group” on page 92

Adding a node to increase the size of your cluster

Steps:

Perform the following steps to add a node to increase the size of your cluster:

1. Click **Work with Nodes -> I/O groups** to determine which I/O group you wish to add the nodes to.
2. Look for the first I/O group listed that has a node count of 0. Write down the I/O group name. You will need it in the following step.
3. **Attention:** If you are re-adding a node to the SAN, ensure that you are adding the node to the same I/O group that it was removed from. Failure to do this can result in data corruption. Use the information that was recorded when the node was originally added to the cluster. If you do not have access to this information, call IBM Service to add the node back into the cluster without corrupting the data.

If you are adding the node into the cluster for the first time, record the following information:

- Card serial number
- Slot the Cisco MDS 9000 Caching Services Module is in
- Node (1 or 2) on the Cisco MDS 9000 Caching Services Module
- WWNN
- All WWPNs
- I/O group that contains the node

This can avoid a possible data corruption exposure if the node must be removed from and re-added to the cluster.

Note: This warning is also displayed on the Console for SAN Volume Controller for Cisco MDS 9000 panel when adding the node.

Add the node back into the cluster by selecting the **Add a node** task from the **Work with Nodes-> Nodes** panel.

4. Select the node from the list of available candidate nodes and select the I/O group from the list. The first I/O group with no nodes is io_grp3. You must add nodes to the lowest numbered group before adding nodes to higher numbered groups.
5. Optionally enter a node name for this node.
6. Verify that the node is online by clicking the **Refresh** button on the Viewing Nodes panel. You may need to close the panel and reopen it for the refresh to take effect.
7. You may also need to modify the configuration of your disk controller systems. If your disk controller system uses a mapping technique to present its RAID arrays or partitions to the cluster, you will need to modify the port groups that belong to the cluster because the WWNN and WWPN's of the node have changed.

Related topics:

- Chapter 8, "Scenario: typical usage for the Console for SAN Volume Controller for Cisco MDS 9000," on page 73
- "Displaying node properties using the Console for SAN Volume Controller for Cisco MDS 9000" on page 79
- Chapter 29, "Configuring and servicing storage subsystems," on page 217

Migrating a VDisk to a new I/O group

You can migrate a VDisk to a new I/O group to manually balance the workload across the nodes in the cluster. You may end up with a pair of nodes that are overworked and another pair that are under-worked. Follow this procedure to migrate a single VDisk to a new I/O group. Repeat for other VDIs as required.

Attention: This is a disruptive procedure, access to the VDisk will be lost while you follow this procedure. Under no circumstances should VDIs be moved to an offline I/O group. You must ensure the I/O group is online before moving the VDIs to avoid data loss scenarios.

Steps:

Perform the following steps to migrate a single VDisk:

1. Quiesce all I/O operations for the VDisk. You may need to determine the hosts that are using this VDisk.
2. Before migrating the VDisk, it is essential that for each vpath presented by the VDisk you intend to move, the SDD configuration is updated to remove the vpaths in question. Failure to do this may result in data corruption. See *IBM TotalStorage Subsystem Device Driver: User's Guide* for details about how to dynamically re-configure SDD for the given host operating system.
3. Any FlashCopy mappings or Remote Copy relationships that use this VDisk should be stopped and deleted. To check if the VDisk is part of a mapping or relationship, perform the following steps:
 - a. Click **Work with Virtual Disks** -> **Virtual Disks** from the portfolio.
 - b. The Filtering Virtual Disks panel is displayed. Specify the filter criteria that you want to use and click **OK**. The Viewing VDIs panel is displayed.
 - c. Click on the VDisk name to view the details.
 - d. Look for the **FlashCopy ID** and **Remote Copy ID** fields. If these are not blank then the VDisk is part of a mapping or relationship.

4. Migrate the VDisk by selecting the VDisk from the Viewing Virtual Disks panel and select **Modify a VDisk** from the drop-down list. Click **Go**. Change only the I/O group to the new I/O group name and click **OK**.
5. It is now necessary to follow the SDD procedure to discover the new vpaths and to check that each vpath is now present with the correct number of paths. See the *IBM TotalStorage Subsystem Device Driver: User's Guide* for details on how to dynamically re-configure SDD for the given host operating system.

Related topics:

- “Advanced function Remote Copy overview” on page 104
- “Determining the relationship between VDisks and MDisks using the Console for SAN Volume Controller for Cisco MDS 9000” on page 89

Replacing a faulty node with a spare node using the Console for SAN Volume Controller for Cisco MDS 9000

You will need the following information to avoid the possibility of data corruption when the node is removed and re-added to the cluster. If you do not have this information, call IBM Service to add the node back into the cluster without corrupting the data.

- Card serial number
- Slot in which the Cisco MDS 9000 Caching Services Module is located
- Node (1 or 2) on the Cisco MDS 9000 Caching Services Module
- WWNN
- All WWPNS
- I/O group that contains the node

Steps:

Perform the following steps to replace a faulty node in the cluster:

1. Verify the name and ID of the node that you wish to remove. Click **Work with Nodes**.
 - a. If the node is faulty, it will be shown as offline. Ensure the partner node in the I/O group is online.
 - 1) If the other node in the I/O group is offline, start Directed Maintenance Procedures to determine the fault.
 - 2) If you were directed here by the Directed Maintenance Procedures and the partner node in the I/O group has failed, see “Recovering from offline VDisks after a node or an I/O group failed” on page 94.
 - b. If you are replacing the node for other reasons, determine the node you wish to replace and again ensure the partner node in the I/O group is online.
 - 1) If the partner node is offline, you will lose access to the VDisks that belong to this I/O group if you continue. Start the Directed Maintenance Procedures and fix the other node before proceeding.
2. Remove the node from the cluster by selecting the node in the list and click **Remove Node** from the drop-down list. Click **Go**.

Attention: Record the following information:

- Card serial number
- Slot in which the Cisco MDS 9000 Caching Services Module is located
- Node (1 or 2) on the Cisco MDS 9000 Caching Services Module
- WWNN
- All WWPNs
- I/O group that contains the node

This can prevent possible data corruption when the node is re-added to the cluster.

3. If the nodes were repaired by replacing it with another node, then the WWNN for the node will change. In this case the following additional steps are required:
 - At the end of the recovery process, it will be necessary to follow the SDD procedure to discover the new paths and to check that each vpath is now presenting the correct number of paths. See the *IBM Subsystem Device Driver (SDD) User's Guide* for details about adding paths to existing paths.
 - You may also need to modify the configuration of your disk controller systems. If your controller uses a mapping technique to present its RAID arrays or partitions to the cluster, you will need to modify the port groups that belong to the cluster because the WWNN or WWPN's of the node have changed. See Chapter 29, "Configuring and servicing storage subsystems," on page 217 for more information.

Note: The WWPN assignments are done through a slot in the switch. If a node is replaced in the same slot, a reassignment of the WWNN or WWPN is not needed. If the node is placed in a different slot, the WWNN or WWPN can be reassigned through a switch command.

4. Add the node back into the cluster. From the Viewing Nodes panel, select **Add a node** from the drop-down list and click **Go**. The Adding a Node to a Cluster panel displays.
5. Select the candidate node from the drop-down list.
6. Select the I/O group to which the node is being added.
7. If you wish you may also provide a name for the node.
8. Click **Refresh** on the Viewing Nodes panel and verify that the node is online.

Related topics:

- "Recovering from offline VDIsks after a node or an I/O group failed"
- "Replacing a faulty node in the cluster using the CLI" on page 177

Recovering from offline VDIsks after a node or an I/O group failed

Prerequisites:

If you have lost both nodes in an I/O group and have therefore, lost access to all the VDIsks that are associated with the I/O group, then you must perform one of the following procedures to regain access to your VDIsks. Depending on the failure type, you may have lost data that was cached for these VDIsks, therefore, they have gone offline.

Context:

Data loss scenario 1 One node in an I/O group failed and failover started on the second node. During this time, the second node in the I/O group fails before the cache has become write through mode. The first node is successfully repaired but its cache data is stale, therefore, it cannot be used. The second node is repaired or replaced and has lost its hard-end data, therefore, the node has no way of recognizing that it is part of the cluster.

Steps:

Perform the following steps to recover from an offline VDisk:

1. Recover the node and include it back into the cluster.
2. Move all the offline VDIs to the recovery I/O group.
3. Move all the offline VDIs back to their original I/O group.

Context:

Data loss scenario 2 Both nodes in the I/O group have failed and have been repaired. The nodes have lost their hard-end data, therefore, the nodes have no way of recognizing that they are part of the cluster.

1. Move all the offline VDIs to the recovery I/O group
2. Move both recovered nodes back into the cluster
3. Move all the offline VDIs back to their original I/O group.

Related topics:

- “Recovering a node and including it back into the cluster”
- “Moving offline VDIs to the recovery I/O group” on page 96
- “Moving offline VDIs to their original I/O group” on page 97

Recovering a node and including it back into the cluster

After a node or an I/O group fails, you can use the following procedure to recover a node and include it back into the cluster.

Steps:

Perform the following steps to recover a node and include it back into the cluster:

1. Verify that the node is offline by viewing the Viewing Nodes Panel. To do this, click **Work with Nodes** -> **Nodes** from the portfolio. The Viewing Nodes panel opens.
2. Remove the old instance of the offline node from the cluster by selecting the node and then selecting **Remove a node** from the drop-down list. Click **Go**.
3. Verify that the node can be seen on the fabric.
4. If the node is repaired by replacing it with another node, then the WWNN for the node will change. In this case, the following additional steps are required:
 - a. At the end of the recovery process it will be necessary to follow the SDD procedure to discover the new paths and to check that each vpath is now presenting the correct number of paths. For information about adding paths to existing vpaths. see the *IBM TotalStorage Subsystem Device Driver: User's Guide*.
 - b. You may also need to modify the configuration of your disk controller systems. If your disk controller system uses a mapping technique to present

its RAID arrays or partitions to the cluster you will need to modify the port groups that belong to the cluster because the WWNN or WWPN's of the node have changed.

Attention: If more than one I/O group is affected, ensure that you are adding the node to the same I/O group that it was removed from. Failure to do this can result in data corruption. Use the information that was recorded when the node was originally added to the cluster. This can avoid a possible data corruption exposure if the node must be removed from and re-added to the cluster. If you do not have access to this information, call IBM Service to add the node back into the cluster without corrupting the data.

Attention: If you are adding the node into the cluster for the first time, record the following information:

- Node serial number
- WWNN
- All WWPNs
- I/O group that contains the node

Note: This warning is also displayed on the Console for SAN Volume Controller for Cisco MDS 9000 panel when adding the node.

5. Add the node back into the cluster. From the **Viewing Nodes** panel, select **Add a Node** from the drop-down list. Click **Go**. The **Adding a Node to a Cluster** panel is displayed. Select the node from the list of candidate nodes and select the I/O group from the list. Optionally enter a node name for this node.
6. Verify that the node is online by clicking the **Refresh** button on the **Viewing Nodes** panel.

Note: You may need to close the panel and reopen it for the refresh to take effect.

Related topics:

- “Recovering from offline VDIs after a node or an I/O group failed” on page 94
- “Moving offline VDIs to the recovery I/O group”
- “Moving offline VDIs to their original I/O group” on page 97

Moving offline VDIs to the recovery I/O group

After a node or an I/O group fails, you can use the following procedure to move offline VDIs to the recovery I/O group.

Steps:

Perform the following steps to move offline VDIs to the recovery I/O group:

1. List all VDIs that are offline and that belong to the I/O group in question by selecting **Work with Virtual Disks** -> **Virtual Disks** from the portfolio. From the Filtering Virtual Disks panel, enter the <iogrname> in the I/O Group filter box and select **offline** as the Status.
2. For each VDI returned, select the VDI and select **Modify a VDI** from the drop-down list box. Click **Go**. In the Modifying Virtual Disk panel, only change the I/O group to **Recovery I/O group**. You may be asked to confirm and force the move, select to force the move.

Related topics:

- “Recovering from offline VDIs after a node or an I/O group failed” on page 94
- “Moving offline VDIs to their original I/O group”
- “Recovering a node and including it back into the cluster” on page 95

Moving offline VDIs to their original I/O group

After a node or an I/O group fails, you can use the following procedure to move offline VDIs to their original I/O group.

Attention: Under no circumstances should VDIs be moved to an offline I/O group. Ensure the I/O group is online before moving the VDIs back to avoid any further data loss.

Steps:

Perform the following steps to move offline VDIs to their original I/O group:

1. For each VDI listed in the Viewing Virtual Disks panel (which you can access by clicking **Work with Virtual Disks** -> **Virtual Disks**). Select the VDI and then select **Modify a VDI** from the drop-down list box. Click **Go**. In the.
2. Verify that the VDIs are now online by closing the Viewing Virtual Disks panel and opening it again. This time, from the Filtering Virtual Disks panel, only enter the <iogrname> in the I/O group filter box. On the Viewing Virtual Disks panel, all of the VDIs should be online.

Related topics:

- “Recovering from offline VDIs after a node or an I/O group failed” on page 94
- “Recovering a node and including it back into the cluster” on page 95
- “Moving offline VDIs to the recovery I/O group” on page 96

Replacing an HBA in a host using the Console for SAN Volume Controller for Cisco MDS 9000

This procedure describes how to notify the SAN Volume Controller Storage Software for Cisco MDS 9000 of a change to a defined host object. It is sometimes necessary to replace the HBA that connects the host to the SAN, at this time you must notify the SAN Volume Controller Storage Software for Cisco MDS 9000 of the new WWPN's that this HBA contains.

Prerequisites:

Ensure your Cisco MDS 9000 is zoned correctly.

Steps:

Perform the following steps to replace an HBA in a host:

1. Locate the host object that corresponds with the host in which you have replaced the HBA. Click **Work with Virtual Disks** -> **Hosts** from the portfolio. From the **Filter Hosts** panel, select the filter criteria. Click **OK**. On the **Viewing Hosts** panel, select the host object and then select **Add Ports** from the drop-down list. Click **Go**.
2. Add the new ports to the existing host object. Select the candidate WWPNs from the Available Ports list and click **Add**. Complete the task by clicking **OK**.

3. Remove the old ports from the host object. From the Viewing Hosts panel, select the host object and select **Delete ports** from the drop-down list. Click **Go**. Select the WWPNs you wish to remove (the ones that correspond with the old HBA that was replaced). Click **Add** to add them to the list of WWPNs to be deleted. Complete the task by clicking **OK**.
4. Any mappings that exist between the host object and VDisks will automatically be applied to the new WWPNs. Therefore, the host should see the VDisks as the same SCSI LUNs as before. See *IBM TotalStorage Subsystem Device Driver: User's Guide* for adding paths to existing vpaths.

Related topics:

- Chapter 29, “Configuring and servicing storage subsystems,” on page 217

Resetting the Node

Steps:

Perform the following steps to reset a node so that it can be added to an IO group:

1. Click Reset Node from the portfolio.
2. Reset the node by entering these values:
 - Fabric switch IP address
 - Fabric switch user ID
 - Fabric switch password
3. From the Slot Number list, select the slot number.
4. From the Port Number list, select the port number.
5. Click OK to reset the node.

Deleting hosts

Prerequisites:

A deletion fails if there are any VDisk-to-host mappings for the host. If you attempt to delete the host and it fails due to the existence of VDisk mappings, then you are presented with the opportunity to perform a forced delete, which will delete VDisk mappings before deleting the host.

Steps:

Perform the following steps to delete a host object:

1. Click **Work with Virtual Disks** in the portfolio.
2. Click **Hosts** in the portfolio. The Filtering Hosts panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The Viewing Hosts panel is displayed.
4. From the Viewing Hosts panel, select the host that you want to delete. Then, from the drop-down list, select **Delete a host**.
5. Click **Go**.
6. On the Deleting Hosts panel, verify that you are deleting the correct host and click **OK**.

Result:

When you delete a host object, all active ports are added to the **Available Ports** list.

Related topics:

- “Host objects” on page 24

Adding a new storage controller to a running configuration using the Console for SAN Volume Controller for Cisco MDS 9000

Prerequisites:

You can add a new storage controller to your SAN at any time. Follow the switch zoning guidelines and also ensure the controller is setup correctly for use with the SAN Volume Controller Storage Software for Cisco MDS 9000.

You should create one or more arrays on the new controller. It is recommend that you use, RAID-5, RAID-1 or RAID-0+1 (sometimes called RAID-10) for maximum redundancy and reliability. Generally 5+P arrays are recommend. If your controller provides array partitioning we recommend that you create a single partition from the entire capacity available in the array, remember the LUN number that you assign to each partition as you will need this later. You should also follow the mapping guidelines (if your storage controller requires LUN mapping) to map the partitions or arrays to the SAN Volume Controller Storage Software for Cisco MDS 9000 ports. You can determine the SAN Volume Controller Storage Software for Cisco MDS 9000 ports by following the procedure for determining WWPNs..

Steps:

Perform the following steps to add a new storage controller to a running configuration:

1. To ensure that the cluster has detected the new storage (MDisks) click **Work with Managed Disks** -> **Managed Disks** from the Portfolio. On the Filtering Managed Disks panel, enter the filter criteria and click **OK**. On the Viewing Managed Disks panel, select **Discover MDisks** from the drop-down list box and click **Go**.
2. The controller itself will have automatically been assigned a default name. If you are unsure which controller is presenting the MDisks click **Work with Managed Disks** -> **Disk Controller Systems**. You should see a new controller listed (the one with the highest numbered default name). You must determine the storage controller name to validate that this is the correct controller.
3. Close and reopen the Viewing Managed Disks panel (to reopen, click on **Work with Managed Disks** -> **Managed Disks** from the Portfolio. On the Filtering Managed Disks panel, select **unmanaged** as the Mode and a Controller Name that corresponds with the new controller's name. The MDisks shown on the Viewing Managed Disks panel should correspond with the RAID arrays or partitions you created. Remember the field controller LUN number, this corresponds with the LUN number you assigned to each of the arrays or partitions.
4. It is recommended that you create a new managed disk group and add only the RAID arrays that belong to the new controller to this MDisk group. You should also avoid mixing RAID types, so for each set of RAID array types (for example, RAID-5, RAID-1) you should create a new MDisk group. Give this MDisk group an appropriate name, so if your controller is called FAST650-fred, and the MDisk group contains RAID-5 arrays, call it something like F600-fred-R5.

5. Click **Work with Managed Disks** -> **Managed Disk Groups** from the Portfolio. On the Filtering Managed Disk Groups panel, choose the filter criteria and click **OK**. On the Viewing Managed Disk Groups panel, select **Create an MDisk** from the drop-down list box and click **Go**. On the Creating Managed Disk Groups panel, enter the name you wish to give this group, select the MDisks you wish to add from the candidate list and click **Add**. Select the extent size you wish this group to have and click **OK**.

Related topics:

- “Determining the WWPNs for a node using the Console for SAN Volume Controller for Cisco MDS 9000” on page 89
- Chapter 29, “Configuring and servicing storage subsystems,” on page 217
- “Determining a storage subsystem name from its SAN Volume Controller Storage Software for Cisco MDS 9000 name using the Console for SAN Volume Controller for Cisco MDS 9000” on page 89

Removing a storage subsystem using the Console for SAN Volume Controller for Cisco MDS 9000

You can replace or decommission an old storage subsystem by following the procedure below. This procedure takes you through adding the new device, migrating the data off of the old device and removing the old MDisks.

This function can also be performed by migrating all the VDIs that are using storage in this MDisk group to another MDisk group. This procedure has an advantage if you wish to consolidate the VDIs in a single or new group. However, you can only migrate a single VDI at a time. The procedure outlined below will migrate all the data through a single command. If you wish to migrate the VDIs however, follow the procedure for all VDIs that are using this group. You can determine the relationship between VDIs and MDisks by following the procedure.

This procedure can also be used to remove or replace a single MDisk in a group. If an MDisk has suffered a partial failure, such as a degraded array, and you can still read the data from the disk but cannot write to it, you can follow this procedure to replace just one MDisk. In steps 1 and 3 on page 101 only add or remove a single MDisk rather than a list of MDisks.

Prerequisites:

All the MDisks that belong to the storage subsystem that is being decommissioned belong to a single MDisk group. You need to repeat this procedure for each MDisk group in turn before removing the old device.

Steps:

Perform the following steps to remove a storage subsystem:

1. Add new storage.
2. Click **Work with Managed Disks** -> **Managed Disk Groups**. On the Filtering Managed Disk Groups panel, enter the filter criteria and click **OK**. On the Viewing Managed Disk Groups panel, select **Add MDisk** from the drop-down list and click **Go**. On the Adding Managed Disks to Managed Disk Groups panel, select the new MDisks from the candidate list and click **Add**. Click **OK** to complete the task.

3. You should now have an MDisk group that contains the old MDisks (those to be decommissioned) and the new MDisks (those that are replacing them). Ensure that the capacity of the new MDisks is the same or exceeds that of the old MDisks before proceeding.
4. Force delete the old MDisks from the group. This will migrate all the data from the old MDisks to the new MDisks. From the drop-down list box, select **Remove MDisks** and click **Go**. From the Deleting Managed Disks from Managed Disk Groups panel, select the MDisks you wish to remove and click **Add**. Click **OK** to complete the task. When prompted click **Forced Delete**. Depending upon the number and size of the MDisks, and the number and size of the VDIsks that are using these MDisks, this operation will take some time to complete although the task will complete immediately.
5. The only way to check progress is by using the command line interface. Issue the following command:


```
svcinfolismigrate
```
6. When all the migration tasks have completed, for example, the command in step 5 returns no output, you can safely remove the old storage subsystem from the SAN.
7. Once you have removed the old storage subsystem from the SAN, rerun the detect MDisks task to remove the entries for the old MDisks.

Related topics:

- “Adding a new storage controller to a running configuration using the CLI” on page 183

Shrinking virtual disks

Context:

VDisks can be reduced in size should it be required. However, if the VDisk contains data that is being used, **under no circumstances should you attempt to shrink a VDisk without first backing up your data**. The SAN Volume Controller Storage Software for Cisco MDS 9000 arbitrarily reduces the capacity of the VDisk by removing a partial, one or more extents from those allocated to the VDisk. You cannot control which extents are removed and so you cannot guarantee that it is unused space that is removed.

Attention: This feature should *only* be used to make a target or auxiliary VDisk the same size as the source or master VDisk when creating FlashCopy mappings or Remote Copy relationships. You should also ensure that the target VDisk is not mapped to any hosts prior to performing this operation.

Steps:

Perform the following steps to shrink a VDisk:

1. Validate that the VDisk is not mapped to any host objects. If the VDisk is mapped, data is displayed.
2. You can determine the exact capacity of the source or master VDisk. Issue the following command:

```
svcinfolsvdisk -bytes <vdiskname>
```

Note: It is not possible to determine the exact size using the Console for SAN Volume Controller for Cisco MDS 9000.

3. Click **Work with Virtual Disks** from the portfolio.
4. Click **Virtual Disks** from the portfolio. The Filtering Virtual Disks panel is displayed.
5. Specify the filter criteria that you want to use. Click **OK**. The Viewing Virtual Disks panel is displayed.
6. Select the VDisk you want to shrink and select **Shrink a VDisk** from the drop-down list. Click **Go**. The Shrinking Virtual Disks panel is displayed.
7. In the Reduce Capacity By field, enter the amount of capacity of the VDisk that you want to reduce by and click **OK**.

Related topics:

- “Virtual disks” on page 22

Migrating virtual disks

Context:

The SAN Volume Controller Storage Software for Cisco MDS 9000 provides various data migration features. These can be used to move the placement of data both within MDisk groups and between MDisk groups. These features can be used concurrent with I/O operations. There are two ways in which you can migrate data:

1. Migrating data (extents) from one MDisk to another (within the same MDisk group). This can be used to remove hot or overutilized MDisks. This can only be performed using the CLI.
2. Migrating VDIs from one MDisk group to another. This can be used to remove hot MDisk groups, for example, you can reduce the utilization of a group of MDisks.

You can determine the usage of particular MDisks by gathering I/O statistics about MDisks and VDIs. Once you have gathered this data, you can analyze it to determine which VDIs or MDisks are hot. This procedure then takes you through migrating VDIs from one MDisk group to another.

When a migrate command is issued, a check is made to ensure that the destination of the migrate has enough free extents to satisfy the command. If it does, the command proceeds, but will take some time to complete. During this time, it is possible for the free destination extents to be consumed by another process, for example, by creating a new VDisk in the destination MDisk group or by starting more migrate commands. In this scenario, when all the destination extents have been allocated the migration commands suspend and an error is logged (error id 020005). There are two methods for recovering from this situation:

1. Add additional MDisks to the target MDisk group. This will provide additional extents in the group and will allow the migrations to be restarted (by marking the error as fixed).
2. Migrate one or more VDIs that are already created from the MDisk group to another group. This will free up extents in the group and allow the original migrations to be restarted (again by marking the error as fixed).

Steps:

Perform the following steps to migrate VDIs between MDisk groups:

1. Isolate any VDIs that are overutilized. You can determine this by requesting an I/O statistics dump and analyzing the output. To start I/O statistics gathering, select **Manage Cluster** from the portfolio, then select the **Start statistics collection** task. Enter 15 minutes for the interval and click **OK**. This will generate a new I/O statistics dump file approximately every 15 minutes. Wait for at least 15 minutes before proceeding to the next step.
2. Select **Service and Maintenance** from the portfolio and then select the **List dumps** task.
3. Click on the **I/O Statistics logs** link in the panel displayed. This will list the I/O statistics files that have been generated. These are prefixed with *m* and *Nm* for MDisk statistics and *v* for VDisk statistics. In order to view the contents of a file, select a filename and then select **View** from the drop-down list. Click **Go**. Click on one of the filenames to view the contents.
4. Analyze the dumps to determine which VDIs are hot. It may be helpful to also determine which MDisks are being heavily utilized as you can spread the data they contain more evenly across all the MDisks in the group.
5. Stop the statistics collection again by selecting **Manage Cluster** from the portfolio and then select **Stop statistics collection** task.

Once you have analyzed the I/O statistics data, you can determine which VDIs are hot. You also need to determine which MDisk group you wish to move this VDisk to. Either create a new MDisk group or determine an existing group that is not yet over utilized. You can do this by checking the I/O statistics files generated above and ensuring that the MDisks or VDIs in the target MDisk group are less utilized than the source group.

1. Click **Work with Virtual Disks** from the portfolio.
2. Click **Virtual Disks** from the portfolio. The Filtering Virtual Disks panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The Viewing Virtual Disks panel is displayed.
4. Select the VDisk you want to migrate and select **Migrate a VDisk** from the drop-down list. Click **Go**. The Migrating Virtual Disks panel is displayed.
5. Select the MDisk group name and the number of threads to devote to this process and then click **OK**.

Related topics:

- “Virtual disks” on page 22

Creating image-mode virtual disks

The SAN Volume Controller Storage Software for Cisco MDS 9000 enables you to import storage that contains existing data and continue to use this storage but make use of the advanced functions, such as, Copy Services, data migration, and the cache. These disks are known as image mode virtual disks.

Make sure you are aware of the following before converting your virtual disks:

1. Managed disks that contain existing data cannot be differentiated from managed disks that are blank. Therefore, it is vital that you control the introduction of these disks to the cluster. It is recommended that you introduce these disks one at a time. For example, map a single LUN from your RAID controller to the cluster and refresh the view of managed disks. The newly detected disk is displayed.

2. *Do not* add a managed disk that contains existing data to a managed disk group manually. If you do, the data will be lost. When you create an image mode virtual disk from this managed disk, it will be automatically added to the managed disk group. However, it will be added in such a way that the cluster can control how it is added to ensure the data is not lost.

Go to the following Web site for more information:

www.ibm.com/storage/support/2062-2300

Steps:

Perform the following steps to convert your virtual disk from image mode to manage mode:

1. Map a single RAID array or LUN from your RAID controller to the cluster. You can do this either through a switch zoning or a RAID controller based on your host mappings.
2. Re-scan the list of managed disks from the Console for SAN Volume Controller for Cisco MDS 9000. Click **Work with Managed Disks** → **Managed Disks**. You can then filter through the unmanaged mode disks.

Optionally, if the new managed disk is not listed you may need to run a fabric level discovery. Select **Work with Managed Disks** and from the Viewing Managed Disks panel choose **Discover MDisks** from the drop-down list. After a few minutes, refresh the view of managed disks and the new managed disk should be displayed.

3. Convert the managed disk into an image mode virtual disk. In the Console for SAN Volume Controller for Cisco MDS 9000, click **Work with Managed Disks** → **Managed Disks** from the Portfolio. On the Filtering Managed Disks panel, specify the filter criteria. Click **OK**. On the Viewing Managed Disks panel, select the specific managed disk and from the drop-down list, click **Create a VDisk in image mode**. Click **Go..** This will bring up the create image mode virtual disk wizard. You can select the managed disk group to add this managed disk to, and the I/O group that will provide the up-stream data path for the virtual disk.
4. Map the new virtual disk to the hosts that were previously using the data that the MDisk contains. In the Console for SAN Volume Controller for Cisco MDS 9000, select **Work with Virtual Disks** → **Virtual Disks**. On the Filtering Virtual Disks (VDisks) panel, enter the filter criteria or click **Bypass filter**. On the Viewing Virtual Disks panel, select the VDisk and then from the drop-down list choose **Map a VDisk to a host**, and click **Go**.

If you wish to convert this virtual disk or managed disk to actually virtualize the storage, you can transform the image mode virtual disk into a striped virtual disk by migrating the data on the managed disk to other managed disks in the same group. This procedure can only be performed using the command-line interface (CLI).

Related topics:

- “Creating an image mode VDisk from an unmanaged MDisk using the CLI” on page 193

Advanced function Remote Copy overview

For detailed information about how to perform advanced FlashCopy and Remote Copy functions, go to the following Web site:

www.ibm.com/redbooks

Advanced function cluster overview

Overview:

The following sections details the advanced cluster functions that you can perform using the Console for SAN Volume Controller for Cisco MDS 9000.

Removing a node from a cluster

You can delete or remove a node from the cluster with the Removing a Node from Cluster panel.

Attention: Before deleting or removing a node from the cluster you should quiesce all I/O operations that are destined for this node. Failure to do so may result in failed I/O operations being reported to your host operating systems.

Prerequisites:

Attention: If you are deleting or removing a single node, and the other node in the I/O group is online, be aware that the cache on the partner node will go into write-through mode and that you are exposed to a single point of failure should the partner node fail. Proceed to step 3 in the following procedure.

Attention: If you are deleting or removing a node, and this is the last node in the I/O group, you will lose access to all VDisks served by this I/O group. Ensure that all VDisks are not being accessed or do not contain data that you wish to continue to access, or ensure that they have been migrated to a different (online) I/O group.

1. Begin by determining the VDisks that are still assigned to this I/O group:
 - a. Determine the VDisks in question by requesting a filtered view of VDisks where the filter attribute is the I/O group in question.
 - b. Once you have a list of VDisks, determine the hosts that they are mapped to by following the procedure called, Determining the hosts that a VDisk is mapped to.
 - c. Once you have determined the hosts and are sure that you do not wish to maintain access to these VDisks proceed to 3 on page 106.
 - d. If you determine that some or all of the VDisks assigned to this I/O group do contain data that you wish to continue to access, you should follow the procedure called, Migrating a VDisk to a new I/O group.
2. Before performing the SDD path removal procedure described in 3 on page 106 you should power off the node that you intend to remove, unless this is the last node in the cluster. This ensures that SDD does not re-discover the paths that are manually removed before you issue the delete node request.

Notes:

- a. If the node you are removing is the configuration node, it might take a minute or more before you can perform the delete node request. You must wait for the configuration node failover to occur.
- b. If the node you are removing is the last node in the cluster, the Console for SAN Volume Controller for Cisco MDS 9000 might seem to hang for up to 3 minutes because you have removed the last access point to the cluster.

Attention: Deleting the configuration node or shutting down the configuration node may result in the SSH command hanging. If this happens, you should either wait for the SSH command to timeout or kill the command and ping the cluster IP address until it responds. At this point the failover has completed and you can start issuing commands again.

Note: If you power back on the node that has been removed and it is still connected to the same fabric or zone it will attempt to re-join the cluster. At this point the cluster will tell the node to remove itself from the cluster and the node will become a candidate for addition to this cluster or another cluster. If you are adding this node back into the cluster, ensure that you add it back to the same I/O group that it was previously a member of. Failure to do so may result in data corruption.

3. Before removing the node, it is essential that for each vpath presented by the VDisks you intend to remove, the SDD configuration is updated to remove the vpaths in question. Failure to do this may result in data corruption. See the *IBM TotalStorage Subsystem Device Driver: User's Guide* for details about how to dynamically reconfigure SDD for the given host operating system.
4. Proceed to 1.

Steps:

Perform the following steps to remove a node from the cluster:

1. Click **Work with Nodes** from the portfolio.
2. Click **Nodes** from the portfolio. The Nodes panel is displayed.

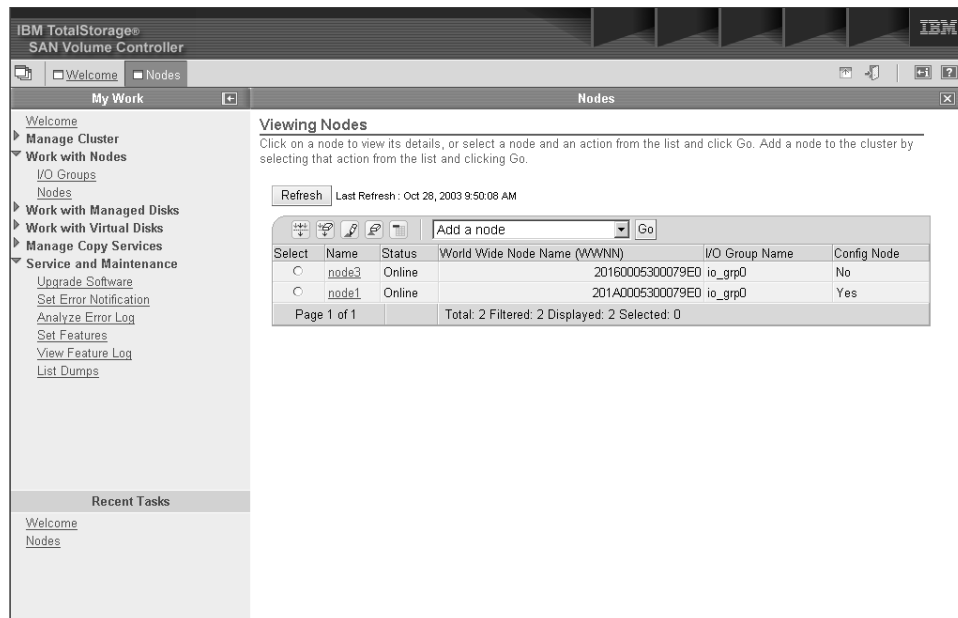


Figure 21. Console for SAN Volume Controller for Cisco MDS 9000 nodes panel

3. Select the node you want to remove and Select **Remove a node** from the list. Click **Go**. The Removing a Node from a Cluster panel is displayed. Click **Yes** to confirm removal of the node from the cluster.

Related topics:

- “Determining the host that a VDisk is mapped to” on page 172

- “Migrating a VDisk to a new I/O group” on page 92

Viewing feature settings and log

Steps:

Perform the following steps to view the feature settings:

1. Click **Service and Maintenance** from the portfolio.
2. To view the feature settings, click **Set Features** from the portfolio. To view the feature log, click **View Feature Log**.

Configuring error notification settings

You can configure the error notification settings for the cluster from the Set Error Notification Settings panel.

Steps:

Perform the following steps to configure the error notification settings:

1. Click **Service and Maintenance** in the portfolio.
2. Click **Set Error Notifications** from the Portfolio to display the existing error notification settings and to change them. The Modifying Error Notification Settings panel is displayed. The Modifying Error Notification Settings panel enables you to update your error notification settings. You can select whether the cluster raises an SNMP trap for entries that are added to the cluster error or event log. Three levels of notification are possible:
 - **None** No error or status changes will be sent.
 - **No State** You will be notified of errors, but you will not be notified of status changes.
 - **All** You will be notified of all errors and status changes.

If you have an SNMP manager installed or if you want to be notified by e-mail of errors or events, you should enable error notification. If you have an SNMP manager installed, you should enable error notification. If you select **All** or **No State** notification, you must specify a destination for the notification.

3. Click **Modify settings** to update the settings.

Resetting the SSH fingerprint

You can reset the SSH fingerprint for a cluster that is managed by the Console for SAN Volume Controller for Cisco MDS 9000 for your configuration by using the Resetting the SSH Fingerprint panel.

Prerequisites:

You must have superuser administrator authority to perform the following procedure.

If you have changed the name of the console workstation, you must also change the console workstation host name in the IBM WebSphere Application Server files.

Context:

The communication between the Console for SAN Volume Controller for Cisco MDS 9000 and the cluster is through the Secure Shell (SSH) protocol. In this protocol, the Console for SAN Volume Controller for Cisco MDS 9000 acts as the SSH client

and the cluster acts as the SSH host server. The SSH protocol requires that credentials are exchanged when communication between the SSH client and server begins. The SSH client places the accepted SSH host server fingerprint in cache. Any change to the SSH server fingerprint in future exchanges results in a challenge to the end user to accept the new fingerprint. When a new code load is performed on the cluster, new SSH server keys can be produced which result in the SSH client flagging the SSH host fingerprint as changed and, therefore, no longer valid.

The Console for SAN Volume Controller for Cisco MDS 9000 displays the status of the cluster SSH server key in the **Availability Status** column of the Viewing Clusters panel.

Steps:

Perform the following steps to reset the SSH fingerprint:

1. Click **Clusters** in the portfolio. The View Clusters panel is displayed.
Attention: Select a cluster that has an availability status of Invalid SSH Fingerprint. In some cases this availability status results from a software upgrade that disrupts normal user operations. In the case of a disruptive software upgrade, follow the procedure for Recovering from a Disruptive Software Upgrade.
2. Select the cluster that you want to reset the SSH fingerprint for and select **Reset SSH Fingerprint** from the list. Click **Go**. The Resetting the SSH Fingerprint panel is displayed.
3. Select **OK** when prompted with the message, CMMVC3201W.

Result:

Availability status is changed to OK

Related topics:

- “Clusters” on page 12

Modifying Internet Protocol (IP) addresses

You can display and change the IP addresses associated with the cluster from the Modify IP Address panel.

Steps:

Perform the following steps to change the IP addresses:

1. Click **Manage Cluster** in the portfolio.
2. Click **Modify IP Address** from the Portfolio to check or change the IP address setting for the cluster. The **Modifying the Cluster IP Address** panel is displayed. You can use the Modifying the Cluster IP Address panel to change the cluster IP address. Fill in all four fields for the IP address that you want to change. Leave the IP address fields blank if you do not want to change them.
Click **OK** to perform the IP address update. When you specify a new cluster IP address, the existing communication with the cluster is broken. You must use the new cluster IP address to reestablish your browser connection.

Related topics:

- “Clusters” on page 12

Listing and saving log and dump files

You can list the various types of log and dump files that are available on the configuration node on the List Dumps panel. Dump data can be saved on any node in the cluster. When you use this procedure to display dump data only the dump files on the configuration node will be displayed. An option is provided on the dumps menu to display data from other nodes. If you choose to display or save data from another node that data will first be copied to the configuration node.

Steps:

Perform the following steps to list the various types of log and dump files:

1. Click **Service and Maintenance** from the portfolio.
2. Click **List Dumps** from the portfolio. The List Dumps panel displays. The List dumps (other nodes) continued panel displays the number of log files or dumps of a particular type that are available on the cluster. If there is more than one node in the cluster (as is usual), the **Check other nodes** button is displayed. If you click this button, the log files and dumps for all nodes that are part of the cluster is displayed. Dumps and logs on all nodes in the cluster can be deleted or copied to the node.

If you click on one of the file types, all the files of that type are listed in a table.

Note: For error logs and software dumps, the file names include the node name and time and date as part of the file name.

You can copy the files to a local workstation by right-clicking on the filename and using the **Save target as** (Netscape) or **Save file as** (Internet Explorer) option from the Web browser.

The file types that the **List dumps** option supports are:

- Error logs
- Configuration logs
- I/O statistic logs
- I/O trace logs
- Feature logs
- Software dumps

Follow the instructions in the right pane to display and save the dumps that you need

The software dump files contain dumps of the SAN Volume Controller Storage Software for Cisco MDS 9000 memory. Your service representative might ask for these dumps to debug problems. The software dumps are large files (approximately 300 MB). Consider copying these files to your host using secure copy methods.

Viewing the feature log

Steps:

Perform the following steps to view the feature log for the cluster:

1. Click **Service and Maintenance**.
2. Click **View Feature Log**. The View Feature Log panel is displayed.

Analyzing the error log

You can analyze the error log from the Analyze Error Log panel.

Steps:

Perform the following steps to analyze the error log:

1. Click **Service and Maintenance** in the portfolio.
2. Click **Analyze Error Log** in the portfolio. The Error log analysis panel is displayed. The Error log analysis panel enables you to analyze the cluster error log. You can display the whole log or filter the log so that only errors, events, or unfixed errors are displayed. In addition, you can request the table to be sorted either by error priority or by time. For error priority, the most serious errors are the lowest-numbered errors. They are, therefore, displayed first in the table. For time, either the oldest or the latest entry can be displayed first in the table. You can also select how many error log entries are to be displayed on each page of the table. The default is set to 10, and the maximum number of error logs that can be displayed on each page is 100.
3. After selecting the options, click **Process** to display the filtered error log in the table. The Analyze error log continued panel is displayed. Forward and Backward scroll buttons are displayed, depending on the existing page number and the total number of pages that are in the table. If the table contains more than two pages of entries, a **Go to** input area is displayed in the table footer. This input area enables you to skip to a particular page number.

If you click on the sequence number of a particular table record, more information about that error log entry is displayed. If the record is an error (instead of an event), you can change the fixed or unfixed status of the record; that is, you can mark an unfixed error as fixed or a fixed error an unfixed.
4. Click **Clear log** to erase the whole cluster error log.

Note: If you click **Clear log**, this will *not* fix the existing errors.

Shutting down a cluster or node

You can shut down a cluster from the Shutting Down cluster panel.

Prerequisites:

If all input power to a SAN Volume Controller Storage Software for Cisco MDS 9000 cluster is to be removed for more than a few minutes, (for example, if the machine room power is to be shutdown for maintenance), it is important that the cluster is shutdown before the power is removed.

Attention: Before shutting down a node or the cluster you should quiesce all I/O operations that are destined for this node or cluster. Failure to do so may result in failed I/O operations being reported to your host operating systems.

Attention: If you are shutting down the entire cluster, you will lose access to all VDIs being provided by this cluster.

Attention: Ensure that you have stopped all FlashCopy, Remote Copy and data migration operations before you attempt a node or cluster shutdown. Also ensure that all asynchronous deletion operations have completed prior to a shutdown operation.

Begin the process of quiescing all I/O to the cluster by stopping the applications on your hosts that are using the VDisks provided by the cluster.

1. If you are unsure which hosts are using the VDisks provided by the cluster, follow the procedure called, Determining the hosts that a VDisk is mapped to.
2. Repeat the previous step for all VDisks.

Context:

If all input power to a SAN Volume Controller Storage Software for Cisco MDS 9000 cluster is to be removed, for example, if the machine room power is to be shutdown for maintenance, you must shut down the cluster before the power is removed. If you do not shut down the cluster before turning off input power to the internal battery backup, the SAN Volume Controller Storage Software for Cisco MDS 9000s will detect the loss of power and continue to run on battery power until all data held in memory is saved to the internal disk drive. This will increase the time required to make the cluster operational when input power is restored and will severely increase the time required to recover from an unexpected loss of power that might occur before the internal battery backup batteries have fully recharged.

A cluster or node can be shutdown by stopping I/O activity and issuing a shutdown command to the cluster.

Steps:

Perform the following steps to shut down a cluster:

1. Click **Manage Clusters** in the portfolio.
2. Click **Shut down Clusters** in the portfolio. The Shutting Down Clusters panel is displayed.

Related topics:

- “Clusters” on page 12

Part 4. Installing or upgrading the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 for Windows

Chapter 11. Installation overview for the Console for SAN Volume Controller for Cisco MDS 9000

You should have some knowledge of how to administer a Windows 2000 Server operating system before you install the Console for SAN Volume Controller for Cisco MDS 9000. You should also become familiar with the command that you use during the installation or upgrade installation of the Console for SAN Volume Controller for Cisco MDS 9000.

You must be aware of the following list of installation and configuration tasks *before* you install or upgrade the Console for SAN Volume Controller for Cisco MDS 9000:

1. Check the hardware and software requirements.
2. If the SSH client software called PuTTY is not yet installed on your system, you must install the SSH client software. For your convenience the PuTTY installation program (putty-o.53b-installer.exe) is on the Console for SAN Volume Controller for Cisco MDS 9000 installation CD-ROM in the SSHClient/PuTTY directory. You can get more information about PuTTY from the PuTTY Web site home page:

<http://www.chiark.greenend.org.uk/~sgtatham/putty/>

3. Install or upgrade the Console for SAN Volume Controller for Cisco MDS 9000 either in graphical mode with the help of an installation wizard or in unattended mode (also known as silent mode), which involves customizing a response file and issuing a command.

Software package:

Cluster software versions comprise a number of software components that are delivered as a single package. The software installation procedure involves copying the new software version to the cluster and then starting an automatic installation process. This installation process might take up to an hour to complete and during the process each of the nodes is restarted in turn. Once all the nodes in the cluster have been successfully restarted with the new software the new software version is automatically committed. While each node is being restarted there might be some degradation in the maximum input/output rate that can be sustained by the cluster.

Installation or upgrade installation operation:

The installation or upgrade installation operation can normally be performed concurrently with normal user I/O operations. If any restrictions apply to the operations that can be performed during the upgrade, then these restrictions will be documented on the SAN Volume Controller Web site from where the upgrade package was obtained. During the upgrade operation, only the SAN Volume Controller commands will be operational from the time the installation process starts to the time that the new software is committed or until the process has been backed-out. For a complete list of SAN Volume Controller commands, enter the following command:

```
svcinfo -?
```

Because of the operational limitations that occur during the software upgrade process, the software installation is a customer task.

4. Verify the following Windows services associated with the Console for SAN Volume Controller for Cisco MDS 9000 are installed and started:
 - Service Location Protocol
 - IBM CIM Object Manager - SVCCISCO

- IBM Websphere Application Server V5 - SVCCISCO
5. Get started using the Console for SAN Volume Controller for Cisco MDS 9000. Use a Web browser to access the Console for SAN Volume Controller for Cisco MDS 9000. You will identify the clusters to be managed to the Console for SAN Volume Controller for Cisco MDS 9000 as well as complete the creation (initialization) of the SAN Volume Controller Storage Software for Cisco MDS 9000 clusters.

To allow nodes to operate as a cluster, you must run all nodes at the same version of software. This rule is enforced by the cluster software itself. When you attempt to add a node to a cluster its software version is examined, and if it is not running the same version of the software as the other nodes in the cluster, the software revisions are automatically copied from one of the other nodes in the cluster before the add operation is completed. If for some reason it is not possible to update the software on the node that you are adding, the operation fails and the cluster logs an error to explain the cause of the failure.

6. Remove the Console for SAN Volume Controller for Cisco MDS 9000. You only need to perform this optional task if you get errors during installation verification.

Related topics:

- Chapter 12, “Console for SAN Volume Controller for Cisco MDS 9000 hardware installation requirements,” on page 117
- Chapter 13, “Console for SAN Volume Controller for Cisco MDS 9000 workstation space requirements,” on page 119
- Chapter 14, “Console for SAN Volume Controller for Cisco MDS 9000 software installation requirements,” on page 121
- Chapter 15, “Installing or upgrading the Console for SAN Volume Controller for Cisco MDS 9000 in graphical mode,” on page 123
- Chapter 16, “Installing or upgrading the Console for SAN Volume Controller for Cisco MDS 9000 in unattended (silent) mode,” on page 129
- Chapter 17, “Verifying the Windows services associated with the Console for SAN Volume Controller for Cisco MDS 9000,” on page 135
- Chapter 18, “Post installation tasks,” on page 137
- Chapter 19, “Removing the Console for SAN Volume Controller for Cisco MDS 9000,” on page 141

Chapter 12. Console for SAN Volume Controller for Cisco MDS 9000 hardware installation requirements

Before starting the installation, ensure that your system satisfies the following hardware installation prerequisites for installing the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 on a Windows 2000 Server operating system.

Hardware prerequisites:

The following hardware is required:

- Any Intel®-based PC running Windows 2000 Server SP 3 or SP4
- Intel Pentium® processor at 1 GHz, or faster
- Support for a communications adapter
- CD-ROM drive
- Minimum 1 GB RAM recommended

Related topics:

- Chapter 13, “Console for SAN Volume Controller for Cisco MDS 9000 workstation space requirements,” on page 119
- Chapter 14, “Console for SAN Volume Controller for Cisco MDS 9000 software installation requirements,” on page 121

Chapter 13. Console for SAN Volume Controller for Cisco MDS 9000 workstation space requirements

Before starting the installation, ensure that your system satisfies the following workstation space prerequisites for installing the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 on a Windows 2000 Server operating system.

Workstation space:

The following space on your workstation is required:

- 350 MB of disk space

Note: You might need to increase the total available disk space on your hard drives if the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 and other associated products are split between more than one logical drive. Also, the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 might require additional memory to operate if you configure it to manage many devices or devices with large configurations.

- Up to 65 MB of temporary disk space for installation purposes

Related topics:

- Chapter 12, “Console for SAN Volume Controller for Cisco MDS 9000 hardware installation requirements,” on page 117
- Chapter 14, “Console for SAN Volume Controller for Cisco MDS 9000 software installation requirements,” on page 121

Chapter 14. Console for SAN Volume Controller for Cisco MDS 9000 software installation requirements

Before starting the installation, ensure that your system satisfies the following software installation prerequisites for installing the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 on a Windows 2000 Server operating system.

Software:

The following software is required:

- Operating systems:
 - Windows 2000 Server SP3 or SP4
- If the SSH client software called PuTTY is not yet installed on your system, you must install the SSH client software. You can get more information about PuTTY from the PuTTY Web site home page:

<http://www.chiark.greenend.org.uk/~sgtatham/putty/>

and download PuTTY from the following Web site download page:

<http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>

For your convenience the PuTTY installation program (putty-o.53b-installer.exe) is on the Console for SAN Volume Controller for Cisco MDS 9000 installation CD-ROM in the SSHClient/PuTTY directory.

- IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000. This is on the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 CD.
- Transmission Control Protocol/Internet Protocol (TCP/IP)
- Adobe Acrobat Reader version 4.0 or later (optional)

You need the Adobe Acrobat Reader to read License Agreement and product information from the Console for SAN Volume Controller for Cisco MDS 9000 LaunchPad. You can download the Adobe Acrobat Reader from the following Web site:

– <http://www.adobe.com/support/downloads/main.html>

Related topics:

- Chapter 12, “Console for SAN Volume Controller for Cisco MDS 9000 hardware installation requirements,” on page 117
- Chapter 13, “Console for SAN Volume Controller for Cisco MDS 9000 workstation space requirements,” on page 119

Chapter 15. Installing or upgrading the Console for SAN Volume Controller for Cisco MDS 9000 in graphical mode

This topic includes the steps to install or upgrade the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 on your Windows system.

If you choose to install or upgrade the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 in unattended mode skip this section. You must satisfy all prerequisites before starting the installation.

Steps:

Perform the following steps to install or upgrade the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000:

1. Log onto your system as a local system administrator.
2. Insert the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 CD into the CD drive.

The IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 program should start within 15 - 30 seconds if you have **autorun** mode set on your system. If the LaunchPad panel does not open, perform either one of the following steps:

- a. Use a Command Prompt to change to the W2K directory on the CD. Type:
LaunchPad
- b. Using Windows Explorer, (**Start->Programs->Accessories->Windows Explorer**), go to the W2K directory located on the CD drive. Then double-click on the **LaunchPad.bat** file.

Note: If you are viewing the folder using the Windows Explorer with the option selected to *Hide file extensions for known file types*, find the LaunchPad file with the file type of MS-DOS Batch File.

3. The following options are displayed when the LaunchPad panel opens:

Product overview	Offers information about the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000.
Readme file	Offers any last minute product information that did not make it into sections concerning the installation of the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000.
Configuration guide	Contains instructions about how to install the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 (a softcopy of this document).
License agreement	Offers information about the license for the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000.
Product Web site	Offers information from the product Web site.
Installation wizard	Starts the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 installation program.

Post installation tasks	Details information about validating the installation, accessing the Console for SAN Volume Controller for Cisco MDS 9000 URL and adding the Console for SAN Volume Controller for Cisco MDS 9000 cluster to the Console for SAN Volume Controller for Cisco MDS 9000 management facility.
Exit	Exits the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 LaunchPad program.

4. Click **Readme file** from the LaunchPad panel or from the **README.txt** file located in the doc or W2K directory on the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 CD to check for information that might supersede the information in this guide.
5. Click **Installation wizard** from the LaunchPad panel to start the installation.

Note: The LaunchPad panel remains open behind the installation wizard so that you can access product information during the installation process. Click **Exit** if you want to close the LaunchPad.

6. There might be a slight delay while the software loads on your system. After the software loads a DOS prompt window opens to display the following message:

```

Initializing InstallShield Wizard...
Preparing Java <tm> Virtual Machine .....
.....
.....
```

7. The Welcome panel opens suggesting what documentation you should review prior to installation. Click **Next** to continue, or click **Cancel** to exit the installation.
8. The License Agreement panel opens. Read the license agreement information. Select **I accept the terms of the license agreement**, then click **Next** to accept the license agreement. Otherwise, keep the selection **I do not accept the terms of the license agreement** (it is the default) and click **Cancel** to exit the installation.
9. The installation wizard verifies that your workstation meets the installation requirements.
 - If you have a Service Location Protocol (SLP) service that is different from the SLP that the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 requires, the installation wizard displays an error and asks you to stop the installation and remove this SLP service from the system.
 - The installation wizard checks if the PuTTY SSH client is installed on your workstation.
 - The installation wizard determines whether this is a new installation, reinstallation or upgrade installation of the SAN Volume Controller Console. If the installation wizard determines that the Console for SAN Volume Controller for Cisco MDS 9000 was previously installed on the system, it does a comparison of the current version, release, modification, and fix code level with that of the code currently installed on the system. If the level is the same, this is a reinstallation. If the new code has a higher level, it is an upgrade. If the new code level is lower than the level on the system, the installation is invalid. In the case of reinstallation or upgrade installation, the installation wizard performs the following actions:

- a. Checks if the Service Location Protocol (SLP), the IBM CIM Object Manager - SVCCISCO service, and IBM Websphere Application Server V5 - SVCCISCO are started. If any of these services are started, the program asks if you want to continue the installation process by clicking **Next**. If you want to exit the installation program click **Cancel**. If you choose to continue, you must stop all the applications that use these services.
 - b. Presents a panel with the check-box option to Preserve Configuration. If you chose to preserve the existent Configuration, the installation program skips the next steps and goes directly to the Installation Confirmation panel discussed below.
10. The Destination Directory panel opens. Select one of the following options:

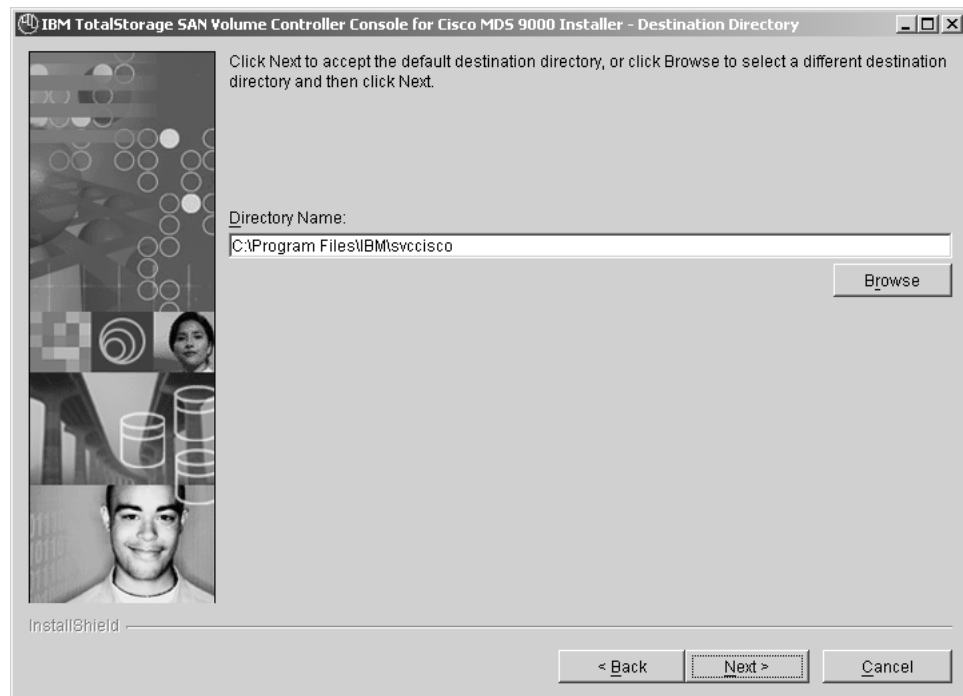


Figure 22. Destination Directory panel

- a. Click **Next** to accept the default directory.
- b. Click **Browse** to select a different directory for installation and then click **Next** to continue the installation process.
- c. Click **Cancel** to exit the installation process.

Notes:

- a. The directory name, including the drive letter, must be a maximum of 44 characters.
- b. If the program detects insufficient space for the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 installation in the chosen destination, an error message is displayed. You can free some space on the destination drive and then click **Next** or you can stop the installation program by clicking **Cancel**. You can also go back by clicking **Back**, and choosing another destination directory for the product.
- c. If the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 is already installed, the Destination Directory panel is not displayed

by the installation program. The IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 is reinstalled or upgraded in the same installation location.

11. The PuTTY configuration panel opens when the product space check is completed.
12. On the Updating CIMOM Ports panel, accept the defaults or update the port assignments for the IBM CIM Object Manager - SVCCISCO Port, IBM CIM Object Manager SVCCISCO Indication Port, and the Communication Protocol. To check ports in use, use the **netstat -a** command. Click **Next** to continue.
13. The Updating Embedded WAS Ports panel is displayed.

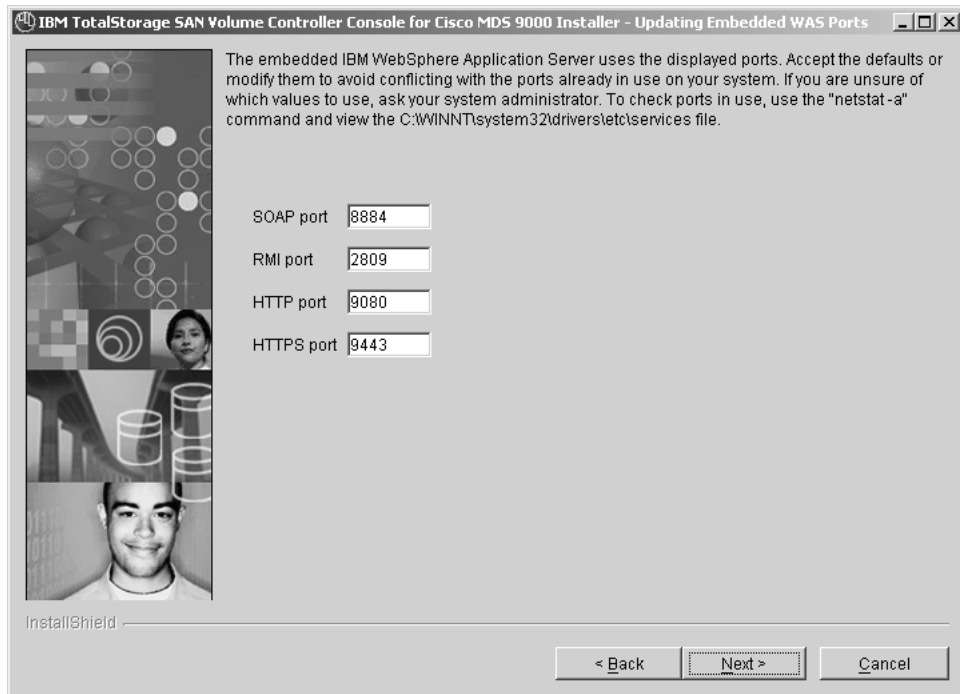


Figure 23. Updating Embedded WAS Ports panel

Update the default ports assignments by typing unique port numbers for the products that have been registered on your system. To check ports in use, use the **netstat -a** command and view the C:\WINNT\system32\drivers\etc\services file. Click **Next** to continue.

14. The Updating CIMOM ports panel is displayed. Update the default port assignments and the default communication protocol by typing the unique port numbers and choosing the desired communication protocol for the products that have been registered on your system. To check ports in use, use the **netstat -a** command and view the C:\WINNT\system32\drivers\etc\services file. Click **Next** to continue.
15. The Installation Confirmation panel opens. Click **Install** to confirm the installation location and file size and to start the final installation, reinstallation or upgrade installation. Click **Cancel** to exit the installation wizard or click **Back** to go to the previous panel.
16. The Installation Progress panel opens indicating how much of the installation has been completed. Installation usually takes 3 - 10 minutes depending on the configuration of your workstation.

Note: If you click **Cancel** a popup panel opens asking you to confirm the cancellation of the installation wizard: Cancel the current operation?. You may confirm the cancellation by clicking **Yes** or continue the installation by selecting **No**. If you confirm the cancellation, the information you entered or selected in the previous panel is not saved. You must start the installation again from the beginning.

After the completion of the successful installation of the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000, the installer attempts to start the following services:

- Service Location Protocol
- IBM CIM Object Manager - SVCCISCO
- IBM Websphere Application Server V5 - SVCCISCO

17. When the Installation Progress panel closes, the Finish panel opens. Before proceeding, you might want to review the log file for any possible error messages. The log file is located in xxx\logs\install.log, where xxx is the destination directory where the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 for Windows was installed. The install.log contains a trace of the installation actions.

Note: At the bottom of the Finish panel is a checkbox labeled **View post installation tasks**. If you check this box and then click **Finish**, the wizard will exit and the post installation tasks text file is displayed. The LaunchPad panel Post installation tasks link also displays this same text file. You can avoid the display of the text file by unchecking the **View post installation tasks** box before you click the **Finish** button.

18. Click **Finish** to exit the installation wizard.

Note: Ordinarily, you do not need to restart your system during or after the installation of the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000. However, the installation wizard might determine that a restart is necessary. Restart your system if required. After you restart the system, the installation wizard continues with the installation.

19. If you have not yet reviewed the post-installation tasks from the installation Finish panel, review the post installation tasks from the LaunchPad program.
 - a. Click **Post installation tasks** on the LaunchPad panel which opens the same file available from the Installation Finish panel.
 - b. Continue with the post-installation tasks for the SAN Volume Controller Storage Software for Cisco MDS 9000 by following the instructions in this file.
20. Exit the LaunchPad program by clicking **Exit** on the LaunchPad panel.
21. Verify that the Windows services associated with your Console for SAN Volume Controller for Cisco MDS 9000 are correctly installed and started.

Related topics:

- Chapter 11, "Installation overview for the Console for SAN Volume Controller for Cisco MDS 9000," on page 115
- Part 4, "Installing or upgrading the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 for Windows," on page 113
- Chapter 17, "Verifying the Windows services associated with the Console for SAN Volume Controller for Cisco MDS 9000," on page 135
- Chapter 16, "Installing or upgrading the Console for SAN Volume Controller for Cisco MDS 9000 in unattended (silent) mode," on page 129

- Chapter 18, “Post installation tasks,” on page 137

Chapter 16. Installing or upgrading the Console for SAN Volume Controller for Cisco MDS 9000 in unattended (silent) mode

This topic provides information about installing or upgrading the Console for SAN Volume Controller for Cisco MDS 9000 in unattended (silent) mode.

The unattended (silent) mode install or upgrade option enables you to run the installation or upgrade installation unattended. Use this method of installation to customize a response file and issue a command from a command prompt window. The response file is a template on the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 CD. You can also create a standard response file to ensure that the product is installed consistently on multiple systems. You must satisfy all prerequisites before starting the installation.

The installation wizard determines whether this is a reinstallation or upgrade of the Console for SAN Volume Controller for Cisco MDS 9000. If the installation wizard determines that the Console for SAN Volume Controller for Cisco MDS 9000 was previously installed on the system, it does a comparison of the current version, release, modification, and fix code level with that of the code currently installed on the system. If the level is the same, this is a reinstallation. If the new code has a higher level, it is an upgrade. If the new code level is lower than the level on the system, the installation is invalid.

Steps:

Perform the following steps to install or upgrade the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 in your Windows environment using the unattended mode:

1. Log on to the system as a local system administrator.
2. Insert the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 CD into the CD drive.
3. If you have autorun mode set on your system, the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 program will start within 15-30 seconds. Click **Exit** from the LaunchPad.
4. Locate the response file named, responsefile, on your IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 CD in the W2K directory.
5. Using Windows Explore or a command prompt, copy the response file to your hard drive.
6. The Console for SAN Volume Controller for Cisco MDS 9000 will connect to the SAN Volume Controller Storage Software for Cisco MDS 9000 using the Secure Shell (SSH) protocol. You need to supply a file containing the SSH2 RSA private key for the SSH client connection. This file will be copied by the installation program to <inst_dir>\cimom\icat.ppk, for example C:\ProgramFiles\IBM\svconconsole\cimom\icat.ppk. If you have not generated the private key file before, you can do it now using the PuTTYgen utility. To generate the private key using the PuTTYgen utility, perform the following steps:
 - a. Click **Start -> Programs -> PuTTY -> PuTTYgen**.
 - b. Select the type of key: **SSH RSA**.
 - c. Click **Generate**.

- d. Move the cursor around the blank area of the Key section to generate a random number.
 - e. Click **Save public key**. You will provide this key to the Console for SAN Volume Controller for Cisco MDS 9000 later.
 - f. Click **Save private key**. You will provide this key to the installation wizard using the option below in the response file.
 - g. Click **Yes** on the PuTTYgen warning. You do not need a passphrase.
 - h. Ensure that you set the value of the `<-W puttyConfiguration.puttyPrivateKeyFile>` option in the response file, to the name of the file containing the PuTTY private key.
7. Using a text editor modify the default options in the response file with the values you want to supply to the installation program:
- Remove the # character from the beginning of a line if you do not want to use the default value. Change the default value to the value that you want for that option. You *must* enclose all values in double quotation marks ("").
 - Depending on whether you are doing a new installation, reinstallation or upgrade, certain response file lines must be active as detailed below. If a response file line is active but inappropriate to the mode (new, reinstall or upgrade) it will be ignored.

New Installation:

- The `<-P product.installLocation>` option defines the default directory where the product is to be installed. To specify a destination directory other than the default, remove the # character from the corresponding line and replace the default directory with the desired directory.
- The `<-G checkPrerequisite>` option checks the prerequisites. If you want to disable this option, remove the # character from the corresponding line and change the value of the option to no.
- Change the default ports values for the embedded IBM Websphere Application Server V5 - SVCCISCO using the update ports variables options. If you want to change a specific port used for a particular WebSphere service, remove the # character from the beginning of the line containing the value of the option and set it to the value you desire. The following are the embedded WebSphere ports options:
 - `<-W ports.portSOAP="8884">`
 - `<-W ports.portRMI="2809">`
 - `<-W ports.portHTTP="9080">`
 - `<-W ports.portHTTPS="9443">`
- Change the default ports values and default server communication type for the IBM CIM Object Manager server using variables options below. If you want to change a specific port or the default server communication type, remove the # character from the beginning of the line containing the option's value and set it to the value you desire. The following are the CIM Object Manager server options:
 - `<-W cimObjectManagerPorts.port="5989">`
 - `<-W cimObjectManagerPorts.indicationPort="5990">`
 - `<-W cimObjectManagerPorts.serverCommunication="HTTPS">`

Reinstallation or Upgrade:

- The `<-G startUpgrade>` option must be enabled to permit the new Console for SAN Volume Controller for Cisco MDS 9000 to be reinstalled (having the

same version) or upgraded (installed at a higher version). Enable this option by removing the # character from the corresponding line and changing the value of the option to yes.

- The `<-G stopProcessesResponse>` option tells the install program whether or not to automatically stop SLP, IBM CIM Object Manager (CIMOM), and WebSphere Application Server - V5 SAN Volume Controller Storage Software for Cisco MDS 9000 services when reinstalling or upgrading the product. By default, this option is set to no. If you do not change this default value, the reinstallation or upgrade stops when these services are running. If you want to automatically stop the SLP and IBM CIM Object Manager (CIMOM), remove the # character from the corresponding line and change its value to yes.
 - The `<-G saveConfiguration>` option specifies whether or not to save the configuration files when reinstalling or upgrading the product. If you do not want to save the configuration files when reinstalling or upgrading, remove the # character from the corresponding line and change the value of the option to no. If you do not choose to save the configuration, you will have to make the following active or accept the default values.
 - Change the default ports values for the embedded WebSphere Application Server - V5 SAN Volume Controller Storage Software for Cisco MDS 9000 using the update ports variables options. If you want to change a specific port used for a particular WebSphere service, remove the # character from the beginning of the line containing the value of the option and set it to the value you desire. The following are the embedded WebSphere ports options:
 - `<-W ports.portSOAP="8884">`
 - `<-W ports.portRMI="2809">`
 - `<-W ports.portHTTP="9080">`
 - `<-W ports.portHTTPS="9443">`
 - Change the default ports values and the default server communication type for the CIM Object Manager server using variables options below. If you want to change a specific port or the default server communication type, remove the # character from the beginning of the line containing the option's value and set it to the value you desire. The following are the CIM Object Manager server options:
 - `<-W cimObjectManagerPorts.port="5989">`
 - `<-W cimObjectManagerPorts.indicationPort="5990">`
 - `<-W cimObjectManagerPorts.serverCommunication="HTTPS">`
 - The `<-W puttyConfiguration.puttyPrivateKeyFile>` options specifies the name and location of the PuTTY private key file that the SAN Volume Controller Console software should use to connect to the SAN Volume Controller clusters. Remove the # character from the corresponding line and add the fully qualified location of the PuTTY private key file. Save the response file without a file extension such as .txt.
8. From a command prompt window, type the following command:
- ```
<CD drive path>\W2K\install -options <response file path>\responsefile
```

where `<CD drive path>` is the path of your CD drive. `<response file path>` is the path of the responsefile file that you copied in step 5 on page 129 and customized in step 7 on page 130.

**Note:** The directory name, including the drive letter, must be a maximum of 44 characters.

9. During the installation, dotted lines are displayed across the screen. When the installation program ends, control returns to the Command Prompt.
10. Check for installation errors in the install.log file. This file is initially created in the system temporary file under the subdirectory named, svccisco. After all the prerequisites checks have been performed, the log file is copied to the <dest-path>\logs directory. The following is an example of an install.log file:

```
(Oct 23, 2003 9:35:11 AM), This summary log is an overview of the sequence
of the installation of the IBM TotalStorage Console for SAN Volume Controller
for Cisco MDS 9000 1.0.0.0
(Oct 23, 2003 9:46:21 AM), IBM TotalStorage Console for SAN Volume Controller
for Cisco MDS 9000 installation process started with the following install
parameters:
Target Directory: "C:\Program Files\IBM\svccisco"
SOAP port: 8886
RMI port: 2829
HTTP port: 9280
HTTPS port: 9643
IBM CIM Object Manager - SVCCISCO Port: 5789
IBM CIM Object Manager - SVCCISCO Indication Port: 5790
Communication Protocol: HTTPS
(Oct 23, 2003 9:46:31 AM), Copying Service Location Protocol Files ...
(Oct 23, 2003 9:46:34 AM), Copying CIM Object Manager Files ...
(Oct 23, 2003 9:50:53 AM), IBM CIM Object Manager ports successfully updated.
(Oct 23, 2003 9:50:53 AM), IBM CIM Object Manager communication protocol successfully updated.
(Oct 23, 2003 9:50:53 AM), The file setupCmdLine.bat successfully updated.
(Oct 23, 2003 9:50:53 AM), Compile MOF files started ...
(Oct 23, 2003 9:51:51 AM), MOF files successfully compiled.
(Oct 23, 2003 9:51:52 AM), Generate a certificate store started ...
(Oct 23, 2003 9:52:06 AM), Certificate store called truststore successfully generated.
(Oct 23, 2003 9:52:06 AM), IBM CIM Object Manager - SVCCISCO successfully installed
(Oct 23, 2003 9:52:07 AM), Installing embedded version of IBM WebSphere Application Server ...
(Oct 23, 2003 9:54:21 AM), IBM WebSphere Application Server V5 - SVCCISCO service successfully installed.
(Oct 23, 2003 9:54:21 AM), Copying Eclipse Files...
(Oct 23, 2003 9:58:49 AM), Copying Console for SAN Volume Controller for Cisco MDS 9000 Help Files...
(Oct 23, 2003 9:58:56 AM), Copying Console for SAN Volume Controller for Cisco MDS 9000 Ear Files...
(Oct 23, 2003 10:05:15 AM), The ICAConsole application successfully installed.
(Oct 23, 2003 10:08:09 AM), The SVCConsole application successfully installed.
(Oct 23, 2003 10:09:25 AM), The help application successfully installed.
(Oct 23, 2003 10:10:14 AM), The "C:\Program Files\IBM\svccisco\console\
embeddedWAS\bin\expressPorts\UpdateExpressMultiPorts.bat" -soap 8886
-boot 2829 -http 9280 -https 9643 -remove" command updated successfully
embedded WAS ports in configuration files.
(Oct 23, 2003 10:10:15 AM), Command to be executed : net start cimomsvccisco
(Oct 23, 2003 10:10:43 AM), Command to be executed : net start "IBMWAS5Service - SVCCISCO"
(Oct 23, 2003 10:12:45 AM), The following services started successfully:
IBM CIM Object Manager - SVCCISCO
IBM WebSphere Application Server V5 - SVCCISCO
(Oct 23, 2003 10:12:45 AM), INSTSUCC: The IBM TotalStorage Console for SAN Volume
Controller for Cisco MDS 9000 has been successfully installed.
```

11. Close the command prompt window by entering a command, for example **exit**.
12. After the completion of the successful installation of the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000, the installer attempts to start the following services:
  - Service Location Protocol
  - IBM CIM Object Manager - SVCCISCO
  - IBM Websphere Application Server V5 - SVCCISCO
13. Continue with the post installation tasks for the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 using the instructions in the following section. You can also view the post installation tasks using the following option:

- a. From a Command Prompt, change directory into the W2K directory on the CD drive. Open the LaunchPad by typing:

LaunchPad

- b. Click **Post installation tasks** on the LaunchPad window. Continue with the post installation tasks for the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 by following the instructions in this file.
14. Verify that the Windows services associated with your Console for SAN Volume Controller for Cisco MDS 9000 are correctly installed and started.

**Related topics:**

- Chapter 11, “Installation overview for the Console for SAN Volume Controller for Cisco MDS 9000,” on page 115
- Chapter 17, “Verifying the Windows services associated with the Console for SAN Volume Controller for Cisco MDS 9000,” on page 135





---

## Chapter 17. Verifying the Windows services associated with the Console for SAN Volume Controller for Cisco MDS 9000

This task verifies that the Windows services associated with your IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 are correctly installed and started.

### Steps:

Perform the following steps to verify your Service Location Protocol (SLP), IBM CIM Object Manager - SVCCISCO, and IBM Websphere Application Server V5 - SVCCISCO services were correctly installed:

1. Verify the installation of the Service Location Protocol (SLP).
  - a. Verify that the Service Location Protocol is started. Select **Start -> Settings -> Control Panel**. Double-click the **Administrative Tools** icon. Double-click the **Services** icon.
  - b. Find **Service Location Protocol** in the **Services** list. For this component, the **Status** column should be marked Started.
  - c. If the Service Location Protocol is not started, right-click on **Service Location Protocol** and select **Start** from the pop-up menu. Wait for the **Status** column to change to Started.
  - d. Do not close the Services window because you will also use it to verify the The CIM Object Manager (CIMOM) service.
2. Verify the installation of the Console for SAN Volume Controller for Cisco MDS 9000.
  - a. Find the **IBM CIM Object Manager - SVCCISCO** in the **Services** list. For this component, the **Status** column should be marked Started.
  - b. If the IBM CIM Object Manager is not started, right click on **IBM CIM Object Manager - SVCCISCO** and select **Start** from the pop-up menu. Wait for the **Status** column to change to Started.
  - c. Do not close the Services window because you will also use it to verify the IBM Websphere Application Server V5 - SVCCISCO service.
3. Verify the installation of the IBM Websphere Application Server V5 - SVCCISCO service.
  - a. Find the **IBM Websphere Application Server V5 - SVCCISCO** in the **Services** list. For this component, the **Status** column should be marked Started.
  - b. If the **IBM Websphere Application Server V5 - SVCCISCO** service is not started, right click on **IBM Websphere Application Server V5 - SVCCISCO** and select **Start** from the pop-up menu. Wait for the **Status** column to change to Started.
  - c. Close the Services window.
  - d. Close the Administrative Tools window.



---

## Chapter 18. Post installation tasks

This section outlines how to get started using the Console for SAN Volume Controller for Cisco MDS 9000 using your Web browser.

Once you have installed the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 and the services (IBM CIM Object Manager - SVCCISCO, IBM Websphere Application Server V5 - SVCCISCO, Service Location Protocol) have started, you will use a browser to access the Web pages of the Console for purposes of administering the Console for SAN Volume Controller for Cisco MDS 9000 as well as configuring SAN Volume Controller Storage Software for Cisco MDS 9000 clusters.

You must enable SSH when using the CIM Agent and Console for SAN Volume Controller for Cisco MDS 9000. To enable SSH on the SAN Volume Controller Storage Software for Cisco MDS 9000, enter the following commands at the switch prompt:

To enter config mode:

```
switch#config-t
```

To enable ssh:

```
switch (config)#ssh server enable
```

### Steps:

1. Start your Web browser to access the Console for SAN Volume Controller for Cisco MDS 9000. You can access the Console for SAN Volume Controller for Cisco MDS 9000 by typing the following:

```
http://localhost:9080/ica
```

**Note:** 9080 is the default HTTP port. If a different port number for HTTP was assigned during the installation process, then you must substitute that port number in the URL.

2. Log onto the Console for SAN Volume Controller for Cisco MDS 9000 using the default super user name and password. The default super user name is superuser and the default super user password is passwd. The first time you log onto the Console for SAN Volume Controller for Cisco MDS 9000 using the default super user name and password, you will be prompted to change the default password.
3. Access user assistance. This is an optional step.  
You can access help for the specific task on which you are working by clicking the small information icon just below the banner in the upper right section of the Web page. The help assistant panel opens on the right-hand side of the page.  
You can also launch a separate user assistance panel by clicking the small question mark icon just below the banner in the upper right section of the Web page. A secondary browser window opens which has icons in the frame labeled **Contents** for you to select to make extensive user assistance information available to you.
4. Identify the SAN Volume Controller Storage Software for Cisco MDS 9000 clusters to the Console for SAN Volume Controller for Cisco MDS 9000. For prerequisites to this step, refer to *IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000: Getting Started Guide*. The steps you might need to perform to add SAN Volume Controller Storage Software for

Cisco MDS 9000 clusters to the Console for SAN Volume Controller for Cisco MDS 9000 collection of managed clusters, depends on the current status of the cluster in which you are interested.

**Note:** If you have never created a SAN Volume Controller Storage Software for Cisco MDS 9000 cluster, you will see the following two buttons on the Welcome panel after you log on:

- a. Add SAN Volume Controller Cluster
- b. Create new SAN Volume Controller Cluster

These buttons are helpful first time usage aids to bring you to the Viewing Clusters panel. You can always access the Viewing Clusters panel if you click **Clusters** in the portfolio frame on the left side of the panel, as described in the steps below. If you have added at least one cluster to the Console for SAN Volume Controller for Cisco MDS 9000, you will not see these buttons again.

Choose one of the following two steps, depending on whether the cluster has completed the cluster creation process:

- a. Previously initialized SAN Volume Controller Storage Software for Cisco MDS 9000 cluster but not yet registered with the Console for SAN Volume Controller for Cisco MDS 9000.

Click **Clusters** in the portfolio frame at the left side of the panel on the Console for SAN Volume Controller for Cisco MDS 9000 Welcome panel. The Viewing Clusters panel is displayed. Select **Add a cluster** from the drop-down list and click **Go**. The Adding a Cluster wizard is displayed. Complete the wizard to add the cluster.

- b. Uninitialized SAN Volume Controller Storage Software for Cisco MDS 9000 cluster.

Click **Clusters** in the portfolio frame on the left side of the Console for SAN Volume Controller for Cisco MDS 9000 Welcome panel. The Viewing Clusters panel is displayed.

Select **Create a cluster** from the drop down list and click **Go**. The Creating a Cluster wizard is displayed. Complete the wizard to create the cluster.

The Creating a Cluster wizard internally adds the SAN Volume Controller Storage Software for Cisco MDS 9000 cluster to the Console for SAN Volume Controller for Cisco MDS 9000 so you do not need to perform an extra step to add the cluster.

5. Launch the secondary Web browser window to manage your specific cluster.

Once you have identified the SAN Volume Controller Storage Software for Cisco MDS 9000 clusters to the Console for SAN Volume Controller for Cisco MDS 9000 you can see a summary of all of the clusters that you added to the Console for SAN Volume Controller for Cisco MDS 9000. From this point, you can select the specific cluster in which you are interested and then launch the browser window specifically for the cluster. Perform the following steps to launch the browser window:

- a. Click **Clusters** in the portfolio section of your browser window in the left-hand frame. The Viewing Clusters panel is displayed in the work area.
- b. Check the small box in the Select column left of the cluster in which you are interested to select that cluster. Select **Launch the SAN Volume Controller application** in the drop down list of the work area and click **Go**. A secondary browser window opens to the SAN Volume Controller Storage

Software for Cisco MDS 9000 Web application. Now you can work with the specific SAN Volume Controller Storage Software for Cisco MDS 9000 cluster which you selected.

**Note:** The ClusterName parameter in the browser location URL, identifies the cluster with which you are working.

For example:

```
http://localhost:9080/svc/Console?Console.login
Token=79334064:f46d035f31:-7ff1&Console.
ClusterName=9.11.111.208
```

Select **Manage Cluster** from the portfolio frame and click **View Cluster Properties**.

**Result:**

This completes the verification of the connection to the SAN Volume Controller Storage Software for Cisco MDS 9000.

**Related topics:**

- Chapter 7, “Overview of creating a cluster using the Console for SAN Volume Controller for Cisco MDS 9000,” on page 63



---

## Chapter 19. Removing the Console for SAN Volume Controller for Cisco MDS 9000

This optional task provides the instructions for removing the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 from your Windows system.

### Steps:

Perform the following steps to remove the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000:

1. Log onto the system as a local system administrator.
2. Stop the IBM CIM Object Manager - SVCCISCO, IBM Websphere Application Server V5 - SVCCISCO, and the Service Location Protocol (SLP) services if they are started.
  - a. Click **Start -> Settings -> Control Panel**. In the Control Panel window, double-click on the **Administrative Tools** icon and then double-click the **Services** icon. The Services window opens.
  - b. Stop the IBM CIM Object Manager - SVCCISCO service:
    - 1) In the Services window, scroll to IBM CIM Object Manager - SVCCISCO. Click on the service to select it.
    - 2) If the **Status** column shows Started, right-click the service, then click **Stop** on the menu.
  - c. Stop the IBM Websphere Application Server V5 - SVCCISCO service:
    - 1) In the Services window, scroll to IBM Websphere Application Server V5 - SVCCISCO. Click on the service to select it.
    - 2) If the **Status** column shows Started, right-click the service, then click **Stop** on the menu.
    - 3) Wait for the service to stop.
  - d. Stop the Service Location Protocol (SLP) service:

**Note:** You must be careful if you have other applications that use the Service Location Protocol (SLP) service. In this case, you must stop these applications before stopping Service Location Protocol (SLP) service, because during the removal process the Service Location Protocol (SLP) service will be deleted. You must also stop the configuration utilities for the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000, if they are running.

- 1) In the Services window, scroll to Service Location Protocol. Click on this service to select it.
- 2) If it is running (the **Status** column shows Started), right-click the service, then click **Stop** on the menu.

(If you did not stop the IBM CIM Object Manager - SVCCISCO service, the system now asks if you want to stop the IBM CIM Object Manager - SVCCISCO . Because the IBM CIM Object Manager - SVCCISCO service is dependent on the Service Location Protocol service which you just stopped, you must click **Yes** to stop the IBM CIM Object Manager - SVCCISCO.)

- 3) Wait for the services to stop.
- 4) Close the Services window.

- 5) Close the Administrative Tools window.
3. Use the Windows Add/Remove Programs facility to remove the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 and the Service Location Protocol components.
  - a. From the Windows menu bar, click **Start -> Settings -> Control Panel**. Double-click **Add/Remove Programs**.
  - b. Click **IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000** from the list of currently installed programs and click **Remove** to remove the product.
4. The Welcome panel for the Uninstaller opens. Click **Next** to continue or click **Cancel** to stop the removal of the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000.
5. The program detects whether the Service Location Protocol, IBM CIM Object Manager - SVCCISCO, and the IBM Websphere Application Server V5 - SVCCISCO services are running.
  - If any of these services are found to be running, the uninstaller will stop these services before proceeding with the uninstallation. You should consider at this point whether applications other than the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 are dependent on the services. You can either:
    - Click **Next** to have the program stop the services for you.
    - Click **Cancel** to exit the removal process if you wish to manually stop the services and any dependent applications. Instructions for stopping the services are described in step 2 on page 141. You must then restart the removal process from the Windows Add/Remove facility.
6. The Confirmation panel opens. Click **Remove** to continue or click **Cancel** to stop the removal of the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000. Click **Back** to return to the previous panel.
7. The Uninstallation Progress panel opens. Wait for the program to remove the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 product.
8. The Finish panel for the Uninstaller opens. This panel indicates the result of the removal process (successful or failed). Click **Finish** to complete the removal process and exit the wizard.

**Note:** If the Uninstaller could not remove some information from the system, you will see a **Next** button instead of a **Finish** button. Click **Next** to open the Reboot panel. If the reboot panel opens, you can choose to either restart your computer now or restart your computer at a later time. Then click **Finish** to complete the removal process and exit the wizard.

9. Close the Add/Remove Programs window.

#### **Post-processing requirements:**

Perform the following steps to complete the removal process:

1. If the system has not been restarted since IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000 was removed, do so now.
2. Log onto the system as a local system administrator.
3. The removal process saves files uniquely related to the configuration in a backup directory under the destination path where you installed the IBM TotalStorage Console for SAN Volume Controller for Cisco MDS 9000. You may



want those files if you intend to reinstall the product. Otherwise you can remove the backup folder and files. An example of the default destination path is:  
C:\Program Files\IBM\svccisco.

4. Perform other cleanup tasks:
  - Empty your Windows Recycle Bin to reclaim the disk space that was made available during the removal process.



---

## Part 5. Command-Line Interface



---

## Chapter 20. Getting started with the Command-Line Interface

### Overview:

The SAN Volume Controller Storage Software for Cisco MDS 9000 cluster Command-Line Interface (CLI) is a collection of commands which enable you to manage the SAN Volume Controller Storage Software for Cisco MDS 9000.

Before using the CLI, you must have created the cluster. Refer to Cisco MDS 9000 publications for information about creating a cluster.

The functions that can be performed with the IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000 Command-Line Interface (CLI) are:

- Setup of the cluster, its nodes, and the I/O groups (or node pairs). This function includes diagnostics and error log analysis of the cluster.
- Setup and maintenance of managed disks and managed disk groups.
- Setup and maintenance of virtual disks.
- Setup of logical host objects.
- Mapping of virtual disks to hosts.
- Navigation from managed hosts to virtual disk groups and to managed disks, and the reverse direction up the chain.
- Setup and trigger of Copy Services:
  - FlashCopy and FlashCopy Consistency groups
  - Synchronous Remote Copy and Remote Copy Consistency groups

Cisco MDS 9000 also provides a CLI that allows you to perform management and service functions. Two commands, create cluster and upgrade, are only available through the Cisco MDS 9000 CLI. Other Cisco MDS CLI commands allow you to reset a node, put a node in service mode, get cluster information, get node information, recover a cluster, and change a WWNN or WWPN. Refer to Cisco MDS 9000 publication for further information. You can also configure the Cisco MDS 9000 Caching Services Module so that you can use telnet instead of SSH.

### Related topics:

- “Issuing CLI commands from a PuTTY SSH Client system”
- “Running the PuTTY and plink utilities” on page 148
- “Configuring the cluster using the CLI” on page 150
- “Related publications” on page xi

---

## Issuing CLI commands from a PuTTY SSH Client system

### Steps:

Perform the following steps to issue CLI commands:

1. Open a Command Prompt to open the SSH connection to issue the CLI commands.
2. Make the PuTTY executables available by performing the following:
  - a. On a host that has installed PuTTY in the default location type the following:

C:\Program Files\Putty

- b. Set the path environment variable to include the PuTTY executable directory. For example, type the following:

```
Set path=c:\Support Utils\putty;%path%
```

3. Use the PuTTY plink utility to connect to the SSH server on the cluster.

**Related topics:**

- “Running the PuTTY and plink utilities”

---

## Running the PuTTY and plink utilities

All CLI commands are run in an SSH session. You can run the commands in either of the two following modes:

- an interactive prompt mode
- in a single line command mode, which is entered one time to include all parameters.

**Interactive mode:**

For interactive mode, you use the PuTTY executable to open the SSH restricted shell. Type the following:

```
C:\support utils\putty>putty admin@<svcconsoleip>
```

If you were to issue the **svcinfolssshkeys** command, which lists the SSH client public keys that are stored on the SAN Volume Controller cluster, the following output is displayed:

```
IBM_2062:admin>svcinfolssshkeys -user all -delim :
id:userid:key identifier
1:admin:smith
2:admin:jones
```

Type `exit` and press **Enter** to escape the interactive mode command.

The SSH protocol specifies that the first access to a new host server will result in a *challenge* to the SSH user to accept the SSH server public key. Because this is the first time that you connect to an SSH server, the server is not included in the SSH client list of known hosts. Therefore, there is a fingerprint challenge, which asks you do you accept the responsibility of connecting with this host. If you type `y`, the host fingerprint and IP address is saved by the SSH client. For PuTTY, you answer by typing `y` to accept this host fingerprint. This information is stored in the registry for the user name which is logged onto Windows.

The following is an example of the host fingerprint challenge when running in interactive mode:

```
C:\Program Files\IBM\svcconsole\cimom>plink admin@9.43.225.208
The server's host key is not cached in the registry. You
have no guarantee that the server is the computer you
think it is.
The server's key fingerprint is:
ssh-rsa 1024 e4:c9:51:50:61:63:e9:cd:73:2a:60:6b:f0:be:25:bf
If you trust this host, enter "y" to add the key to
```

PuTTY's cache and carry on connecting.  
If you want to carry on connecting just once, without adding the key to the cache, enter "n".  
If you do not trust this host, press Return to abandon the connection.  
Store key in cache? (y/n) y  
Using username "admin".  
Authenticating with public key "imported-openssh-key"  
IBM\_2062:admin>

### Single line command:

For single line command mode, you can type the following all on one command line:

```
C:\Program Files\IBM\svconconsole\cimom>
plink admin@9.43.225.208 svcinfo lssshkeys
-user all -delim :
Authenticating with public key "imported-openssh-key"
/bin/ls: id:userid:key identifier
1:admin:smith
2:admin:jones
```

```
C:\Program Files\IBM\svconconsole\cimom>
```

The SSH protocol specifies that the first access to a new host server will result in a *challenge* to the SSH user to accept the SSH server public key. Because this is the first time that you connect to an SSH server, the server is not included in the SSH client list of known hosts. Therefore, there is a fingerprint challenge, which asks you do you accept the responsibility of connecting with this host. If you type y, the host fingerprint and IP address is saved by the SSH client. For PuTTY, you answer by typing y to accept this host fingerprint. This information is stored in the registry for the user name which is logged onto Windows.

The following is an example of the host fingerprint challenge when running in single line command mode:

```
C:\Program Files\IBM\svconconsole\cimom>
plink admin@9.43.225.208 svcinfo lssshkeys
-user all -delim :
The server's host key is not cached in the registry. You
have no guarantee that the server is the computer you
think it is.
The server's key fingerprint is:
ssh-rsa 1024 e4:c9:51:50:61:63:e9:cd:73:2a:60:6b:f0:be:25:bf
If you trust this host, enter "y" to add the key to
PuTTY's cache and carry on connecting.
If you want to carry on connecting just once, without
adding the key to the cache, enter "n".
If you do not trust this host, press Return to abandon the
connection.
Store key in cache? (y/n) y
Authenticating with public key "imported-openssh-key"
/bin/ls: /proc/20282/exe: Permission denied
dircolors: `etc/DIR_COLORS': Permission denied
id:userid:key identifier
```

```
1:admin:smith
2:admin:jones
```

```
C:\Program Files\IBM\svconconsole\cimom>
```

**Note:** If you are submitting a CLI command with all parameters in single line command mode, you will be challenged upon first appearance of the SSH server host fingerprint. Be careful to ensure that the SSH server host fingerprint is accepted before submitting a batch script file.

The SSH protocol also specifies that once the SSH server public key is accepted, another challenge will be presented if the fingerprint of an SSH server changes from the one previously accepted. In this case, you will need to decide whether to accept this changed host fingerprint. For PuTTY, you answer by typing *y* to accept this host fingerprint. PuTTY stores this information in the registry for the user name which is logged onto Windows.

**Note:** The SSH server keys on the SAN Volume Controller will be regenerated when a microcode load is performed on the cluster. Due to this behavior, you will see a *challenge* presented because the fingerprint of the SSH server has changed.

**Related topics:**

- “Configuring the PuTTY session for the command-line interface” on page 54

---

## Configuring the cluster using the CLI

**Steps:**

Perform the following steps to configure the cluster:

1. Open a command prompt window.
2. If you choose you can review and modify the initial configuration of the cluster using the following steps:
3. Enable your featurization settings by issuing the **svctask chlicense** command. You will need to specify whether you want the FlashCopy or Remote Copy enabled or disabled. You can also specify your size capacity for virtualization.
4. If you choose to set up error notifications to manage your errors from the cluster, you can issue the **svcservicemodetask setevent** command to set up SNMP traps.

**Related topics:**

- “Displaying cluster properties using the CLI” on page 151
- “Reviewing and setting the cluster features using the CLI”
- “Setting up error notifications using the CLI” on page 196

---

## Reviewing and setting the cluster features using the CLI

**Steps:**

Perform the following steps to set up the cluster features:

1. Open a command prompt window.



2. Issue the **svctask lslicense** command to return the current license (featurization) settings for the cluster. The output displayed lists the feature functions in a list and displays whether they are enabled or disabled.
3. Issue the **svcinfo lslicense** command to change the licensed settings of the cluster. Because the feature settings are entered when the cluster is first created, you need to update the settings only if you have changed your license. You can change the following values:
  - FlashCopy: disabled or enabled
  - Remote Copy: disabled or enabled
  - Virtualization limit: number, in gigabytes (1073741824 bytes)

---

## Displaying cluster properties using the CLI

### Steps:

Perform the following steps to display cluster properties:

1. Open a command prompt window.
2. Issue the **svcinfo lscluster** command to display a concise view of the cluster.

```
svcinfo lscluster -delim : 10030a007e5
```

where *10030a007e5* is the name of the cluster. The output from this command will display the following: The output from this command will include the following for each cluster on the fabric:

- cluster ID
- cluster name
- cluster IP address



---

## Chapter 21. Scenario: typical usage for the command-line interface

The main focus of the following example is to provide storage to your host system.

For example, you wish to provide a host system with two disks and create a FlashCopy of these two disks. The copy is to be made available to a second host. These two hosts require that the host objects that are created, correspond with the group of WWPNs presented by their fibre-channel HBAs to the SAN. You also need to create four virtual disks, one for each of the disks that are to be presented to the hosts. Once the VDIs are created, you can map two of them to each host. In order to create the VDIs you need to have a managed disk group to be able to create them from. You wish to spread the 8 managed disks across two groups and create the source VDIs from one and the target VDIs from the other. In order to create any of these objects you need to create a cluster and at least add one more node to the cluster.

The following steps illustrates how this can be done:

1. Create a cluster.
2. Configure the cluster with an IP address of 9.20.123.456. Name the cluster `examplecluster`.
3. Add nodes
  - `knode` and `lnode` to the I/O group called `io_grp0` in the `examplecluster` cluster
  - `mnode` and `nnode` to the I/O group called `io_grp1` in the `examplecluster` cluster
4. Create the MDisk groups `maindiskgroup` and `bkpdiskgroup`
5. Create 4 VDIs
  - 2 VDIs from `maindiskgroup`
  - 2 VDIs from `bkpdiskgroup`
6. Create 2 host objects
  - a host object called `demohost1` with HBAs that have WWPNs of `10000000C92AD7E5`, and `10000000C92F5123`
  - a host object called `demohost2` with HBAs that have WWPNs of `210000E08B0525D4`, and `210100E08B2525D4`
7. Create the VDI-to-host mappings
  - Map the two VDIs from `maindiskgroup` to `demohost1`
  - Map the two VDIs from `bkpdiskgroup` to `demohost2`
8. Create FlashCopy mappings
  - Create a FlashCopy mapping called `main1copy` that has a background copy rate of 75
  - Create a FlashCopy mapping called `main2copy` that has a background copy rate of 50
9. Create a FlashCopy consistency group called `maintobkpfcopy` and add the 2 FlashCopy mappings to it
10. Prepare and trigger (start) the FlashCopy Consistency Group that contains these mappings

**Note:** Once this step is complete, you have created and allocated storage to your host systems. You have made two VDisks available to demohost1 and then used FlashCopy to make backup copies on two VDisks which are accessible to demohost2.

**Related topics:**

- “Configuring a cluster using the Console for SAN Volume Controller for Cisco MDS 9000” on page 67
- “Adding nodes to a cluster using the IBM TotalStorage CLI”
- “Displaying node properties using the CLI” on page 156
- “Discovering MDisks using the CLI” on page 158
- “Creating managed disk (MDisk) groups using the CLI” on page 159
- “Adding MDisks to MDisk groups using the CLI” on page 161
- “Create virtual disks (VDisks)” on page 162
- “Creating host objects using the CLI” on page 165
- “Create VDisk-to-host mappings using the CLI” on page 166
- “Creating a FlashCopy consistency group and adding mappings using the CLI” on page 168
- “Create FlashCopy mappings using the CLI” on page 167
- “Preparing and triggering a FlashCopy Consistency Group using the CLI” on page 169

---

## Adding nodes to a cluster using the IBM TotalStorage CLI

**Context:**

Before you begin adding nodes, you must have already created a cluster, performed the initial configuration using the Console for SAN Volume Controller for Cisco MDS 9000 and performed the necessary setup to use the Command Line Interface (CLI).

**Prerequisites:**

| Enter a mode where you can issue the IBM TotalStorage CLI commands. Issue the  
| Cisco switch command **svc-ibmcli cluster-name SampCluster**. In this example,  
| **SampCluster** is the name of the cluster.

**Steps:**

Perform the following steps to add nodes to a cluster:

1. Open a command prompt window.
2. Determine the I/O group to which the node should be added. Type the **svcinfo lsnode** command to list the nodes that are currently part of the cluster.

**Example:**

```
svcinfo lsnode -delim :

id:name:card_serial_number:WWNN:status:IO_group_id:IO_group_name:
 config_node:card_unique_id
1:node1:JAB0649055d:2541000530003C20:online:
 0:io_grp0:yes:0100064930353564
```

3. Type the **svcinfolnodecandidate** command to list nodes that are not assigned to a cluster and to verify that when adding a second node to an I/O group that it is attached to a different Cisco MDS 9000 Caching Services Module.

**Example:**

```
svcinfolnodecandidate -delim :

id:panel_name:card_serial_number:card_unique_id
2543000530003C20:switchname.slot.node:JAB0650081e:0100065030383165
254F000530003C20:switchname.slot.node:JAB071105kr:0100071130356B72
2559000530003C20:switchname.slot.node:JAB06500812:0100065030383132
```

4. Issue the command **svctask addnode** to add a node to the cluster.

You can specify a name for the node. To change the name of a node that is already part of a cluster, use the **svctask chnode** command.

If you don't specify the name for the node, the node can later be identified by using the world wide node name of that node or the name as it appears in the output of **svcinfolnodecandidate**.

For the new node, record the following information:

- Node serial number
- WWNN
- All WWPNs
- I/O group that contains the node

This can avoid a possible data corruption exposure if the node must be removed from and re-added to the cluster.

If you are re-adding a node to the cluster, ensure that you are adding the node to the same I/O group that it was removed from. Failure to do this can result in data corruption. Use the information that was recorded when the node was originally added to the cluster. If you do not have access to this information, call IBM Service to add the node back into the cluster without corrupting the data.

**Example:**

Add a second node to the first I/O group by issuing the following command.

```
svctask addnode -panelname switchname.slot.node -iogrp io_grp0 -name group1node2
```

This command adds the node to the cluster. The node will be added to I/O group, `io_grp0`, and called `group1node2`.

Next add two nodes to the second I/O group. Make sure that each node is attached to a different Cisco MDS 9000 Caching Services Module.

```
svctask addnode -wwnodename 2559000530003C20 -iogrp io_grp1 -name group2node1
svctask addnode -wwnodename 254F000530003C20 -iogrp io_grp1 -name group2node2
```

These commands will add the nodes, identified by the WWNN 2559000530003C20 and the WWNN 254F000530003C20 to the cluster. The nodes will be added to I/O group, `io_grp1` and called `group2node1` and `group2node2`.

Finally change the name of the first node from the default name `node1` so that it conforms with your naming convention.

```
svctask chnode -name group1node1 node1
```

5. Verify the final configuration using the **svcinfolnode** command.

**Example:**

In our hypothetical scenario, the command to list the nodes is:

```
svcinfolnode -delim :
```

```
id:name:card_serial_number:WWNN:status:IO_group_id:
 IO_group_name:config_node:card_unique_id
1:group1node1:JAB0649055d:2541000530003C20:online:0:io_grp0:yes:0100064930353564
2:group1node2:JAB0650081e:2543000530003C20:online:0:io_grp0:no:0100065030383165
3:group2node1:JAB06500812:2559000530003C20:online:1:io_grp1:no:0100065030383132
4:group2node2:JAB071105kr:254F000530003C20:online:1:io_grp1:no:0100071130356B72
```

**Note:** If this command is issued quickly after adding nodes to the cluster the status of the nodes may be adding rather than online indicating that the process of adding the nodes to the cluster is still in progress. You do not however have to wait for all the nodes to become online before continuing with the configuration process.

**Remember:** Record the following information:

- WWNN
- All WWPNs
- I/O group that contains the node
- Card serial number
- The slot where the Cisco MDS 9000 Caching Services Module resides
- The Cisco MDS 9000 Caching Services Module node (1/2) the node is defined as

This can avoid a possible data corruption exposure if the node must be removed from and re-added to the cluster.

**Result:**

You have now added four nodes to one cluster. The nodes are split into two I/O groups.

**Related topics:**

- Chapter 21, “Scenario: typical usage for the command-line interface,” on page 153
- “Displaying node properties using the CLI”

---

## Displaying node properties using the CLI

**Steps:**

Perform the following steps to display the node properties:

1. Open a command prompt window.
2. Issue the **svcinfolnode** command to display a concise list of nodes in the cluster.

**Example:**

Type the following command:

```
svcinfolnode -delim :
```

This command displays the following:

```
id:name:card_serial_number:WWNN:status:IO_group_id:IO_group_name:
 config_node:card_unique_id
1:node1:JAB0649055d:2541000530003C20:online:0:io_grp0:yes:0100064930353564
```

3. Issue the **svcinfolnode** command again, however, this time, specify the node ID or name of a node to receive the detailed output.

**Example:**

For example, to provide a detailed view of the node named 1 type the following:

```
svcinfolnode -delim = 1
```

This command displays the following:

```
id=1
name=node1
card_serial_number=JAB071007HB
WWNN=2548000530003920 status=online
IO_group_id=0
IO_group_name=io_grp0
partner_node_id=2
partner_node_name=node2
config_node=yes
card_unique_id=0100071030374842
port_id=1E39003005000220
port_status=active
```

The output includes:

- node ID
- node name
- WWNN
- details about the node status
- details about the I/O group of which the node is a member
- detailed fibre channel port status information.

---

## Removing managed disks that represent de-configured LUs

When deconfiguring or removing LUs from your storage subsystem, the managed disks (MDisks) that represent those LUs may still exist in the cluster. Use the following procedure to remove those MDisks.

**Context:**

MDisks exist in the cluster that are no longer accessible to it. That is because the LUs that these MDisks represent have been deconfigured or removed from the storage subsystem. You should remove these MDisks.

**Steps:**

Perform the following steps to remove the MDisks:

1. Run the **svctask includemdisk** command on all the affected MDisks.
2. Run the **svctask rmmdisk** command on all affected MDisks. This puts the MDisks into the unmanaged mode.
3. Run the **svctask detectmdisk** command. The cluster detects that the MDisks no longer exist in the storage subsystem.

**Result:**

All of the MDisks that represent de-configured LUs are removed from the cluster.

**Related topics:**

- “Discovering MDisks using the CLI”

---

## Discovering MDisks using the CLI

**Context:**

When backend controllers are added to the fibre-channel SAN and are included in the same switch zone as a SAN Volume Controller Storage Software for Cisco MDS 9000 Cluster the cluster will automatically discover the backend controller and will integrate the controller to determine what storage it is presented to the SAN Volume Controller Storage Software for Cisco MDS 9000. The SCSI LUs presented by the backend controller will be displayed as unmanaged MDisks. If however the configuration of the backend controller is modified after this has occurred then the SAN Volume Controller Storage Software for Cisco MDS 9000 may be unaware of these configuration changes. This task allows a user to request the SAN Volume Controller Storage Software for Cisco MDS 9000 to re-scan the fibre-channel SAN to update the list of unmanaged MDisks.

**Note:** The automatic discovery performed by SAN Volume Controller Storage Software for Cisco MDS 9000 does not write anything to a unmanaged MDisk. It is only when a the user instructs the SAN Volume Controller Storage Software for Cisco MDS 9000 to add a MDisk to a managed disk group or use a MDisk to create an image mode virtual disk that the storage will actually be used.

**Steps:**

Perform the following steps to display MDisks:

1. Open a command prompt window.
2. Check to see which MDisks are available by issuing the **svctask detectmdisk** command to manually scan the fibre-channel network for any MDisks.
3. Issue the **svcinfo lsmdiskcandidate** command to show the unmanaged MDisks. These MDisks have not been assigned to an MDisk group. Alternatively, you can issue the **svcinfo lsmdisk** command to view all of the MDisks.

**Example:**

In our hypothetical scenario we have a single backend controller that is presenting eight SCSI LUs to the SAN Volume Controller Storage Software for Cisco MDS 9000. Issue the following command to cause the cluster to discover all available MDisks on the SAN:

```
svctask detectmdisk
```

```
svcinfo lsmdiskcandidate
```

This command displays the following:

```
id
0
1
```



2  
3  
4  
5  
6  
7

Issue the following command:

```
svcinfolsmdisk -delim : -filtervalue mode=unmanaged
```

This command displays the following:

```
id:name:status:mode:mdisk_grp_id:mdisk_grp_name:
 capacity:ctrl_LUN#:controller_name
0:mdisk0:online:unmanaged:::273.3GB:0000000000000000:controller0
1:mdisk1:online:unmanaged:::273.3GB:0000000000000001:controller0
2:mdisk2:online:unmanaged:::273.3GB:0000000000000002:controller0
3:mdisk3:online:unmanaged:::273.3GB:0000000000000003:controller0
4:mdisk4:online:unmanaged:::136.7GB:0000000000000004:controller0
5:mdisk5:online:unmanaged:::136.7GB:0000000000000005:controller0
6:mdisk6:online:unmanaged:::136.7GB:0000000000000006:controller0
7:mdisk7:online:unmanaged:::136.7GB:0000000000000007:controller0
```

### Result:

You have now shown that the backend controllers and switches have been setup correctly and that the SAN Volume Controller Storage Software for Cisco MDS 9000 can see the storage being presented by the backend controller.

### Related topics:

- Chapter 21, “Scenario: typical usage for the command-line interface,” on page 153

---

## Creating managed disk (MDisk) groups using the CLI

**Attention:** If you add a MDisk to a MDisk group as a managed disk, any data on the MDisk will be lost. If you want to keep the data on a MDisk (for example because you want to import storage that was previously not managed by a SAN Volume Controller Storage Software for Cisco MDS 9000) then you should create image mode VDisks instead.

### Context:

Assume that the cluster has been setup and that a backend controller has been configured to present some new storage to the SAN Volume Controller Storage Software for Cisco MDS 9000.

### Prerequisites:

Before creating managed disk groups consider how you are going to use your storage. The SAN Volume Controller Storage Software for Cisco MDS 9000 allows

you to create up to 128 managed disks groups and to add up to 128 MDisks to an MDisk group. Consider the following factors when deciding how many managed disk groups to create:

- A virtual disk can only be created using the storage from one managed disk group. Therefore if you create small managed disk groups then you may lose the benefits provided by virtualization, namely more efficient management of free space and a more evenly distributed workload to provide better performance.
- If any managed disk in a managed disk group goes offline then all the virtual disks in the managed disk group will go offline. Therefore you might want to consider using different managed disk groups for different backend controllers or for different applications.
- If you anticipate regularly adding and removing backend controllers or storage then this task will be made simpler by grouping all the managed disks presented by a backend controller into one managed disk group.
- All the managed disks in a managed disk group should have similar levels of performance or reliability, or both. If a managed disk group contains managed disks with different levels of performance then the performance of the virtual disks in this group will be limited by the performance of the slowest managed disk. If a managed disk group contains managed disks with different levels of reliability then the reliability of the virtual disks in this group will be that of the least reliably managed disk in the group.

Even with the best planning, circumstances may change and you may wish to re-configure your managed disk groups after they have been created. The data migration facilities provided by the SAN Volume Controller Storage Software for Cisco MDS 9000 enable you to move data without disrupting I/O.

**Choosing a managed disk group extent size:** You must specify the extent size when you create a new MDisk group. You cannot change the extent size later; it must remain constant throughout the lifetime of the MDisk group. MDisk groups may have different extent sizes however this will place restrictions on the use of data migration. The choice of extent size affects the total amount of storage that can be managed by a SAN Volume Controller Storage Software for Cisco MDS 9000 Cluster. Table 12 shows the maximum amount of storage that can be managed by a cluster for each extent size. Because the SAN Volume Controller Storage Software for Cisco MDS 9000 allocates a whole number of extents to each virtual disk that is created, using a larger extent size may increase the amount of wasted storage at the end of each virtual disk. Larger extent sizes also reduces the ability of the SAN Volume Controller Storage Software for Cisco MDS 9000 to distribute sequential I/O workloads across many managed disks and hence may reduce the performance benefits of virtualization.

*Table 12. Extent size*

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

**Attention:** You can specify different extent sizes for different managed disk groups, however you will not be able to migrate virtual disks between managed disk groups with different extent sizes. Therefore if possible create all your managed disk groups with the same extent size.

### Steps:

Perform the following steps to create a MDisk group:

1. Open a command prompt window.
2. Type the **svctask mkmdiskgrp** command to create a MDisk group.

#### Example:

In our hypothetical scenario, the command to create a MDisk group is:

```
svctask mkmdiskgrp -name maindiskgroup -ext 32 -mdisk msk0:msk1:msk2:msk3
```

This command will create a MDisk group called *maindiskgroup*. The extent size used within this group will be 32 MB, and there are four MDisks *msk0*, *msk1*, *msk2*, *msk3* added to the group.

#### Example:

In our hypothetical scenario, the command to create a second MDisk group is:

**Note:** In this example we will create a second MDisk group first and add MDisks later.

```
svctask mkmdiskgrp -name bkpdiskgroup -ext 32
```

This command will create a MDisk group called *bkpdiskgroup*. The extent size used within this group will be 32 MB.

#### Example:

To add MDisks to the MDisk group, issue the **svctask addmdisk** command. In our hypothetical scenario, the command to add MDisks to the MDisk group is:

```
svctask addmdisk -mdisk msk4:msk5:msk6:msk7 bkpdiskgroup
```

This command will add four MDisks *msk4*, *msk5*, *msk6*, *msk7* to the MDisk group called *bkpdiskgroup*.

### Related topics:

- Chapter 21, “Scenario: typical usage for the command-line interface,” on page 153
- “Adding MDisks to MDisk groups using the CLI”
- “Creating an image mode VDisk from an unmanaged MDisk using the CLI” on page 193

---

## Adding MDisks to MDisk groups using the CLI

The managed disks must be in unmanaged mode. Disks that already belong to a group cannot be added to another group until they have been deleted from their current group. You can delete a managed disk from a group under the following circumstances:

- If the managed disk does not contain any extents in use by a virtual disk

- If you can first migrate the extents in use onto other free extents within the group.

**Steps:**

Perform the following steps to add MDisks to MDisk groups:

1. Open a command prompt window.
2. Type the **svcinfo lsmdiskgrp** command to list the existing MDisk groups.

**Example:**

In our hypothetical scenario, we have two MDisk groups, one with four managed disks and one with no managed disks. Type the following command:

```
svcinfo lsmdiskgrp -delim :
```

This command displays the following:

```
id:name:status:mdisk_count:vdisk_count:
capacity:extent_size:free_capacity
0:mainmdiskgroup:online:4:0:1093.2GB:32:1093.2GB
1:bkpmdiskgroup:online:0:0:0:32:0
```

3. To add MDisks to the MDisk group, issue the **svctask addmdisk** command.

**Example:**

In our hypothetical scenario, the command to add MDisks to the MDisk group is:

```
svctask addmdisk -mdisk mdisk4:mdisk5:mdisk6:mdisk7 bkpmdiskgroup
```

This command will add four MDisks mdisk4, mdisk5, mdisk6 and mdisk7 to the MDisk group called bkpmdiskgroup.

**Related topics:**

- Chapter 21, “Scenario: typical usage for the command-line interface,” on page 153

## Create virtual disks (VDisks)

**Note:** If you want to keep the data on a MDisk (for example because you want to want to import storage that was previously not managed by a SAN Volume Controller Storage Software for Cisco MDS 9000) then you should create image mode VDIsks instead. This task only deals with creating VDIsks with a striped virtualization policy. For details of other virtualization policies refer to the *IBM TotalStorage Subsystem Device Driver: User's Guide*.

**Context:**

Assume that the cluster has been setup and that you have created managed disk groups. You must establish an empty managed disk group to hold the MDisks used for image mode VDIsks.

**Steps:**

Perform the following steps to create VDIsks:

1. Open a command prompt window.
2. Decide which managed disk group will provide the storage for the VDisk. Use the **svcinfolsmdiskgrp** command to list the available managed disk groups and the amount of free storage in each group.

**Example:**

In our hypothetical scenario, issue the following:

```
svcinfolsmdiskgrp -delim :
```

This command displays the following:

```
id:name:status:mdisk_count:vdisk_count:
capacity:extent_size:free_capacity
0:mainmdiskgroup:online:4:0:1093.2GB:32:1093.2GB
1:bkpmdiskgroup:online:4:0:546.8GB:32:546.8GB
```

3. Decide which I/O group the VDisk should be assigned to. This determines which SAN Volume Controller Storage Software for Cisco MDS 9000 nodes in the cluster process the I/O requests from the host systems. If you have more than one I/O group then make sure you distribute the VDIsks between the I/O groups so that the I/O workload is shared evenly between all SAN Volume Controller Storage Software for Cisco MDS 9000 nodes. Use the **svcinfolsiogrp** command to show the I/O groups and the number of virtual disks assigned to each I/O group.

**Note:** It is normal for clusters with more than one I/O group to have MDisk groups that have VDIsks in different I/O groups. FlashCopy can be used to make copies of VDIsks regardless of whether the source and destination VDisk are in the same I/O group. If however you plan to use intra-cluster remote copy then make sure that both the master and auxiliary VDisk are in the same I/O group.

**Example:**

In our hypothetical scenario, there are two I/O groups each with two nodes. Neither I/O group has any virtual disks yet. Issue the following command:

```
svcinfolsiogrp -delim :
```

This command displays the following:

```
id:name:node_count:vdisk_count
0:io_grp0:2:0
1:io_grp1:2:0
2:io_grp2:0:0
3:io_grp3:0:0
4:recovery_io_grp:0:0
```

4. Type the **svctask mkvdisk** command to create a virtual disk (VDisk).

**Example:**

In our hypothetical scenario, the command to create a VDisk is:

```
svctask mkvdisk -name mainvdisk1 -iogrp 0
-mdiskgrp 0 -vtype striped -size 256 -unit gb
```

This command will create a VDisk called mainvdisk1, the VDisk will use I/O group 0 and MDisk group 0 (the ID of maindiskgroup as shown in the output from step 2). The VDisk capacity is 256GB and will be made up of extents from the MDisks in the MDisk group.

**Example:**

In our hypothetical scenario, the command to create a second VDisk is:

**Note:** This command is the same as the above example, however, here we are specifying the names of the objects instead of the IDs.

```
svctask mkvdisk -name mainvdisk2 -iogrp io_grp0
-mdiskgrp maindiskgroup -vtype striped -size 256 -unit gb
```

This command will create a VDisk called mainvdisk2, the VDisk will use the I/O group named io\_grp0 and the MDisk group named maindiskgroup. The VDisk capacity is 256GB and will be made up of extents from the MDisks in the MDisk group.

**Example:**

In our hypothetical scenario, the commands to create a third VDisk is:

**Note:** This virtual disk is created with an ordered list of MDisks within the MDisk group to allocate extents from.

The following command lists the managed disks in the MDisk group with ID 1 (named bkpmdiskgroup):

```
svcinfolsmdisk -delim : -filtervalue mdisk_grp_id=1
```

This command displays the following:

```
id:name:status:mode:mdisk_grp_id:
mdisk_grp_name:capacity:ctrl_LUN_#:
controller_name
4:mdisk4:online:managed:1:bkpmdiskgroup:
136.7GB:0000000000000004:controller0
5:mdisk5:online:managed:1:bkpmdiskgroup:
136.7GB:0000000000000005:controller0
6:mdisk6:online:managed:1:bkpmdiskgroup:
136.7GB:0000000000000006:controller0
7:mdisk7:online:managed:1:bkpmdiskgroup:
136.7GB:0000000000000007:controller0
```

Issue the following command:

```
svctask mkvdisk -name bkpvdisk1 -iogrp io_grp1
-mdiskgrp bkpmdiskgrp -vtype striped -size 256
-unit gb -mdisk 4:5
```

This command will create a VDisk called bkpvdisk1, the Vdisk will use the I/O group named io\_grp1 and the MDisk group named bkpmdiskgrp. The VDisk capacity is 256GB and will be made up of extents allocated from the MDisks with IDs 4 and 5.

**Example:**

In our hypothetical scenario, the command to create a fourth VDisk is:

```
svctask mkvdisk -name bkpvdisk2 -iogrp io_grp1
-mdiskgrp bkpmdiskgrp -vtype striped -size 256 -unit
gb -mdisk mdisk6:mdisk7
```

This command will create a VDisk called bkpvdisk2, the VDisk will use the I/O group named io\_grp1 and the MDisk group named bkpmdiskgrp. The VDisk capacity is 256GB and will be made up of extents allocated from the MDisks with names mdisk6 and mdisk7.

5. To list all the virtual disks that have been created use the **svcinfolsvdisk** command.

**Example:**

In our hypothetical scenario we have created four VDisks. Issue the following command:

```
svcinfolsvdisk -delim :
```

This command displays the following:

```
id:name:I0_group_id:I0_group_name:status:
mdisk_grp_id:mdisk_grp_name:capacity:type:FC_id:
FC_name:RC_id:RC_name
0:mainvdisk1:0:io_grp0:online:0:mainmdiskgroup:
512.0GB:striped:::
1:mainvdisk2:0:io_grp0:online:0:mainmdiskgroup:
512.0GB:striped:::
2:bkpvdisk1:1:io_grp1:online:1:bkpmdiskgroup:
512.0GB:striped:::
3:bkpvdisk2:1:io_grp1:online:1:bkpmdiskgroup:
512.0GB:striped:::
```

**Related topics:**

- Chapter 21, “Scenario: typical usage for the command-line interface,” on page 153
- “Creating an image mode VDisk from an unmanaged MDisk using the CLI” on page 193

---

## Creating host objects using the CLI

**Steps:**

Perform the following steps to create host objects:

1. Open a command prompt window.
2. Type the **svctask mkhost** command to create a logical host object. Assign your WWPN for the HBAs in the hosts.

**Example:**

In our hypothetical scenario, the command to create a host is:

```
svctask mkhost -name demohost1 -hbawwpn 210100e08b251dd4
```

This command will create a host called *demohost1* with the HBA WWPN of *210100e08b251dd4*.

3. Type the **svctask addhostport** command to add ports to the host.

**Example:**

In our hypothetical scenario, the command to add a port to the host is:

```
svctask mkhost -name demohost2 -hbawwpn 210100e08b251dd5
```

This command will add another HBA WWPN called *210100e08b251dd5* to the host that we created in step 2 on page 165.

**Example:**

In our hypothetical scenario, the command to create a second host is:

```
svctask mkhost -hbawwpn 210100e08b251dd6:210100e08b251dd7 -name demohost2
```

This command will create a second host called *demohost2* with the HBA WWPN of *210100e08b251dd6*, *210100e08b251dd7*.

**Note:** If you were to add a host with a faulty WWPN, or the WWPN had been assigned to the wrong host, you will need to issue the **svctask addhostport** command to add that same host with the correct WWPN, then issue the **svctask rmhostport** command to delete the host with the wrong or faulty WWPN. For example, if you had a host called *demohost1* and its WWPN stopped working, you would need to issue the following:

```
svctask addhostport -hbawwpn 210100e08b251dd4 demohost1
```

This would add the host called *demohost1* with the WWPN, *210100e08b251dd4*. You would then need to issue the **svctask rmhostport** command to delete the host with the WWPN that had stopped working. For example, you would issue the following:

```
svctask rmhostport -hbawwpn 210100e08b251dd5 demohost1
```

From these two commands, you have deleted the host with the WWPN *210100e08b251dd5*, and have added the same host with the WWPN *210100e08b251dd4*.

**Related topics:**

- Chapter 21, “Scenario: typical usage for the command-line interface,” on page 153

---

## Create VDisk-to-host mappings using the CLI

**Prerequisites:**

We are going to map the VDisks named, *mainvdisk1* and *mainvdisk2*, to the host named *demohost1*. We are also going to map the VDisks named, *bkpvdisk1* and *bkpvdisk2*, to the host named *demohost2*. The VDisks, *mainvdisk1* and *mainvdisk2*,



are contained in the managed disk (MDisk) group, mainmdiskgroup; while the VDIs, bkpvdisk1 and bkpvdisk2, are contained in the MDisk group, bkpmdiskgroup.

### Steps:

Perform the following steps to create VDisk-to-host mappings:

1. Open a command prompt window.
2. Type the **svctask mkvdiskhostmap** to create VDisk-to-host mappings.

#### Example:

In our hypothetical scenario, the commands to create VDisk-to-host mappings are:

```
svctask mkvdiskhostmap -host demohost1 mainvdisk1
svctask mkvdiskhostmap -host demohost1 mainvdisk2
svctask mkvdiskhostmap -host demohost2 bkpvdisk1
svctask mkvdiskhostmap -host demohost2 bkpvdisk2
```

The above set of commands map each VDisk to a host.

### Related topics:

- Chapter 21, “Scenario: typical usage for the command-line interface,” on page 153

---

## Create FlashCopy mappings using the CLI

### Prerequisites:

We are going to create mappings that enable us to copy the VDisk, mainvdisk1, to bkpvdisk1 and the VDisk, mainvdisk2 to bkpvdisk2.

The mapping specifies the source and destination virtual disks. The destination must be identical in size to the source, or the mapping will fail. Issue the **svcinfo lsvdisk -bytes** command to find the exact size of the source VDisk that you want to create a target disk of the same size. The source and destination cannot be in an existing mapping. That is, a virtual disk can be either a source or a destination disk in **only one** mapping. A mapping is triggered at the point in time when the copy is required.

### Steps:

Perform the following steps to create FlashCopy mappings:

1. Open a command prompt window.
2. Type the **svctask mkfcmap** command to create a FlashCopy mapping.

#### Example:

In our hypothetical scenario, the commands to create FlashCopy mappings are:

```
svctask mkfcmap -source mainvdisk1 -target bkpvdisk1
 -name main1copy -copyrate 75
svctask mkfcmap -source mainvdisk2 -target bkpvdisk2
 -name main2copy
```

The above commands create two FlashCopy mappings. For main1copy the background copy rate is 75; for main2copy, because the rate is not specified in the **mkfcmap** command, the priority is the default, 50.

3. To check the attributes of the mappings that have been created, issue the following **svcinfolsfcmmap** command:

```
svcinfolsfcmmap -delim :
```

This command displays the following:

```
id:name:source vdisk id:source vdisk name:target
 vdisk id:target vdisk name:group id:group
 name:status:progress:copy rate
0:main1copy:0:mainvdisk1:1:bkpvdisk1:::idle_copied::75
1:main2copy:2:mainvdisk2:3:bkpvdisk2:::idle_copied::50
```

**Related topics:**

- Chapter 21, “Scenario: typical usage for the command-line interface,” on page 153

---

## Creating a FlashCopy consistency group and adding mappings using the CLI

If you have created several FlashCopy mappings for a group of VDisks that contain elements of data for the same application, you may find it convenient to assign these mappings to a single FlashCopy Consistency Group. Then you can issue a single prepare or trigger command for the whole group, so that, for example, all the files for a particular database are copied at the same time.

**Steps:**

Perform the following steps to create a FlashCopy mappings:

1. Open a command prompt window.
2. Issue the **svctask mkfconsisgrp** command to create a FlashCopy Consistency Group.

**Example:**

In our hypothetical scenario, the command to create a FlashCopy Consistency group called *maintobkpfcopy* is:

```
svctask mkfconsisgrp -name maintobkpfcopy
```

Use the **svcinfolsfconsisgrp** command to display the attributes of the group you have created.

```
svcinfolsfconsisgrp -delim :
```

This command displays the following:

```
id:name:status
1:maintobkpfcopy:idle_copied
```

3. Use the **svctask chfcmmap** command to add the two FlashCopy mappings created in the previous section to the new consistency group.

**Example:**

In our hypothetical scenario, the commands to add the mappings called *main1copy* and *main2copy* to the consistency group called *maintobkpfcopy* are:

```
svctask chfcmmap -consistgrp maintobkpfcopy main1copy
svctask chfcmmap -consistgrp maintobkpfcopy main2copy
```

Use the **svcinfolsfcmmap** command to display the new attributes of the mappings.

```
svcinfolsfcmmap -delim :
id:name:source_vdisk_id:source_vdisk_name:target_vdisk_id:
target_vdisk_name:group_id:group_name:state:progress:copy_rate
0:main1copy:28:maindisk1:29:bkpdisk1:1:maintobkpfcopy:idle_copied::75
1:main2copy:30:maindisk2:31:bkpdisk2:1:maintobkpfcopy:idle_copied::50
```

Notice that the `group_name` field displays `maintobkpfcopy` for both mappings.

Use the **`svcinfolsfconsistgrp`** command with the name of the consistency group to display the detailed attributes of the group. This now includes a list of the IDs and names of the mappings that are in the group.

```
svcinfolsfconsistgrp -delim : maintobkpfcopy
id:1
name:maintobkpfcopy
status:idle_copied
FC_mapping_id:0
FC_mapping_name:main1copy
FC_mapping_id:1
FC_mapping_name:main2copy
```

#### Related topics:

- Chapter 21, “Scenario: typical usage for the command-line interface,” on page 153

---

## Preparing and triggering a FlashCopy Consistency Group using the CLI

Starting the FlashCopy process will create a point-in-time copy of the data on the source VDisk and write it to the target VDisk for each mapping in the group.

#### Steps:

Perform the following steps to prepare and trigger a FlashCopy consistency group:

1. Open a command prompt window.
2. Issue the **`svctask prestartfcconsistgrp`** command to prepare the FlashCopy Consistency Group before the copy process can be started (triggered). When you have assigned several mappings to a FlashCopy Consistency Group, you only have to issue a single prepare command for the whole group, to prepare all the mappings at once.

#### Example:

In our hypothetical scenario, the command to prepare a FlashCopy Consistency group called *maintobkpfcopy* is:

```
svctask prestartfcconsistgrp maintobkpfcopy
```

The group will enter the preparing state, and then move to the prepared state when it is ready. Issue the **`svcinfolsfconsistgrp`** command to check:

```
svcinfolsfconsistgrp -delim :
id:name:status
1:maintobkpfcopy:prepared
```

3. Issue the **`svctask startfcconsistgrp`** command to start (trigger) the FlashCopy Consistency Group to make the copy. You only have to issue a single start command for the whole group, to trigger all the mappings at once.

**Example:**

In our hypothetical scenario, the command to trigger a FlashCopy Consistency group called *maintobkpfcopy* is:

```
svctask startfcconsistgrp maintobkpfcopy
```

The group will enter the copying state, and then return to the *idle\_copied* state when complete. You can issue the **svcinfc lsfcconsistgrp** command to check the state of the group:

```
svcinfc lsfcconsistgrp -delim :
id:name:state
1:maintobkpfcopy:copying
```

Use the **svcinfc lsfcmapprogress** command to check the progress of each mapping, *main1copy* and *main2copy*:

```
svcinfc lsfcmapprogress -delim : main1copy
id:progress
0:100
```

```
svcinfc lsfcmapprogress -delim : main2copy
id:progress
1:23
```

Finally issue the **svcinfc lsfcconsistgrp** command to display the detailed view of the group *maintobkpfcopy*, which returns to *idle\_copied* state when both mappings have reached 100% progress:

```
svcinfc lsfcconsistgrp -delim : maintobkpfcopy
id:1
name:maintobkpfcopy
state:idle_copied
FC_mapping_id:0
FC_mapping_name:main1copy
FC_mapping_id:1
FC_mapping_name:main2copy
```

You have now made a point-in-time copy of the data on *mainvdisk1* which has been written to *bkpvdisk1*, and a copy of the data on *mainvdisk2* which has been written to *bkpvdisk2*. The data on *bkpvdisk1* and *bkpvdisk2* will be visible to *demohost2* because these VDisks are only mapped to *demohost2*.

**Related topics:**

- Chapter 21, “Scenario: typical usage for the command-line interface,” on page 153

---

## Chapter 22. Advanced functions with the CLI

---

### Determining a nodes WWPNs using the CLI

**Steps:**

Perform the following steps to determine a nodes WWPNs:

1. List the nodes in the cluster by issuing the following command:

```
svcinfolnode
```

**Note:** Remember the node name or ID as you will need it in the next step.

2. For the node or nodes in question, issue the following command:

```
svcinfolnode <nodename/id>
```

where *<nodename/id>* is the node name or ID.

**Note:** Remember the four port ID's (WWPNs).

---

### Determining a storage subsystem name from its SAN Volume Controller Storage Software for Cisco MDS 9000 name using the CLI

**Steps:**

Perform the following steps to determine a storage subsystem name:

1. List the storage subsystem by issuing the following command:

```
svcinfolcontroller
```

Remember the name or ID for the storage subsystem you want to determine.

2. For the device in question, issue the following command:

```
svcinfolcontroller <controllername/id>
```

where *<controllername/id>* is the name or ID. Remember the WWNN for the device. Make a written record of it. The WWNN can be used to determine the actual storage subsystem by launching the native user interface or using the command line tools it provides to verify the actual storage subsystem that has this WWNN.

---

### Determining the VDisk name from the vpath number on the host

Each VDisk exported by the SAN Volume Controller Storage Software for Cisco MDS 9000 is assigned a unique vpath number. This number uniquely identifies the VDisk and can be used to determine which VDisk corresponds to the volume that the hosts sees. This procedure can only be performed using the command line interface.

**Steps:**

Perform the following steps to determine the VDisk name from the vpath number:

1. For the volume in question, find the vpath serial number by issuing the following command:

```
datapath query device
```

2. Find the host object defined to the SAN Volume Controller Storage Software for Cisco MDS 9000 that corresponds with the host you are working with.
  - a. The WWPNs are an attribute of the HBA. You can find these by looking at the device definitions stored by your operating system. For example, on AIX they will be in the ODM, in Windows they will be in the Device Manager details for the given HBA.
  - b. Verify which host object defined to the SAN Volume Controller Storage Software for Cisco MDS 9000 that these ports belong to. The ports are stored as part of the detailed view, so you will need to list each host in turn by issuing the following:

```
svcinfolshost <name/id>
```

where *<name/id>* is the name or ID of the host. Check for matching WWPNs.

**Note:** You should name your hosts accordingly, for example, if the actual host is called *orange* you should also name the host object defined to the SAN Volume Controller Storage Software for Cisco MDS 9000 as *orange*.

3. Now that you have the *<host name>* as defined to the SAN Volume Controller Storage Software for Cisco MDS 9000 and the *<vpath serial number>*, issue the following command:

```
svcinfolshostvdiskmap <hostname>
```

where *<hostname>* is the name of the host. A list is displayed.

4. Look for the VDisk UID that matches the *<vpath serial number>* and remember the VDisk name or ID.

---

## Determining the host that a VDisk is mapped to

### Steps:

Perform the following steps to determine the host that the VDisk is mapped to:

1. Find the VDisk name or ID that you wish to check.
2. List the hosts that this VDisk is mapped, by issuing the following command:

```
svcinfolsvdiskhostmap <vdiskname/id>
```

where *<vdiskname/id>* is the name or ID of the VDisk. A list is displayed.

3. Look for the host name or ID to determine which host this VDisk is mapped to. If no data is returned, the VDisk is not mapped to any hosts.

---

## Determining the relationship between VDIs and MDIs using the CLI

Every VDisk is constructed from one or more MDIs. At times you may need to determine the relationship between the two objects. The following procedure allows you to determine the relationships.

### Steps:

Perform the following steps to determine the relationship between VDIs and MDIs:

1. For a given VDisk `<vdiskname/id>`, issue the following command:

```
svcinfolsvdiskmember <vdiskname/id>
```

where `<vdiskname/id>` is the name or ID of the VDisk. This will return a list of IDs that correspond to the MDIs that make up the VDisk.

### Steps:

Perform the following steps to determine the relationship between VDIs and MDIs and the number of extents provided by each MDI:

If you wish more details, you can also determine the number of extents that are being provided by each MDI. This procedure can only be performed using the command line interface.

1. For a given VDisk `<vdiskname/id>`, issue the following command:

```
svcinfolsvdiskextent <vdiskname/id>
```

where `<vdiskname/id>` is the name or ID of the VDisk. This will return a table of MDI IDs and the corresponding number of extents each MDI is providing as storage for the given VDisk.

### Steps:

Perform the following steps to determine the relationship between MDIs and VDIs:

1. For a given MDI `<mdiskname/id>`, issue the following command:

```
svcinfolsmdiskmember <mdiskname/id>
```

where `<mdiskname/id>` is the name or ID of the MDI. This will return a list of IDs that correspond to the VDIs that are using this MDI.

### Steps:

Perform the following steps to determine the relationship between MDIs and VDIs and the number of extents used by each VDisk:

If you wish more details, you can also determine the number of extents that this MDI is providing for each VDisk. This procedure can only be performed using the command line interface.

1. For a given MDI `<mdiskname/id>`, issue the following command:

```
svcinfo lsmdiskextent <mdiskname/id>
```

where *<mdiskname/id>* is the name or ID of the MDisk. This returns a table of VDisk IDs and the corresponding number of extents being used by each VDisk.

---

## Determining the relationship between MDisks and RAID arrays or LUNs using the CLI

Each MDisk corresponds with a single RAID array, or a single partition on a given RAID array. Each RAID controller will define a LUN number for this disk. The LUN number and controller name or ID are needed to be able to determine the relationship between MDisks and RAID arrays or partitions.

### Steps:

Perform the following steps to determine the relationship between MDisks and RAID arrays:

1. Show the detailed view of the given MDisk *<mdiskname>*, by issuing the following command:

```
svcinfo lsmdisk <mdiskname>
```

where *<mdiskname>* is the name of the MDisk.

**Note:** Remember the controller name or controller ID and controller LUN number.

2. Show the detailed view of the controller determined in by issuing the following command:

```
svcinfo lscontroller <controllername>
```

where *<controllername>* is the name of the controller.

**Note:** Remember the vendor ID, product ID, and WWNN. Use these to determine what is being presented to the MDisk.

3. From the native user interface for the given controller, list the LUNs it is presenting and match the LUN number with that noted in 1. This will tell you the exact RAID array or partition that corresponds with the MDisk.

---

## Increasing the size of your cluster using the CLI

To increase the size of your cluster you need to add nodes in pairs to a new I/O group. Your existing cluster may have become a bottleneck and so you wish to increase throughput by adding more nodes to the cluster.

### Steps:

Perform the following steps to increase the size of your cluster:

1. Perform the steps in the section and repeat this procedure for the second node.
2. If you wish to balance the load between the existing I/O groups and the new I/O groups, follow the procedure. Repeat this procedure for all VDIsks you want to assign to the new I/O group.

### Related topics:



- “Increasing the size of your cluster using the CLI” on page 174
- “Adding a node to increase the size of your cluster using the CLI”

## Adding a node to increase the size of your cluster using the CLI

### Steps:

Perform the following steps to add a node to increase the size of your cluster:

1. Issue the following command to verify that the node can be seen on the fabric:

```
svcinfolnodecandidate
```

You should see the node listed as a candidate.

**Note:** Remember the WWNN's. You will need it in the following step.

2. Issue the following command to determine the I/O group you wish to add the nodes to:

```
svcinfolsiogrp
```

3. Select the first I/O group listed that has a node count = 0.

**Note:** Remember the I/O group name or ID. You will need it in the following step.

4. **Attention:** If you are re-adding a node to the SAN, ensure that you are adding the node to the same I/O group that it was removed from. Failure to do this can result in data corruption. Use the information that was recorded when the node was originally added to the cluster. If you do not have access to this information, call IBM Service to add the node back into the cluster without corrupting the data.

If you are adding the node into the cluster for the first time, record the following information:

- Card serial number
- Slot in which the Cisco MDS 9000 Caching Services Module is located
- Node (1 or 2) on the Cisco MDS 9000 Caching Services Module
- WWNN
- All WWPNS
- I/O group that contains the node

This can avoid a possible data corruption exposure if the node must be removed from and re-added to the cluster.

**Note:** This warning also is displayed on the Console for SAN Volume Controller for Cisco MDS 9000 panel when adding the node.

Issue the following command to add the node into the cluster. The <newnodename> is the name you wish to assign to this node.

```
svctask addnode -wwnodename <WWNN> -iogrp <newiogrpname/id>
[-name <newnodename>]
```

5. Issue the following command to verify that the node is online:

```
svcinfolnode
```

You may also need to modify the configuration of your disk controllers. If your controller uses a mapping technique to present its RAID arrays or partitions to the cluster you will need to modify the port groups that belong to the cluster because the WWNN and WWPNN's of the node have changed.

**Related topics:**

- Chapter 29, “Configuring and servicing storage subsystems,” on page 217

## Migrating a VDisk to a new I/O group

You can migrate a VDisk to a new I/O group to manually balance the workload across the nodes in the cluster. You may end up with a pair of nodes that are overworked and another pair that are under-worked. Follow this procedure to migrate a single VDisk to a new I/O group. Repeat for other VDIs as required.

**Attention:** This is a disruptive procedure, access to the VDisk will be lost while you follow this procedure. Under no circumstances should VDIs be moved to an offline I/O group. You must ensure the I/O group is online before moving the VDIs to avoid data loss scenarios.

**Steps:**

Perform the following steps to migrate a single VDisk:

1. Quiesce all I/O operations for the VDisk. You may need to determine the hosts that are using this VDisk.
2. Before migrating the VDisk, it is essential that for each vpath presented by the VDisk you intend to move, the SDD configuration is updated to remove the vpaths in question. Failure to do this may result in data corruption. See *IBM TotalStorage Subsystem Device Driver: User's Guide* for details about how to dynamically re-configure SDD for the given host operating system.
3. Any FlashCopy mappings or Remote Copy relationships that use this VDisk should be stopped or deleted. Issue the following command, to check if the VDisk is part of a relationship or mapping:

```
svcinfolsvdisk <vdiskname/id>
```

where <vdiskname/id> is the name or ID of the VDisk.

4. Look for the **FC\_id** and **RC\_id** fields. If these are not blank then the VDisk is part of a mapping or relationship. *IBM TotalStorage Subsystem Device Driver: User's Guide* for details on how to stop or delete the mapping or relationship.
5. Issue the following command to migrate the VDisk:

```
svctask chvdisk -iogrp <newiogrpname/id> <vdiskname/id>
```

6. Follow the procedure to discover the new vpaths and to check that each vpath is now presenting the correct number of paths. See the *IBM TotalStorage Subsystem Device Driver: User's Guide* for details on how to dynamically re-configure SDD for the given host operating system.

**Related topics:**

- “Determining the host that a VDisk is mapped to” on page 172
- “Determining the relationship between VDIs and MDIs using the CLI” on page 173
- “Advanced function FlashCopy and Remote Copy overview for CLI” on page 194

---

## Replacing a faulty node in the cluster using the CLI

**Note:** You will need to use the information that was recorded when the node was originally added to the cluster. If you do not have access to this information, call IBM Service to add the node back into the cluster without corrupting the data.

If you are adding the node into the cluster for the first time, record the following information:

- Card serial number
- Slot in which the Cisco MDS 9000 Caching Services Module is located
- Node (1 or 2) on the Cisco MDS 9000 Caching Services Module
- WWNN
- All WWPNs
- I/O group that contains the node

**Notes:**

1. This can avoid a possible corruption of data if the node must be removed from and re-added to the cluster. This warning also is displayed on the Console for SAN Volume Controller for Cisco MDS 9000 panel when adding the node.
2. The WWPN assignments are done through the slot in the switch. If a node is replaced in the same slot, a reassignment of the WWNN and WWPN is needed. IF a node is placed in a different slot, the WWNN and WWPN can be reassigned through a switch command.

Perform the following steps to replace a faulty node in the cluster:

1. Verify the name or ID of the node you want to remove by issuing the **svcinfo lsnode** command.
2. Select one of the following actions based on why you are replacing the node:
  - a. If the node was faulty, it will be shown as `offline`.
    - 1) Ensure that the partner node in the I/O group is online.
    - 2) If the other node in the I/O group is offline, start the directed maintenance procedures to determine the fault.
  - b. If you have been directed here by the directed maintenance procedures because the partner node in the I/O group has subsequently failed, follow the procedure for recovering from offline VDisks after a node or an I/O group failed.
  - c. If you are replacing the node for other reasons:
    - 1) Determine which node you want to replace.
    - 2) Ensure that the partner node in the I/O group is online.

**Note:** If the partner node is offline and you continue, you will lose access to the VDisks that belong to this I/O group. Start the directed maintenance procedures, and fix the other node before proceeding.

3. Remove the node from the cluster by issuing the **svctask rmnode <nodename/id>** command. (First, determine if you need to shut down the node. If this is the last node in the cluster, you do not need to shut down the node. Update the SDD configuration for each virtual path (vpath) that is presented by the VDisks that you intend to remove. Updating the SDD configuration removes the vpaths from the VDisks. Failure to update the configuration can result in

data corruption. See the *IBM TotalStorage Subsystem Device Driver: User's Guide* for details about how to dynamically reconfigure SDD for the given host operating system. If this is *not* the last node in the cluster, shut down the node you intend to remove using the stopcluster node command.)

4. Verify that the node can be seen on the fabric by issuing the **svcinfo lsnodecandidate** command. You should see the node listed as a candidate.
5. Make a note of the WWNNs for each node.
6. If the node was repaired by replacing it with another node, the WWNN for the node will change. You can modify the WWNN of the replacement node so that it matches that of the old node. To change the WWNN and WWPN so they match the old node, see your *Cisco MDS 9000 Family Command Reference*. After a node has been replaced, the service procedures guide you through changing the WWNN. If this is the case, complete the following additional steps:
  - a. At the end of the recovery process, follow the SDD procedure to discover the new paths and to check that each vpath is now presenting the correct number of paths. See the *IBM TotalStorage Subsystem Device Driver: User's Guide* sections on dynamic reconfiguration, specifically adding paths to existing vpaths.
  - b. You might also need to modify the configuration of your disk controllers. If your controller uses a mapping technique to present its RAID's or partitions to the cluster, you need to modify the port groups that belong to the cluster because the WWNN or WWPNs of the node have changed.
7. Issue the following command to add the node back into the cluster:
8. Issue the following command to verify that the node is online:

```
svcinfo lsnode
```

#### Related topics:

- Chapter 29, "Configuring and servicing storage subsystems," on page 217
- "Deleting a node from a cluster using the IBM TotalStorage CLI" on page 194
- "Adding nodes to a cluster using the IBM TotalStorage CLI" on page 154
- "Recovering from offline VDisks after a node or an I/O group failed using the CLI"

---

## Recovering from offline VDisks after a node or an I/O group failed using the CLI

### Prerequisites:

If you have lost both nodes in an I/O group and have therefore, lost access to all the VDisks that are associated with the I/O group, then you must perform one of the following procedures to regain access to your VDisks. Depending on the failure type, you may have lost data that was cached for these VDisks, therefore, they have gone offline.

### Context:

**Data loss scenario 1** One node in an I/O group failed and failover started on the second node. During this time, the second node in the I/O group fails before the cache has become write-through mode. The first node is successfully repaired but

its cache data is stale, therefore, it cannot be used. The second node is repaired or replaced and has lost its hard-end data, therefore, the node has no way of recognizing that it is part of the cluster.

#### Steps:

Perform the following steps to recover from an offline VDisk:

1. Recover the node and include it back into the cluster.
2. Move all the offline VDIs to the recovery I/O group.
3. Move all the offline VDIs back to their original I/O group.

#### Context:

**Data loss scenario 2** Both nodes in the I/O group have failed and have been repaired. The nodes have lost their hard-end data, therefore, the nodes have no way of recognizing that they are part of the cluster.

1. Move all the offline VDIs to the recovery I/O group
2. Move both recovered nodes back into the cluster
3. Move all the offline VDIs back to their original I/O group.

#### Related topics:

- “Recovering a node and including it back into the cluster”
- “Moving offline VDIs to the recovery I/O group” on page 180
- “Moving offline VDIs to their original I/O group using the CLI” on page 181

## Recovering a node and including it back into the cluster

After a node or an I/O group fails, you can use the following procedure to recover a node and include it back into the cluster.

#### Steps:

Perform the following steps to recover a node and include it back into the cluster:

1. Verify that the node is offline. Issue the following command:

```
svcinfolnode
```

2. Remove the old instance of the offline node from the cluster. Issue the following command:

```
svctask rmnode <nodename/id>
```

where <NODENAME> is the name of the node.

3. Verify that the node can be seen on the fabric. Issue the following command:

```
svcinfolnodecandidate
```

You should see the nodes listed as a candidate.

**Note:** Remember the WWNNs for each node, you will need it in the following step.

4. If the node is repaired by moving the CSM to a different slot, then the WWNN for the node will change. In this case, the following additional steps are required:

- a. At the end of the recovery process it will be necessary to follow the SDD procedure to discover the new paths and to check that each vpath is now presenting the correct number of paths. See the *IBM TotalStorage Subsystem Device Driver: User's Guide* sections on dynamic reconfiguration, specifically adding paths to existing vpaths.
- b. You may also need to modify the configuration of your disk controllers. If your controller uses a mapping technique to present its RAID arrays or partitions to the cluster you will need to modify the port groups that belong to the cluster because the WWNN or WWPN's of the node have changed.

**Attention:** If more than one I/O group is affected, ensure that you are adding the node to the same I/O group that it was removed from. Failure to do this can result in data corruption. Use the information that was recorded when the node was originally added to the cluster. This can avoid a possible data corruption exposure if the node must be removed from and re-added to the cluster. If you do not have access to this information, call IBM Service to add the node back into the cluster without corrupting the data.

**Attention:** If you are adding the node into the cluster for the first time, record the following information:

- Card serial number
- Slot in which the Cisco MDS 9000 Caching Services Module is located
- Node (1 or 2) on the Cisco MDS 9000 Caching Services Module
- WWNN
- All WWPNs
- I/O group that contains the node

**Note:** This warning also is displayed on the Console for SAN Volume Controller for Cisco MDS 9000 panel when adding the node.

5. Add the node back into the cluster. Issue the following command:

```
svctask addnode -wwnodename <WWNN> -iogrp
<IOGRPNAME/ID> [-name <NODENAME>]
```

where <WWNN> is the worldwide node name <IOGRPNAME/ID> is the I/O group name or ID.

6. Verify that the node is online. Issue the following command:

```
svcinfolnode
```

**Related topics:**

- “Recovering from offline VDIsks after a node or an I/O group failed using the CLI” on page 178
- “Moving offline VDIsks to the recovery I/O group”
- “Moving offline VDIsks to their original I/O group using the CLI” on page 181

## Moving offline VDIsks to the recovery I/O group

After a node or an I/O group fails, you can use the following procedure to move offline VDIsks to the recovery I/O group.

**Steps:**

Perform the following steps to move offline VDIs to the recovery I/O group:

**Notes:**

1. <IOGRPNAME> = the name of the I/O group that failed.
2. <vdiskname/ID> = the name of one of the VDIs that are offline.
1. List all VDIs that are offline and belong to the I/O group in question. Issue the following command:

```
svcinfolsvdisk -filtervalue IO_group_name=
<IOGRPNAME/ID>;status=offline
```

2. For each VDI returned, move the VDI to the recovery I/O group. Issue the following command:

```
svctask chvdisk -iogrp recovery_io_grp -force
<vdiskname/ID>
```

**Related topics:**

- “Recovering from offline VDIs after a node or an I/O group failed using the CLI” on page 178
- “Recovering a node and including it back into the cluster” on page 179
- “Moving offline VDIs to their original I/O group using the CLI”

## Moving offline VDIs to their original I/O group using the CLI

After a node or an I/O group fails, you can use the following procedure to move offline VDIs to their original I/O group.

**Attention:** Under no circumstances should VDIs be moved to an offline I/O group. Ensure the I/O group is online before moving the VDIs back to avoid any further data loss.

**Steps:**

Perform the following steps to move offline VDIs to their original I/O group:

**Notes:**

1. <IOGRPNAME> = the name of the I/O group that failed.
2. <vdiskname/ID> = the name of one of the VDIs that are offline.
1. For each VDI, move the VDI back into the original I/O group. Issue the following command:

```
svctask chvdisk -iogrp <IOGRPNAME/ID> -force
<vdiskname/ID>
```

2. Verify that the VDIs are now online. Issue the following command:

```
svcinfolsvdisk -filtervalue IO_group_name=
<IOGRPNAME/ID>
```

**Related topics:**

- “Recovering from offline VDIs after a node or an I/O group failed using the CLI” on page 178
- “Recovering a node and including it back into the cluster” on page 179
- “Moving offline VDIs to the recovery I/O group” on page 180

---

## Replacing an HBA in a host using the CLI

This procedure describes how to notify the SAN Volume Controller Storage Software for Cisco MDS 9000 of a change to a defined host object. It is sometimes necessary to replace the HBA that connects the host to the SAN, at this time you must notify the SAN Volume Controller Storage Software for Cisco MDS 9000 of the new WWPN's that this HBA contains.

### Prerequisites:

Ensure your switch is zoned correctly.

### Steps:

Perform the following steps to replace an HBA in a host using the CLI:

1. Issue the following command to list the candidate HBA ports:

```
svcinfolshbaportcandidate
```

You should see a list of the HBA ports that are available to be added to host objects. One or more of these should correspond with the one or more WWPNs that belong to the new HBA.

2. Locate the host object that corresponds with the host in which you have replaced the HBA. The following command lists all the defined host objects:

```
svcinfolshost
```

To list the WWPNs currently assigned to the host, issue the following:

```
svcinfolshost <hostobjectname>
```

where *<hostobjectname>* is the name of the host object.

3. Add the new ports to the existing host object by issuing the following command:

```
svctask addhostport -hbawwpn <one or more existing WWPNs separated by :> <hostobjectname/ID>
```

where *<one or more existing WWPNs separated by :>* correspond with those listed in step 1 and *<hostobjectname/id>* corresponds with the host object located in step 2.

4. Remove the old ports from the host object by issuing the following command:

```
svctask rmhostport -hbawwpn <one or more existing WWPNs separated by :> <hostobjectname/ID>
```

where *<one or more existing WWPNs separated by :>* correspond with those listed in step 2 that belong to the old HBA that has been replaced.

5. Any mappings that exist between the host object and VDisks will automatically be applied to the new WWPNs. Therefore, the host should see the VDisks as the same SCSI LUNs as before.
6. See the *IBM TotalStorage Subsystem Device Driver: User's Guide* for additional information about dynamic reconfiguration.

### Related topics:



---

## Adding a new storage controller to a running configuration using the CLI

### Prerequisites:

You can add a new disk controller system to your SAN at any time. Follow the switch zoning guidelines in the section about switch zoning. Also, ensure the controller is setup correctly for use with the SAN Volume Controller Storage Software for Cisco MDS 9000 .

You should create one or more arrays on the new controller. It is recommend that you use, RAID-5, RAID-1 or RAID-0+1 (sometimes called RAID-10) for maximum redundancy and reliability. Generally 5+P arrays are recommend. If your controller provides array partitioning we recommend that you create a single partition from the entire capacity available in the array, remember the LUN number that you assign to each partition as you will need this later. You should also follow the mapping guidelines (if your disk controller system requires LUN mapping) to map the partitions or arrays to the SAN Volume Controller Storage Software for Cisco MDS 9000 ports.

### Steps:

Perform the following steps to add a new disk controller system to a running configuration:

1. To ensure that the cluster has detected the new storage (MDisks) issue the following command:

```
svctask detectmdisk
```

2. The controller itself will have automatically been assigned a default name. If you are unsure which controller is presenting the MDisks, list the controllers by issuing the following command:

```
svcinfolcontroller
```

You should see a new controller listed (the one with the highest numbered default name). Remember the controller name and follow the instructions in the section about determining a disk controller system name.

3. You should give this controller a name that you can easily use to identify it. Issue the following command:

```
svctask chcontroller -name <newname> <oldname>
```

4. List the unmanaged MDisks by issuing the following command:

```
svcinfolmdisk -filtervalue mode=unmanaged:controller_name=<new_name>
```

These MDisks should correspond with the RAID arrays or partitions you have created. Remember the field controller LUN number. This corresponds with the LUN number you assigned to each of the arrays or partitions.

5. It is recommended that you create a new managed disk group and add only the RAID arrays that belong to the new controller to this MDisk group. You should also avoid mixing RAID types, so for each set of RAID array types (for example, RAID-5, RAID-1) you should create a new MDisk group. Give this MDisk group

an appropriate name, so if your controller is called FAST650-fred, and the MDisk group contains RAID-5 arrays, call it something like F600-fred-R5). Issue the following command:

```
svctask mkmdiskgrp -ext 16 -name <mdisk_grp_name>
-mdisk <colon separated list of RAID-x mdisks returned
in step 4.
```

**Note:** This will create a new MDisk group with an extent size of 16MB.

**Related topics:**

- “Determining a nodes WWPNs using the CLI” on page 171
- “Determining a storage subsystem name from its SAN Volume Controller Storage Software for Cisco MDS 9000 name using the CLI” on page 171
- Chapter 29, “Configuring and servicing storage subsystems,” on page 217

---

## Removing a storage subsystem using the CLI

You can replace or decommission an old storage subsystem by following the procedure below. This procedure takes you through adding the new device, migrating the data off of the old device and removing the old MDisks.

This function can also be performed by migrating all the VDIs that are using storage in this MDisk group to another MDisk group. This procedure has an advantage if you wish to consolidate the VDIs in a single or new group. However, you can only migrate a single VDI at a time. The procedure outlined below will migrate all the data through a single command.

This procedure can also be used to remove or replace a single MDisk in a group. If an MDisk has suffered a partial failure, such as a degraded array, and you can still read the data from the disk but cannot write to it, you can follow this procedure to replace just one MDisk. In steps 1 and 3 on page 185 only add or remove a single MDisk rather than a list of MDisks.

**Prerequisites:**

All the MDisks that belong to the storage subsystem that is being decommissioned belong to a single MDisk group. You need to repeat this procedure for each MDisk group in turn before removing the old device.

**Steps:**

Perform the following steps to remove a storage subsystem:

1. Add the new storage subsystem to your cluster configuration.
2. Issue the following command:

```
svctask addmdisk -mdisk <colon separated mdisk
list as determined in step 4> <mdisk_grp_name>
```

Where *<mdisk\_grp\_name>* is the name of the MDisk group that contains the MDisks that are being decommissioned.

3. You should now have an MDisk group that contains the old MDisks (those to be decommissioned) and the new MDisks (those that are replacing them). Ensure that the capacity of the new MDisks is the same or exceeds that of the old MDisks before proceeding.
4. Force delete the old MDisks from the group. This will migrate all the data from the old MDisks to the new MDisks. Issue the following command:

```
svctask rmmdisk -force -mdisk <colon separated
mdisk list of all the old mdisks> <mdisk_grp_name>
```

Depending upon the number and size of the MDisks, and the number and size of the VDisks that are using these MDisks, this operation will take some time to complete although the command will return immediately.

5. Check progress by issuing the following command:

```
svcinfolismigrate
```

6. When all the migration tasks have completed, for example, the command in step 4 returns no output, you can safely remove the old device from the SAN.
7. Once you have removed the old device from the SAN, rerun the **svctask detectmdisk** command to remove the entries for the old MDisks.

#### Related topics:

- “Adding a new storage controller to a running configuration using the CLI” on page 183
- “Migrating VDisks between MDisk groups using the CLI” on page 191
- “Removing a storage subsystem using the CLI” on page 184
- “Determining the relationship between VDisks and MDisks using the CLI” on page 173

---

## Expanding VDisks

You can expand a VDisk from the Expanding VDisks panel or by using the command-line interface (CLI). This topic lists the supported platforms and requirements if this feature is to be used

A VDisk that is not yet mapped to any hosts and hence does not contain any customer data can be expanded at any time. However, if the VDisk contains data that is being used, only AIX and Windows 2000 hosts can cope with a VDisk being expanded.

The following matrix shows the supported platforms and requirements if this feature is to be used:

*Table 13. Supported platforms and requirements*

| Platform     | Supported | Requirement          |
|--------------|-----------|----------------------|
| AIX          | Yes       | AIX 5.2 onwards only |
| HP-UX        | No        |                      |
| Linux        | No        |                      |
| SUN Solaris  | No        |                      |
| Windows NT   | No        |                      |
| Windows 2000 | Yes       |                      |

**Related topics:**

- “Virtual disks” on page 22
- “Expanding a Virtual disk that is mapped to an AIX host”
- “Expanding a Virtual disk that is mapped to a Windows 2000 host” on page 187

## Expanding a Virtual disk that is mapped to an AIX host

**Prerequisites:**

VDisks that are participating in Flash Copy mappings or in Remote Copy relationships cannot be expanded.

Determine the exact size of the source or master VDisk by issuing the following command-line interface (CLI) command:

```
svcinfolsvdisk -bytes <vdiskname>
```

**Context:**

This feature can be used in two ways:

- To increase the capacity available on a particular VDisk that is already mapped to a host.
- To increase the size of a VDisk so that it matches the size of the source or master VDisk and can be used in a FlashCopy mapping or Remote Copy relationship.

**Steps:**

Perform the following steps to expand a VDisk that is mapped to an AIX host:

1. Determine the VDisk you wish to expand and remember its <vdiskname>.
2. Verify that this VDisk is mapped to an AIX host.
3. Determine the volume group that contains the VDisk (you must know the VDisk to hdisk relationship)
4. Quiesce all I/O operations to **all** volumes that belong to the volume group and sync the file-systems mounted on this volume group.
5. Check the current type of the VDisk by viewing the VDisk details in the Work with VDIsks panel.

**Notes:**

- a. If the VDisk has a type of image, it cannot be expanded.
  - b. If the VDisk has a type of sequential, it becomes a striped VDisk when you expand it.
6. Deactivate the volume group that contains this VDisk. Issue the following command from the command prompt:  

```
varyoffvg <volume_group>
```
  7. Expand the VDisk using either of the following methods:
    - From the Work with VDIsks panel, select the VDisk and select the Expand task. Enter the capacity by which you wish to extend this VDisk and the select the appropriate units. Select one, more, or all of the MDisks from the list. These will be the MDisks that provide the extra capacity. Optionally, select the format check-box if you want this extra capacity to be formatted before use.
    - From the command prompt issue the following command:

```
svctask expandvdisksize
```

8. 8. Reactivate the volume group so that the change in size is detected by the HBA device driver. Issue the following command from the command prompt:  

```
varyonvg <volume_group>
```
9. Run the **change volume group** command to notify the LVM that the size has changed. Issue the following command from the command prompt:  

```
chvg -g <volume_group>
```
10. Expand all the file-systems that are mounted on this VDisk (or use the new capacity as required)

#### **Post-processing requirements:**

Restart I/O operations to the volume group.

#### **Related topics:**

- “Virtual disks” on page 22

## **Expanding a Virtual disk that is mapped to a Windows 2000 host**

#### **Prerequisites:**

Vdisks that are in participating in Flash Copy mappings or in Remote Copy relationships cannot be expanded.

Ensure that you have run Windows Update and have applied all recommended updates to your system prior to attempting to expand a VDisk that is mapped to a Windows 2000 host

Determine the exact size of the source or master VDisk by issuing the following command-line interface (CLI) command:

```
svcinfolsvdisk -bytes <vdiskname>
```

#### **Context:**

This feature can be used in two ways:

- To increase the capacity available on a particular VDisk that is already mapped to a host.
- To increase the size of a VDisk so that it matches the size of the source or master VDisk and can be used in a FlashCopy mapping or Remote Copy relationship.

VDisks can be expanded under Windows 2000 concurrently with I/O operations.

#### **Steps:**

Perform the following steps to expand a VDisk that is mapped to a Windows 2000 host:

1. Expand the VDisk using either of the following methods:
  - From the Work with VDisks panel, select the VDisk and select the Expand task. Enter the capacity by which you wish to extend this VDisk and the select the appropriate units. Select one, more, or all of the MDisks from the

list. These will be the MDisks that provide the extra capacity. Optionally, select the format check-box if you want this extra capacity to be formatted before use.

- From the command prompt issue the following command:

```
svctask expandvdisksize
```

2. On the Windows Host, start the Computer Management application and open the Disk Management window under the Storage branch.

### Result:

You will see the VDisk that you expanded now has some unallocated space at the end of the disk.

Dynamic disks can be expanded without stopping I/O operations in most cases. However, in some applications the operating system may report I/O errors. When this problem occurs, either of the following entries may be recorded in the System event log:

```
Event Type: Information
Event Source: dmio
Event Category: None
Event ID: 31
Description: dmio:
Harddisk0 write error at block ##### due to
disk removal
```

```
Event Type: Information
Event Source: dmio
Event Category: None
Event ID: 34
Description: dmio:
Harddisk0 is re-online by PnP
```

**Attention:** This is a known problem with Windows 2000 and is documented at the Microsoft knowledge base as article Q327020. If either of these errors are seen, run Windows Update and apply the recommended fixes to resolve the problem.

### Post-processing requirements:

Restart the Computer Management application if it was opened prior to expanding the VDisk.

If the disk is a Windows basic disk you can create a new primary or extended partition from the unallocated space.

If the disk is a Windows dynamic disk you can use the unallocated space to create a new volume (simple, striped, mirrored) or add it to an existing volume.

### Related topics:

- “Virtual disks” on page 22

---

## Shrinking a VDisk using the CLI

VDisks can be reduced in size should it be required. However, if the VDisk contains data that is being used, **under no circumstances should you attempt to shrink a VDisk without first backing up your data.** The SAN Volume Controller Storage Software for Cisco MDS 9000 arbitrarily reduces the capacity of the VDisk by

removing a partial, one or more extents from those allocated to the VDisk. You cannot control which extents are removed and so you cannot guarantee that it is unused space that is removed.

**Attention:** This feature should *only* be used to make a target or auxiliary VDisk the same size as the source or master VDisk when creating FlashCopy mappings or Remote Copy relationships. You should also ensure that the target VDisk is not mapped to any hosts prior to performing this operation.

### Steps:

Perform the following steps to shrink a VDisk:

1. Validate that the VDisk is not mapped to any host objects. If the VDisk is mapped, data is displayed.
2. You can determine the exact capacity of the source or master VDisk. Issue the following command:

```
svcinfolsvdisk -bytes <vdiskname>
```

3. Shrink the VDisk by the required amount. Issue the following command:

```
svctask shrinkvdisksize -size <capacitytoshrinkby> -unit
<unitsforreduction> <vdiskname/ID>
```

### Related topics:

- “Determining the host that a VDisk is mapped to” on page 172

---

## Migrating extents using the CLI

The SAN Volume Controller Storage Software for Cisco MDS 9000 provides various data migration features. These can be used to move the placement of data both *within* MDisk groups and *between* MDisk groups. These features can be used concurrent with I/O operations. There are two ways in which you can migrate data:

1. Migrating data (extents) from one MDisk to another (within the same MDisk group). This can be used to remove hot or overutilized MDisks.
2. Migrating VDIs from one MDisk group to another. This can be used to remove hot MDisk groups, for example, reduce the utilization of a group of MDisks.

You can determine the usage of particular MDisks by gathering I/O statistics about MDisks and VDIs. Once you have gathered this data, you can analyze it to determine which MDisks are hot. The procedure then takes you through querying and migrating extents to elsewhere in the same MDisk group. This procedure can only be performed using the command line tools.

To migrate extents to remove possible problems, perform the following:

1. Isolate any MDisks that are overutilized. You can determine this by requesting an I/O statistics dump and analyzing the output. To start I/O statistics gathering, issue the following:

```
svctask startstats -interval 15
```

2. This will generate a new I/O statistics dump file approximately every 15 minutes. Wait for at least 15 minutes after issuing the **svctask startstats** command and then issue the following:

```
svcinfolsiostatsdumps
```

This will list the I/O statistics files that have been generated. These are prefixed with `m` and `Nm` for MDisk statistics and `v` for VDisk statistics.

3. Use secure copy (`scp`) to retrieve the dumps files to analyze. For example, issue the following:

```
<AIX HOST PROMPT#>scp <clusterip>:/dumps/iostats/m_*
```

This will copy all the MDisk statistics files to the AIX host in the current directory.

4. Analyze the dumps to determine which MDisks are hot. It may be helpful to also determine which VDIsks are being heavily utilized as you can spread the data they contain more evenly across all the MDisks in the group using the procedure below.
5. Stop the statistics collection again by issuing the following command:

```
svctask stopstats
```

### Steps:

Once you have determined which MDisks are hot, you can migrate some of the data onto some less hot MDisks within the same MDisk group.

1. Determine the number of extents that are in use by each VDisk for the given MDisk. Issue the following command:

```
svcinfo lsmdiskextent <mdiskname>
```

This will return the number of extents that each VDisk is using on the given MDisk. You should pick some of these to migrate elsewhere in the group.

2. Determine the other MDisks that reside in the same MDisk group.
  - a. To determine the MDisk group that the MDisk belongs to, issue the following command:

```
svcinfo lsmdisk <mdiskname/ID>
```

Look for the `mdisk_grp_name` attribute.

- b. List the MDisks in the group by issuing the following command:

```
svcinfo lsmdisk -filtervalue mdisk_grp_name=<mdiskgrpname>
```

3. Select one of these MDisks as the target MDisk for the extents. You can determine how many free extents exist on an mdisk by issuing the following command:

```
svcinfo lsfreeextents <mdiskname>
```

You can issue the **`svcinfo lsmdiskextent <newmdiskname>`** command for each of the target MDisks to ensure that you are not just moving the over-utilization to another MDisk. Check that the VDisk that owns the set of extents to be moved, (see step1), does not already own a large set of extents on the target MDisk.

4. For each set of extents, issue the following command to move them to another MDisk:



```
svctask migrateextents -source <mdiskname/ID> -exts
<num_extents_from_step1> -target <newmdiskname/ID>
-threads 4 <vdiskid_returned_from_step1>
```

where *<num\_extents\_from\_step1>* is the number of extents on the *<vdiskid\_returned\_from\_step1>*, that is, the data that is returned from the command issued in step 1 on page 190. *<newmdiskname/ID>* is the name or ID of the MDisk to which you want to migrate this set of extents.

5. Repeat steps 2 on page 190 to 4 on page 190 for all the sets of extents you wish to move.
6. You can check the progress of the migration(s) by issuing the following command:

```
svcinfo lsmigrate
```

---

## Migrating VDIs between MDisk groups using the CLI

You can determine the usage of particular MDisks by gathering I/O statistics about MDisks and VDIs. Once you have gathered this data, you can analyze it to determine which VDIs or MDisks are hot. This procedure then takes you through migrating VDIs from one MDisk group to another.

When a migrate command is issued, a check is made to ensure that the destination of the migrate has enough free extents to satisfy the command. If it does, the command proceeds, but will take some time to complete. During this time, it is possible for the free destination extents to be consumed by another process, for example, by creating a new VDisk in the destination MDisk group or by starting more migrate commands. In this scenario, when all the destination extents have been allocated the migration commands suspend and an error is logged (error id 020005). There are two methods for recovering from this situation:

1. Add additional MDisks to the target MDisk group. This provides additional extents in the group and allows the migrations to be restarted. You will need to mark the error as fixed in order to reattempt the migration.
2. Migrate one or more VDIs that are already created from the MDisk group to another group. This will free up extents in the group and allow the original migrations to be restarted (again by marking the error as fixed).

### Steps:

Perform the following steps to migrate VDIs between MDisk groups:

1. Isolate any VDIs that are overutilized. You can determine this by requesting an I/O statistics dump and analyzing the output. To start I/O statistics gathering, issue the following command:

```
svctask startstats -interval 15
```

2. This will generate a new I/O statistics dump file approximately every 15 minutes. Wait for at least 15 minutes after issuing the **svctask startstats** command and then issue the following command:

```
svcinfo lsiostatsdumps
```

This will list the I/O statistics files that have been generated. These are prefixed with *m* and *Nm* for MDisk statistics and *v* for VDisk statistics.

3. Use secure copy (scp) to retrieve the dumps files for analyzing. For example, issue the following:

```
<AIX HOST PROMPT#>scp <clusterip>:/dumps/iostats/v_*
```

This will copy all the VDisk statistics files to the AIX host in the current directory.

4. Analyze the dumps to determine which VDIs are hot. It may be helpful to also determine which MDisks are being heavily utilized as you can spread the data they contain more evenly across all the MDisks in the group by migrating the extents.
5. Stop the statistics collection again. Issue the following command:

```
svctask stopstats
```

Once you have analyzed the I/O statistics data, you can determine which VDIs are hot. You also need to determine which MDisk group you wish to move this VDisk to. Either create a new MDisk group or determine an existing group that is not yet over utilized. You can do this by checking the I/O statistics files generated above and ensuring that the MDisks or VDIs in the target MDisk group are less utilized than the source group.

6. After having determined which VDisk you wish to migrate, and the new MDisk group you wish to migrate it to, issue the following command:

```
svctask migratevdisk -vdisk <vdiskname/ID> -mdiskgrp
<newmdiskgrpname/ID> -threads 4
```

7. You can check the progress of the migration by issuing the following command:

```
svcinfo lsmigrate
```

**Related topics:**

- “Migrating extents using the CLI” on page 189

---

## Migrating a VDisk between I/O groups using the CLI

**Attention:** These migration tasks are disruptive, in that the cached data held within the cluster must first be written to disk, then the allocation of the VDisk can be changed.

Modifying the I/O group that services the virtual disk cannot be done concurrently with I/O operations. It also requires a re-scan at the host level to ensure that SDD gets notified that the allocation of the preferred node has changed and the ports by which the virtual disk is accessed has changed. This should only be done in the situation where one pair of nodes has become over utilized.

**Steps:**

Perform the following steps to migrate a VDisk between I/O groups:

1. Sync all file-systems that are mounted on the given virtual disk.
2. Stop all I/O operations to the virtual disk.
3. Type the following:

```
svctask chvdisk -iogrp <new_io_grp_name_or_id>
<vdisk>
```

4. Issue the SDD command to resync the VDisk to host mapping. See the *IBM TotalStorage Subsystem Device Driver: User's Guide* for more information.
5. Restart the I/O operations to the virtual disk.

---

## Creating an image mode VDisk from an unmanaged MDisk using the CLI

The SAN Volume Controller Storage Software for Cisco MDS 9000 enables you to import storage that contains existing data and continue to use this storage but make use of the advanced functions, such as, Copy Services, data migration, and the cache. These disks are known as image mode virtual disks.

Make sure you are aware of the following before converting your virtual disks:

1. Unmanaged disks that contain existing data cannot be differentiated from unmanaged disks that are blank. Therefore, it is vital that you control the introduction of these disks to the cluster. It is recommended that you introduce these disks one at a time. For example, map a single LUN from your RAID controller to the cluster and refresh the view of managed disks. The newly detected disk is displayed.
2. *Do not* add a unmanaged disk that contains existing data to a unmanaged disk group manually. If you do, the data will be lost. When you create an image mode virtual disk from this unmanaged disk, it will be automatically added to the unmanaged disk group. However, it will be added in such a way that the cluster can control how it is added to ensure the data is not lost.

Go to the following Web site for more information:

[www.ibm.com/storage/support/2062-2300](http://www.ibm.com/storage/support/2062-2300)

### Steps:

Perform the following steps to convert your virtual disk from image mode to unmanaged mode:

1. Map a single RAID array or LUN from your RAID controller to the cluster. You can do this either through a switch zoning or a RAID controller based on your host mappings.
2. Re-scan the list of unmanaged disks from the Console for SAN Volume Controller for Cisco MDS 9000. Issue the **svcinfolismdisk** command to list the available unmanaged disks.  
Optionally, if the new unmanaged disk is not listed you may need to run a fabric level discovery. Issue the **svctask detectmdisk** command to manually re-scan the fibre-channel network for any new unmanaged disks that might have been added.
3. Convert the unmanaged disk into an image mode virtual disk. Issue the **svctask mkvdisk** command to create an image mode virtual disk object. Once mapped to a host object, these virtual disks are seen as disk drives with which the host can perform I/O operations.
4. Map the new virtual disk to the hosts that were previously using the data that the MDisk contains. Issue the **svctask mkvdiskhostmap** command to create a new mapping between a virtual disk and a host. That is, the virtual disk is made accessible for I/O operations to the specified host.

If you wish to convert this virtual disk or managed disk to actually virtualize the storage, you can transform the image mode virtual disk into a striped virtual disk by

migrating the data on the image mode disk to other managed disks in another MDisk group. This procedure can be performed using the command-line interface (CLI). Issue the **svctask migratevdisk** command to migrate an entire virtual disk from one managed disk group to another managed disk group.

---

## Advanced function FlashCopy and Remote Copy overview for CLI

For detailed information about how to perform advanced FlashCopy and Remote Copy functions, go to the following Web site:

[www.ibm.com/redbooks](http://www.ibm.com/redbooks)

---

## Advanced function cluster overview using the CLI

### Overview:

The following sections details the advanced cluster functions that you can perform using the CLI.

## Deleting a node from a cluster using the IBM TotalStorage CLI

**Attention:** Before you delete a node from the cluster, quiesce all I/O operations that are destined for this node. Failure to do so can result in failed I/O operations that are reported to your host operating systems. On hosts that have multiple paths to both nodes in the I/O group and are running IBM Subsystem Device Driver (SDD), you do not have to quiesce I/O operations.

### Prerequisites:

Issue the Cisco switch command **svc-ibmcli** to enter a mode where you can issue the IBM TotalStorage CLI commands.

**Attention:** If you are deleting a single node, and the other node in the I/O group is online, be aware that the cache on the partner node will go into write-through mode and that you are exposed to a single point of failure should the partner node fail.

If you are deleting a node, and this is the last node in the I/O group, you will lose access to all VDIsks served by this I/O group. Ensure that all VDIsks are not being accessed or contain data that you wish to continue to access, or ensure that they have been migrated to a different, online I/O group.

1. Begin by determining the VDIsks that are still assigned to this I/O group:
  - a. Determine the VDIsks in question by requesting a filtered view of VDIsks where the filter attribute is the I/O group in question. This can be done using the following command:

```
svcinfolsvdisk -filtervalue IO_group_name=<name>
```

where <name> is the name of the I/O group in question.
  - b. Once you have a list of VDIsks, determine the hosts that they are mapped to by following the procedure called, Determining the hosts that a VDisk is mapped to.
  - c. Once you have determined the hosts and are sure that you do not wish to maintain access to these VDIsks proceed to 3 on page 195.

- d. If you determine that some or all of the VDisks assigned to this I/O group do contain data that you wish to continue to access, you should follow the procedure called, Migrating a VDisk to a new I/O group.
2. Before performing the SDD path removal procedure described in 3 you should stop the node that you intend to remove using the **svctask stopcluster** command, unless this is the last node in the cluster. This ensures that SDD does not rediscover the paths that are manually removed before you issue the delete node request.
 

**Attention:** If the node being removed is the configuration node, it may take a minute or so before you can perform the delete node request as the configuration node failover has to occur. If the node being removed is the last node in the cluster, the Console for SAN Volume Controller for Cisco MDS 9000 may appear to hang for up to 3 minutes because you have removed the last access point to the cluster. Removing the last node in the cluster destroys the cluster. Ensure that this is what you want to do before performing this task.

**Note:** If you power back on the node that has been removed and it is still connected to the same fabric or zone it will attempt to rejoin the cluster. At this point the cluster will tell the node to remove itself from the cluster and the node will become a candidate for addition to this cluster or another cluster. If you are adding this node back into the cluster, ensure that you add it back to the same I/O group that it was previously a member of. Failure to do so may result in data corruption.
  3. Before deleting the node, it is essential that for each vpath presented by the VDisks you intend to remove, the SDD configuration is updated to remove the vpaths in question. Failure to do this may result in data corruption. See the *IBM TotalStorage Subsystem Device Driver: User's Guide* for details about how to dynamically re-configure SDD for the given host operating system.

### Steps:

Perform the following steps to delete a node:

1. Access the IBM SAN Volume Controller for Cisco MDS 9000 command line interface.

#### Notes:

- a. Before removing a node, be sure that this is what you want to do. All of the VDisks in the node will be assigned to the remaining node in the I/O group. That is, the preferred node will be changed. You cannot change this setting once this is done. Also, all VDisks will go into write-through cache mode because there is not a redundant node available to duplicate the cached information.
  - b. If this is the last node in the I/O group or the last node in the cluster, you will be prompted to force the deletion.
  - c. If this is the last node in the cluster or it is the configuration node, all connections to the cluster will be lost. The user interface and any open CLI sessions will become unavailable. This can result in a timeout from the command, because the command cannot be completed before the node is deleted.
2. Enter **svc-ibmcli cluster-name SampCluster**, where **SampCluster** is the name of the cluster.
  3. Issue the **svctask rmnode** command to delete a node from the cluster. You can enter this command any time after a cluster has been created.

## Performing the cluster maintenance procedure using the CLI

### Steps:

Perform the following steps for cluster maintenance:

1. Open a command prompt window.
2. Issue the **svctask finderr** command to analyze the error log for the highest severity of unfixed errors. This command scans the error log for any unfixed errors. Given a priority ordering defined within the code, the highest priority of unfixed errors is returned.
3. Issue the **svctask dumperrlog** command to dump the contents of the error log to a text file.
4. Locate the error and fix.
5. Issue the **svctask clearerrlog** command to clear all entries from the error log including status events and any unfixed errors.

**Note:** Clearing the error log will not fix the errors.

**Attention:** You should only use this command when you have either rebuilt the cluster, or have fixed a major problem that has caused many entries in the error log that you do not want to fix individually.

6. Issue the **svctask cherrstate** command to change the state of an error. The state can be changed from unfixed to fixed, or fixed to unfixed.

## Setting up error notifications using the CLI

### Steps:

Perform the following steps to set up error notifications:

1. Open a command prompt window.
2. Issue the **svctask setevent** command to specify what you like to happen when an error or event is logged to the error log. You can select whether the cluster raises an SNMP trap, issues an e-mail notification for entries that are added to the cluster error or event log, or both. Three levels of notification are possible:
  - **None** No error or status changes will be sent.
  - **Hardware\_only** You will be notified of errors, but you will not be notified of status changes.
  - **All** You will be notified of all errors and status changes.

If you have an SNMP manager installed or if you want to be notified by e-mail of errors or events, you should enable error notification. The notification levels for SNMP and e-mail alerts can be set independently. If you choose **All** or **Hardware\_only** notification, you must select a destination for the notification.

## Listing log or dump files using the CLI

### Steps:

Perform the following steps to list log or dump files:

1. Open a command prompt window.
2. You can issue any of the following commands to list error log files:
  - **svcinfo lserrlogbydisk**

- **svcinfo lserrlogbymdiskgroup**
- **svcinfo lserrlogbyvdisk**
- **svcinfo lserrlogbyhost**
- **svcinfo lserrlogbynode**
- **svcinfo lserrlogbyiogrp**
- **svcinfo lserrlogbyfcconsistgrp**
- **svcinfo lserrlogbyfcmap**
- **svcinfo lserrlogbyrcconsistgrp**
- **svcinfo lserrlogbyrcrelationship**

These commands will list the error log by type. These commands will return a list of dumps in the appropriate directory. For example, issue the **svcinfo lserrlogbymdisk** command, displays the error log by MDisks.

You can display the whole log or filter the log so that only errors, events, or unfixed errors are displayed. In addition, you can request the output to be sorted either by error priority or by time. For error priority, the most serious errors are the lowest-numbered errors. They are, therefore, displayed first in the table. For time, either the older or the latest entry can be displayed first in the output.

3. You can issue any of the following command to list dump files.

- **svcinfo lsconfigdumps**
- **svcinfo lserrlogdumps**
- **svcinfo lsfeaturedumps**
- **svcinfo lsiostatsdumps**
- **svcinfo lsiotracedumps**
- **svcinfo lssoftwaredumps**
- **svcinfo ls2145dumps**

These commands will list the dump file by type. These commands will return a list of dumps in the appropriate directory. For example, issue the **svcinfo lsconfigdumps** command, a list of dumps for configurations will be stored in the /dumps/configs destination directory.

The software dump files contain dumps of the SAN Volume Controller Storage Software for Cisco MDS 9000 memory. Your service representative might ask for these dumps to debug problems. The software dumps are large files (approximately 300 MB).

## Viewing the feature log using the CLI

### Steps:

Perform the following steps to view the feature log:

1. Open a command prompt window.
2. Issue the **svcinfo lsfeaturedumps** command to return a list of dumps in the /dumps/feature destination directory. The feature log is maintained by the cluster. The feature log records events that are generated when license parameters are entered or when the current license settings have been breached.
3. Issue the **svcservicemodeinfo lsfeaturedumps** command to return a list of the files that exist of the type specified on the given node.

## Analyzing the error log using the CLI

### Steps:

Perform the following steps to analyze the error log:

1. Open a command prompt window.
2. You can issue any of the following commands to list error log files:
  - **svcinfo lserrlogbymdisk**
  - **svcinfo lserrlogbymdiskgroup**
  - **svcinfo lserrlogbyvdisk**
  - **svcinfo lserrlogbyhost**
  - **svcinfo lserrlogbynode**
  - **svcinfo lserrlogbyiogrp**
  - **svcinfo lserrlogbyfcconsistgrp**
  - **svcinfo lserrlogbyfcmap**
  - **svcinfo lserrlogbyrconsistgrp**
  - **svcinfo lserrlogbyrcrelationship**

These commands will list the error log by type. These commands will return a list of dumps in the appropriate directory. For example, issue the **svcinfo lserrlogbymdisk** command, displays the error log by MDisk.

You can display the whole log or filter the log so that only errors, events, or unfixed errors are displayed. In addition, you can request the output to be sorted either by error priority or by time. For error priority, the most serious errors are the lowest-numbered errors. They are, therefore, displayed first in the table. For time, either the older or the latest entry can be displayed first in the output.

## Shutting down a cluster or single node using the CLI

### Prerequisites:

If all input power to a SAN Volume Controller Storage Software for Cisco MDS 9000 cluster is to be removed for more than a few minutes, (for example, if the machine room power is to be shutdown for maintenance), it is important that the cluster is shutdown before the power is removed.

**Attention:** Before shutting down a node or the cluster you should quiesce all I/O operations that are destined for this node or cluster. Failure to do so may result in failed I/O operations being reported to your host operating systems.

**Attention:** If you are shutting down the entire cluster, you will lose access to all VDIs being provided by this cluster.

### Shutting down the cluster:

#### Steps:

Perform the following steps to shut down a cluster:

1. Begin the process of quiescing all I/O to the cluster by stopping the applications on your hosts that are using the VDIs provided by the cluster.



- a. If you are unsure which hosts are using the VDIs provided by the cluster, follow the procedure called, Determining the hosts that a VDisk is mapped to.
  - b. Repeat the previous step for all VDIs.
2. Open a command prompt window.
  3. When all I/O has been stopped, issue the **svctask stopcluster** to shut down a single node or the entire cluster in a controller manner. If you specify the node ID or node name, you can shut down a single node.

When you enter this command either a node ID or node name argument, the node in question is shut down. After the command completes, the other node in the I/O group de-stages the contents of its cache and goes into write-through mode until the power to the node is returned and the node rejoins the cluster.

**Attention:** If this is the last node in an I/O group, you will lose all access to the virtual disks in the I/O group. Before you enter this command, ensure that this is what you want to do. You must specify the force flag.

4. Close the SSH session if you are using SSH in interactive mode.

### Shutting down a single node:

**Attention:** If you are shutting down a single node, and the other node in the I/O group is online, be aware that the cache on the partner node will go into write-through mode and that you are exposed to a single point of failure should the partner node fail while this node is shut down. Proceed to 2.

**Attention:** If you are shutting down a single node, and this is the last node in the I/O group, you will lose access to all VDIs being served by this I/O group.

### Steps:

Perform the following steps to shut down a single node:

1. Begin the process of quiescing all I/O to the VDIs being served by this nodes I/O group.
  - a. Determine the VDIs in question by requesting a filtered view of VDIs where the filter attribute is the I/O group in question. This can be done using the following command:

```
svcinfolsvdisk -filtervalue IO_group_name=<name>
```

where <name> is the name of the I/O group in question.

- b. Once you have a list of VDIs, determine the hosts that these are mapped to by following the procedure called, Determining the hosts that a VDisk is mapped to.

2. When all I/O has been stopped issue the following command to shut down the node:

```
svctask stopcluster <nodename/ID>
```

where <nodename/ID> is the name or ID of the node that you want to shut down.

**Note:** If this is the last node in the I/O group you also need to specify the -force parameter. For example to force the shutdown of node1:

```
svctask stopcluster -force node1
```

### Related topics:

- “Determining the host that a VDisk is mapped to” on page 172

---

## Part 6. Software upgrade strategy

You can upgrade your software while your day-to-day operations are running. You must, however, expect performance to be degraded while the software is being installed.

**Note:** Applying a software update takes approximately one hour. This is in part due to the 30 minute delay which is inserted to allow the multi-pathing software to recover.

Software and microcode for the SAN Volume Controller Storage Software for Cisco MDS 9000 and its attached adapters is tested and released as a single package. The package number is increased each time a new release is made, although only some of the components might have changed.

If you are upgrading through more than one level; for example, from level 1 to level 3, under some circumstances, you might need to install an intermediate level. For example, if you are upgrading from level 1 to level 3, you might need to install level 2 before you install level 3. Details of any prerequisite levels are provided with the source files.

**Attention:** Applying a software upgrade while the node is in service mode results in deleting the node from the cluster. Status information stored within the node will be deleted, and this causes data loss if the cluster is dependent solely on this node.

When upgrading cluster software where the cluster participates in one or more intercluster relationships only one cluster should be upgraded at a time. That is, both clusters should not be upgraded concurrently. The software upgrade should be allowed to complete one cluster before it is started on the other cluster. If both clusters are upgraded concurrently, it may lead to a loss of synchronization. It may further lead to a loss of availability.

**Attention:** Ensure that you have no unfixed errors in the log, and that the Cluster Time/Date is correctly set.



---

## Chapter 23. Upgrading the SAN Volume Controller Storage Software for Cisco MDS 9000 firmware using the Console for SAN Volume Controller for Cisco MDS 9000

Firmware upgrade files can be quite large, if you experience problems when uploading upgrade files to the cluster reliably, you should disable proxies on the Web browser from where you will upload the file. This should also shorten the file upload time.

**Note:** If you disable proxies, you may not be able to connect to external Web sites. It is therefore advised that prior to disabling proxies, you make a record of your existing settings in case you need to restore access to other Web sites.

### Prerequisites:

If you are using Internet Explorer, perform the following:

1. Click on **Tools** in the menu.
2. Select **Internet Options** → **Connections** tab.
3. Click on **LAN Settings...** and ensure that the box marked **Use a proxy server** is unchecked. Click **OK** twice to accept the settings.

If you are using Netscape, perform the following:

1. Click on **Edit** in the menu.
2. Click on **Preferences....** Expand the Advanced section and select **Proxies**.
3. Select the radio button marked **Direct connection to the Internet**. Click **OK** to accept the settings.

**Note:** You cannot download the upgrade. You need to download the file to your local directory so the package can be uploaded in the process.

### Steps:

Perform the following steps to upgrade the firmware:

1. Click **Service and Maintenance** from the portfolio.
2. Click **Upgrade Software** to check the installed firmware level or to install a new level of firmware on the cluster. The Software Upgrade panel is displayed.

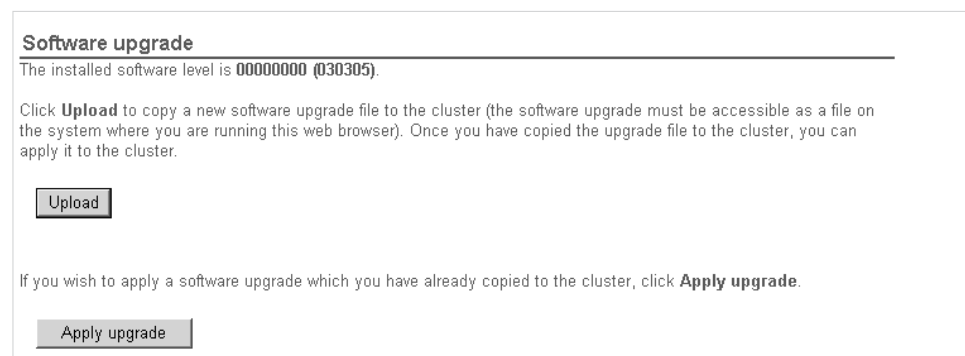



Figure 24. Software upgrade panel

Specify the location of the firmware upgrade file. This could be a URL.

3. Click **Apply upgrade** to upgrade the software. This page enables you to select the upgrade and to apply it to the cluster. This page displays a list of the software levels that you can apply to the cluster.



*Figure 25. Applying software upgrade panel*

**Result:**

When a new code level is applied, it is automatically installed on all the nodes that are in the cluster. Note this process can take up to 30 minutes per node.

**Post-processing requirements:**

To check the status of the upgrade, issue the following commands from the CISCO MDS 9000 switch:

```
svc -config
show nodes local
```

**Related topics:**

- Chapter 24, “Automatic upgrade,” on page 205

---

## Chapter 24. Automatic upgrade

New nodes introduced to the cluster normally have software packages downloaded to them from the cluster without any manual intervention. A new node requiring a code version higher than that currently available on the cluster or a node that already contains a code version higher than that on the cluster will not be configured into the cluster. If a node is added to the network that has no code installed, for example because the disk drive has been replaced, or it has such an old code version installed that it cannot advertise itself to the clusters, a reinstall of the software is forced by using the Node Rescue procedure.

When new nodes are added to the cluster, the upgrade packages are usually automatically downloaded to them from the SAN Volume Controller Storage Software for Cisco MDS 9000 cluster. No manual intervention is needed.

If you add a new SAN Volume Controller Storage Software for Cisco MDS 9000 node that has a code version that is higher than the one that is available on the cluster, that node is *not* configured into the cluster. It will join the cluster, however the node will be downgraded to the cluster level.

**Error counts:** During the SAN Volume Controller Storage Software for Cisco MDS 9000 software upgrade, you can expect to see either I/O error counts displayed by the data-path query adapter, or an increase in the number of **datapath query device** commands if active I/O operations exist between hosts and the SANs. See the *IBM TotalStorage Subsystem Device Driver: User's Guide* for more information about **datapath query** commands.

During the software upgrade, each SAN Volume Controller Storage Software for Cisco MDS 9000 node of a working pair is upgraded sequentially. The SAN Volume Controller Storage Software for Cisco MDS 9000 node that is being upgraded is temporarily unavailable, and all I/O operations to that SAN Volume Controller Storage Software for Cisco MDS 9000 fail. As a result, I/O error counts increase. However, failed I/O operations are directed to the other SAN Volume Controller Storage Software for Cisco MDS 9000 node of the working pair, and applications should not see any I/O failures.





---

## Chapter 25. Automatic recovery from upgrade problems

The cluster will automatically terminate the upgrade process if any of the nodes fail to upgrade to the new software level. In this case, any nodes that have already upgraded to the new software level will downgrade back to the original code level. You should check the error log to determine the reason for the failure before attempting to upgrade the cluster again.



---

## Chapter 26. Installing the upgrade using the CLI

If you want to use secure copy, perform the following:

1. Use switch commands to upgrade the software package. Refer to Cisco MDS 9000 publications for information about installing the upgraded software package.

**Note:** The installation process will *only* fail when some paths between the host systems and the cluster are not available. Data access can be lost temporarily during the upgrading process. You can prevent this if, before you start the installation, you issue a data-path query device on each host system to ensure that all paths are available. See the *IBM TotalStorage Subsystem Device Driver: User's Guide* for more information about data-path query commands.

**Attention:** The order in which the nodes are upgraded depends on the following:

- The position of the nodes. The code will be transferred to all the nodes in an I/O group.
  - The I/O group ID. The code will be transferred from the lowest I/O group ID that includes nodes on it.
2. To verify that the upgrade was successful, you can perform any one of the following steps:
    - The code level is distributed to all the nodes that are in the cluster. The nodes, in turn, are then restarted. If all the nodes successfully restart with the new code level, the new version is committed and the cluster vital product data (VPD) is updated to new level of code.
    - The software upgrade is complete when the cluster verifies the upgrade package. To determine whether the upgrade has completed, you must either display the software version in the cluster VPD, or look for the Software upgrade complete event in the SAN Volume Controller Storage Software for Cisco MDS 9000 error or event log. If the node does not restart automatically during the upgrade, you should repair or manually delete that node from the cluster to complete the backout process.
    - Alternatively, you can also either perform the following steps:
      - a. Issue the **svctask dumperrlog** command to dump the contents of the error log to a text file. You can also use this command to delete unwanted error log dumps from the cluster.
      - b. Once you have the contents of the error log dumped into a text file, verify that there were no errors in the text file. If there are no errors, you have successfully upgraded the software and output similar to the following is displayed in the log file:

```
Upgrade completed successfully
```
    - c. Issue the **svcinfo lsnodevpd** command for each node. You should see that the software version field has been updated.

### Related topics:

- Chapter 23, “Upgrading the SAN Volume Controller Storage Software for Cisco MDS 9000 firmware using the Console for SAN Volume Controller for Cisco MDS 9000,” on page 203

- “Related publications” on page xi

---

## Chapter 27. Installing the software

**Software package:** The software is delivered to you as a single package. Cluster software versions comprise a number of software components that are delivered as a single package. The size of the software update package depends on the number of components that are being replaced by that upgrade package. The software installation procedure involves copying the new software version to the cluster and then starting an automatic installation process. This installation process might take up to two hours to complete. If the installation fails for any reason, the process might take up to four hours, two hours to install it and two hours to uninstall it. During the process, the software version is checked to ensure compatibility with other software present in the switch and each of the nodes is restarted in turn. Once all the nodes in the cluster have been successfully restarted with the new software the new software version is automatically committed. While each node is being restarted there might be some degradation in the maximum input/output rate that can be sustained by the cluster.

### **Installation operation:**

The installation operation can normally be performed concurrently with normal user I/O operations. If any restrictions apply to the operations that can be performed during the upgrade, then these restrictions will be documented on the SAN Volume Controller Storage Software for Cisco MDS 9000 web site from where the upgrade package was obtained. During the upgrade operation, only the following SAN Volume Controller Storage Software for Cisco MDS 9000 commands will be operational from the time the install process starts to the time that the new software is committed or until the process has been backed-out. All other commands will fail with a message indicating that a software upgrade is in progress. In the following commands, xxxx is the object type.

- **svcinfo lsxxxx**
- **svcinfo lsxxxxcandidate**
- **svcinfo lsxxxxprogress**
- **svcinfo lsxxxxmember**
- **svcinfo lsxxxxextent**
- **svcinfo lsxxxxdumps**
- **svcinfo caterrlog**
- **svcinfo lserrlogbyxxxx**
- **svcinfo caterrlogbyseqnum**
- **svctask rmnode**
- **svcservicetask rmnode**

Because of the operational limitations that occur during the upgrade process the software installation is a customer task.



---

## Chapter 28. Manual recovery from software upgrade problems

**Attention:** This procedure causes a loss of *all* data currently configured in the cluster. This is a last resort and should only be done if you have recently backed-up your data.

When a revised version of software is committed, you might not be able to return to a previous software version because some data structures might have been changed such that they cannot be used with the previous software version. Therefore, if you have any problems, you must go forward to a later version of the code. In extreme conditions where you cannot wait for a software update and you need to return to the previous software version, you can use the following procedure.

**Attention:** This procedure, however, causes the total loss of the SAN Volume Controller Storage Software for Cisco MDS 9000 cluster. This should only be done as a last resort.

### Steps:

Perform the following steps to reset from software upgrade problems:

1. Power-off all but one of the nodes that are in the cluster.
2. Set the powered-on node to the service access mode.
3. Use the service access functions to force the download of the older software package.
4. Repeat the action for each of the failed nodes.
5. From a node that has the new code, create a new cluster.





---

## **Part 7. Configuring other SAN devices for use with the SAN Volume Controller Storage Software for Cisco MDS 9000**

### **Related topics:**

- Chapter 29, “Configuring and servicing storage subsystems,” on page 217



---

## Chapter 29. Configuring and servicing storage subsystems

### Overview:

Although virtualization provides many benefits over direct attached, or direct SAN attached storage, it is also more liable to and is not restricted by, the creation of performance hot-spots. These hot-spots could lead to I/O errors being returned to your hosts, and potentially a loss of access to data may occur.

Follow the guidelines and procedures outlined in this section to make the most of the performance available from your storage subsystems and to avoid potential I/O problems.

### Related topics:

- Chapter 30, “Configuring the EMC CLARiiON controller,” on page 221
- Chapter 31, “Configuring the Enterprise Storage Server,” on page 231
- Chapter 32, “Configuring the FASTT disk controller system,” on page 235
- Chapter 33, “Configuring the HDS Lightning disk controller system,” on page 243
- Chapter 34, “Configuring the HP StorageWorks subsystem,” on page 247

---

## Configuring a balanced storage subsystem

The attachment of a given storage subsystem to a SAN Volume Controller Storage Software for Cisco MDS 9000 requires that some specific settings be applied to the device, some limitations are also listed for each storage type. There are 2 major steps in this process:

1. Setting the characteristics of the SAN Volume Controller Storage Software for Cisco MDS 9000 to storage connection(s)
2. Mapping logical unit(s) to these connections such that the SAN Volume Controller Storage Software for Cisco MDS 9000 can access them.

The virtualization features of the IBM Total Storage SAN Volume Controller Storage Software for Cisco MDS 9000 enable you to choose how your storage is divided up and presented to hosts. While virtualization provides you with a great deal of flexibility, it also offers the potential for setting up a storage subsystem which can be overloaded. A storage subsystem is overloaded if the quantity of I/O transactions that are being issued by the host systems exceeds the capability of the storage to process those transactions. If a storage subsystem is overloaded, then, at best, it causes delays in the host systems and, at worst, it causes I/O transactions to be timed out in the host which leads to errors being logged by the hosts and I/Os being failed back to applications.

As an extreme example of an overloaded storage subsystem it would be possible to use an IBM Total Storage SAN Volume Controller Storage Software for Cisco MDS 9000 to virtualize a single RAID array and to divide this storage among sixty-four host systems. Clearly, if all host systems attempt to access this storage at the same time the single RAID array will be overloaded. The following guidelines are provided to help you configure balanced storage subsystems.

### Steps:

Perform the following steps to configure balanced storage subsystems:

1. Calculate the I/O rate for an array. For each RAID array in the storage subsystem use the following table to calculate the approximate number of I/O operations per second that can be processed by the RAID array. Note that the actual number of I/O operations per second that can be processed will vary depending on the location and length of each I/O, whether the I/O is a read or a write operation and on the specifications of the component disks of the RAID array.

Table 14. Calculate the I/O rate

| Type of RAID Array                                      | Number of component disks in the RAID Array | Approximate I/O rate |
|---------------------------------------------------------|---------------------------------------------|----------------------|
| RAID-1 (mirrored) arrays                                | 2                                           | 300                  |
| RAID-3, RAID-4, RAID-5 (striped + parity) arrays        | $N + 1$ parity                              | $150 * N$            |
| RAID-10, RAID 0+1, RAID 1+0 (striped + mirrored) arrays | $N$                                         | $150 * N$            |

For example, a RAID-5 array with eight component disks has an approximate I/O rate of  $150 * 7 = 1050$ .

2. Calculate the I/O rate for a managed disk. If there is a one-to-one relationship between backend arrays and managed disks (this is the recommended configuration) then the I/O rate for a managed disk is the same as the I/O rate of the corresponding array. If an array is divided up into multiple managed disks then the I/O rate per managed disk is the I/O rate of the array divided by the number of managed disks that are using that array.
3. Calculate the I/O rate for a managed disk group. The I/O rate for a managed disk group is simply the sum of the I/O rates of the managed disk within that group.

For example, a managed disk group contains eight managed disks each of which corresponds to a RAID-1 array. From the table above the I/O rate for each managed disks can be calculated as 300. The I/O rate for the managed disk group is therefore  $300 * 8 = 2400$ .

4. Calculate the impact of FlashCopy relationships. If you are using the FlashCopy feature provided by the IBM Total Storage SAN Volume Controller Storage Software for Cisco MDS 9000 then you need to consider how much additional I/O will be generated by using this feature as this will reduce the rate at which I/O from host systems that can be processed. When a FlashCopy relationship is copying data any write I/Os from host systems to areas of the source or target virtual disk that have not yet been copied will cause extra I/Os to be generated by the IBM Total Storage SAN Volume Controller Storage Software for Cisco MDS 9000 to copy the data before the write I/O is actually performed. The effect of using FlashCopy depends on the type of I/O workload being generated by an application:

Table 15. Calculate the impact of FlashCopy relationships

| Type of application                       | Impact to I/O rate         | Additional Weighting for FlashCopy |
|-------------------------------------------|----------------------------|------------------------------------|
| Application is doing no I/O               | Insignificant impact       | 0                                  |
| Application is only reading data          | Insignificant impact       | 0                                  |
| Application is only issuing random writes | Up to 50 times as much I/O | 49                                 |

Table 15. Calculate the impact of FlashCopy relationships (continued)

|                                                   |                            |    |
|---------------------------------------------------|----------------------------|----|
| Application is issuing random reads and writes    | Up to 15 times as much I/O | 14 |
| Application is issuing sequential reads or writes | Up to 2 times as much I/O  | 1  |

For each virtual disk that is either the source or target of an active FlashCopy relationship consider the type of application that will be using that virtual disk and record the additional weighting for that virtual disk.

For example, a FlashCopy relationship is being used to provide point in time backups. During the FlashCopy process, a host application generates an I/O workload of random reads and writes to the source virtual disk. A second host application reads the target virtual disk and writes the data to tape to create a backup. The additional weighting for the source virtual disk is 14. The additional weighting for the destination virtual disk is 0.

- Calculate the I/O rate for virtual disks in a managed disk group. Calculate the number of virtual disks in the managed disk group. Add the additional weighting for each virtual disk that is either the source or a target of an active FlashCopy relationship. Divide the I/O rate of the managed disk group by this number to give an I/O rate per VDisk.

Example 1: A managed disk group has an I/O rate of 2400 and contains 20 virtual disks. There are no FlashCopy relationships. The I/O rate per virtual disk is  $2400 / 20 = 120$ .

Example 2: A managed disk group has an I/O rate of 5000 and contains 20 virtual disks. There are two active FlashCopy relationships which have source virtual disks in this managed disk group. Both source virtual disks are being accessed by applications issuing random reads and write and hence the additional weighting for each of these virtual disks is 14. The I/O rate per virtual disk is  $5000 / (20 + 14 + 14) = 104$ .

- Determine whether the storage subsystem is overloaded. The figure determined at step 4 on page 218 provides some indication of how many I/O operations per second can be processed by each virtual disk in the managed disk group. If you know how many I/O operations per second your host applications generate you can compare these figures to determine if the system is overloaded. If you do not know how many I/O operations per second your host applications generate then you can either measure this (for example by using the I/O statistics facilities provided by the IBM Total Storage SAN Volume Controller Storage Software for Cisco MDS 9000 to measure the I/O rate of your virtual disks) or use the following table as a guideline:

Table 16. Determine if the storage subsystem is overloaded

| Type of Application                              | I/O rate per virtual disk |
|--------------------------------------------------|---------------------------|
| Applications that generate a high I/O workload   | 200                       |
| Applications that generate a medium I/O workload | 80                        |
| Applications that generate a low I/O workload    | 10                        |

- Interpret the result. If the I/O rate generated by the application exceeds the I/O rate you calculated per virtual disk this indicates that you could be overloading your storage subsystem and you should monitor the system carefully to see if the backend storage is actually limiting the overall performance of your system. It is also possible that the calculation above is too simplistic to model your use

of storage, for example the calculation assumes that your applications generate the same I/O workload to all virtual disks.

One method you can use to monitor the performance of your storage subsystem is to use the I/O statistics facilities provided by the IBM Total Storage SAN Volume Controller Storage Software for Cisco MDS 9000 to measure the I/O rate of your managed disks. Alternatively you could use the performance and I/O statistics facilities provided by your backend controllers.

If you find your storage subsystem is over loaded there are several actions that can be take to resolve the problem:

- a. Adding more backend storage to the system will allow you to increase the quantity of I/O that can be processed by your storage subsystem. The virtualization and data migration facilities provided by the IBM Total Storage SAN Volume Controller Storage Software for Cisco MDS 9000 can be used to redistribute the I/O work load of virtual disks across a greater number of managed disks without having to take the storage offline.
- b. Stop any unessential FlashCopy relationships as this will reduce the amount of I/O operations submitted to the backend storage. If you are making many FlashCopy's in parallel then consider starting less FlashCopy relationships in parallel.
- c. The I/O workload generated by a host can often be limited by adjusting the queue depth (for example, the maximum number of I/O operations that are submitted in parallel). Depending on the type of host and type of host bus adapters it may be possible to limit the queue depth per virtual disk and/or limit the queue depth per host bus adapter. An alternative method of limiting the I/O workload generated by a host would be to use the I/O governing features provided by the IBM Total Storage SAN Volume Controller Storage Software for Cisco MDS 9000. These techniques may be particularly applicable if using a mixture of different host systems to prevent one host system from saturating an I/O subsystem to the detriment of the other host systems. Note that although these techniques may be used to avoid I/O time-outs it still means the performance of your system is being limited by the amount of storage.

---

## Chapter 30. Configuring the EMC CLARiiON controller

### Related topics:

- Chapter 29, “Configuring and servicing storage subsystems,” on page 217

---

### Configuring the EMC CLARiiON controller with Access Logix installed

#### Prerequisites:

It is assumed at this point that the EMC CLARiiON controller is not connected to the SAN Volume Controller Storage Software for Cisco MDS 9000. It is also assume that you already have a RAID controller with logical units and you've identified which LUs will be presented to the SAN Volume Controller Storage Software for Cisco MDS 9000.

With Access Logix installed, the SAN Volume Controller Storage Software for Cisco MDS 9000 will not have access to the storage controllers LUs. In order to give SAN Volume Controller Storage Software for Cisco MDS 9000 access to a LU, the SAN Volume Controller Storage Software for Cisco MDS 9000 and LU must be associated using the EMC CLARiiON configuration tools. The association is formed by creating a Storage Group containing the LU and the SAN Volume Controller Storage Software for Cisco MDS 9000.

#### Related topics:

- “Registering the SAN Volume Controller Storage Software for Cisco MDS 9000 Ports with your EMC CLARiiON”
- “Configuring your storage groups” on page 222

### Registering the SAN Volume Controller Storage Software for Cisco MDS 9000 Ports with your EMC CLARiiON

#### Prerequisites:

At this point, the EMC CLARiiON controller is not connected to the SAN Volume Controller Storage Software for Cisco MDS 9000. It is assume you already have a RAID controller with logical units (LUs) and you've identified which LUs will be presented to the SAN Volume Controller Storage Software for Cisco MDS 9000.

#### Steps:

Perform the following steps to register SVC ports with the EMC CLARiiON controller:

1. Connect the fibre channel and zone the MDS9000 as required.
2. Issue an `svctask detectmdisk` command .
3. From the Enterprise Storage window, right-click on the storage subsystem.
4. Select **Connectivity Status**. The Connectivity Status window is displayed.
5. Click **New**. The Create Initiator Record window is displayed.
6. Wait for the list of SVC ports to appear in the dialog box. (Identify them from the WWPN.) This can take several minutes.
7. Click **Group Edit**.

8. Select all instances of all the SVC ports in the Available dialog box. Then click the right arrow to move them to the selected box.
9. Fill in the HBA WWN field. You must know the following information :
  - WWNN of each SAN Volume Controller Storage Software for Cisco MDS 9000 in the cluster
  - WWPN of each port ID for each node on the cluster

The HBA WWN field is made up of the WWNN and the WWPN for the SAN Volume Controller Storage Software for Cisco MDS 9000 port. The following output shows an example:  
50:05:07:68:01:00:8B:D8:50:05:07:68:01:20:8B:D8
10. Select **A** in the field marked SP and **0** in the SP Port field.
11. For the Initiator Type field, select **CLARiiON Open** in the drop down list.
12. Deselect the ArrayCommPath checkbox if it has been selected.
13. Select **2** from the Failover Mode field drop down list.
- 14.

**Note:** Failure to select failover mode 2 will prevent SVC from being able to failover I/O. Your data may become unavailable in the event of a single failure.

Assign a host name in the Host Name field.

15. Click **OK**.
16. Specify the IP address of your MDS9000 switch. The EMC CLARiiON will not use this IP address. However it must be unique (within the EMC CLARiiON) to prevent errant behavior by Navisphere.
17. Perform step 10 for all possible combinations. The following example shows the different combinations of a subsystem with four ports:
  - SP: A SP Port: 0
  - SP: A SP Port: 1
  - SP: B SP Port: 0
  - SP: B SP Port: 1

**Result:**

All your WWPNs are registered against the host name that you specified.

## Configuring your storage groups

**Steps:**

Perform the following steps to configure your storage groups:

1. From the Enterprise Storage window, right-click on the storage subsystem.
2. Select **Create Storage Group**. The Create Storage Group window is displayed.
3. Choose a name for your storage group. Enter this name in the Storage Group Name field.
4. Select **Dedicated** in the Sharing State field.
5. Click **OK**. The storage group has been created.
6. From the Enterprise Storage window, right-click the storage group that was just created.
7. Select **Properties**. The Storage Group Properties window is displayed.



8. From the Storage Group Properties window, perform the following steps:
  - a. Select the **LUNs** tab.
  - b. Select the LUNs you want SAN Volume Controller Storage Software for Cisco MDS 9000 to manage in the Available LUNs table.
 

**Attention:** Ensure that the logical units that you have selected are not used by another storage group.
  - c. Click the forward arrow button.
  - d. Click **Apply**. A Confirmation window is displayed.
  - e. Click **Yes** to continue. A Success window is displayed.
  - f. Click **OK**.
  - g. Select the **Hosts** tab.
  - h. Select the host you created from step 14 on page 222.
 

**Attention:** Ensure only SAN Volume Controller Storage Software for Cisco MDS 9000 hosts (initiator ports) are in the storage group.
  - i. Click the forward arrow button.
  - j. Click **OK**. The Confirmation window is displayed.
  - k. Click **Yes** to continue. A Success window is displayed.
  - l. Click **OK**.

---

## Configuring the EMC CLARiiON controller (Access Logix not installed)

If Access Logix is not installed on an EMC CLARiiON controller, all LUs that were created on the controller may be used by the SAN Volume Controller Storage Software for Cisco MDS 9000. No further configuration of the EMC CLARiiON controller is necessary.

Configure the switch zoning such that no hosts can access these LUs.

**Note:** The EMC CLARiiON must not be shared with direct attached hosts.

**Related topics:**

- “Switch zoning limitations for the EMC CLARiiON” on page 225

---

## Supported models of the EMC CLARiiON

The following table lists the supported models of the EMC CLARiiON.

*Table 17. Supported models of the EMC CLARiiON*

| Model    |
|----------|
| FC4700–1 |
| FC4700–2 |
| CX200    |
| CX400    |
| CX600    |

---

## Supported firmware levels for the EMC CLARiiON

See the following Web site for specific firmware levels and the latest supported hardware: <http://www.ibm.com/storage/support/2145/>

---

## Concurrent maintenance on the EMC CLARiiON

Concurrent maintenance of the EMC CLARiiON is supported with the SAN Volume Controller Storage Software for Cisco MDS 9000.

Concurrent maintenance is the ability to perform I/O operations to a controller while simultaneously performing maintenance on it. This topic lists the components of the EMC CLARiiON that can be maintained in this way.

The EMC CLARiiON model FC series allows for concurrent replacement of the following components:

- Disk drives
- Controller fans (Fans must be replaced within 2 minutes or controllers will shut down.)
- Disk enclosure fans (Fans must be replaced within 2 minutes or controllers will shut down.)
- Controller (Service Processor: cache must be disabled first)
- Fibre Channel Bypass cards (LCC)
- Power supplies (Fans must be removed first.)

The EMC CLARiiON model CX series allows for concurrent replacement of the following components:

- Disk drives
- Controller (Service processor or drawer controller)
- Power/cooling modules (Modules must be replaced within 2 minutes or controllers will shut down.)

### Notes:

1. EMC CLARiiON procedures for concurrent upgrade must be followed in all cases.
2. The CX Series also has a feature called Data In Place Upgrade which allows you to upgrade from one model to another (for example, from the CX200 to the CX600) with no data loss or migration required. This is **not** supported by the SAN Volume Controller Storage Software for Cisco MDS 9000.

---

## Sharing the EMC CLARiiON between a host and the SAN Volume Controller Storage Software for Cisco MDS 9000

The EMC CLARiiON can be shared between a host and a SAN Volume Controller Storage Software for Cisco MDS 9000. This topic briefly discusses that restriction.

- Split controller access is only supported when Access Logix is installed and enabled.
- A host cannot be connected to both the SAN Volume Controller Storage Software for Cisco MDS 9000 and EMC CLARiiON at the same time.
- LUs must not be shared between a host and a SAN Volume Controller Storage Software for Cisco MDS 9000.
- Partitions in a RAID group must not be shared between a host and a SAN Volume Controller Storage Software for Cisco MDS 9000.

---

## Switch zoning limitations for the EMC CLARiiON

The number of connections (process logins) consumed by the SAN Volume Controller Storage Software for Cisco MDS 9000 cluster and the EMC CLARiiON must be carefully considered. For a single fabric, the number of connections is:

- Number of SAN Volume Controller Storage Software for Cisco MDS 9000 ports \* number of EMC CLARiiON ports

The EMC CLARiiON CX200 provides 2 ports and supports 30 connections. Using a single SAN fabric, a four-node cluster requires 32 connections ( $4 * 4 * 2$ ). This exceeds the CX200 capability and exposes the SAN Volume Controller Storage Software for Cisco MDS 9000 cluster integrity.

The solution is to either zone the fabric, or have two fabrics such that only two ports on each SAN Volume Controller Storage Software for Cisco MDS 9000 node are visible to each CX200 port. Each zone/fabric therefore consumes 8 connections ( $4 * 2 * 1$ ) and the CX200 only deals with 16 connections.

A maximum configuration, using the EMC CLARiiON CX200, is a 2 node SAN Volume Controller Storage Software for Cisco MDS 9000 cluster with 7 dual-ported hosts.

EMC CLARiiON FC4700 and CX400 systems provide 4 ports and support 64 connections. Using a single SAN fabric, a four-node cluster requires 64 connections ( $4 * 4 * 4$ ). This equals the EMC CLARiiON capabilities and is therefore only a problem if split support with other hosts is required.

---

## Quorum disks on the EMC CLARiiON

Quorum disks are not supported for the EMC CLARiiON.

---

## Advanced functions for the EMC CLARiiON

### Flash Copy & SnapView:

The EMC CLARiiON's form of FlashCopy called SnapView is not supported by the SAN Volume Controller Storage Software for Cisco MDS 9000. In a split controller configuration, SnapView is not supported even for the LUs that are controlled by the host.

### Remote Copy & MirrorView:

The EMC CLARiiON's form of Remote Copy called MirrorView is not supported by the SAN Volume Controller Storage Software for Cisco MDS 9000. In a split controller configuration, MirrorView is not supported even for the LUs that are controlled by the host.

### SAN Copy:

The EMC CLARiiON provides a form of Flash Copy called SAN Copy which is not supported by the SAN Volume Controller Storage Software for Cisco MDS 9000. In a split controller configuration, SAN Copy is not supported even for the LUs that are controlled by the host.

### MetaLUN:

MetaLUN allows an LU to be expanded using LUs in other RAID groups. The SAN Volume Controller Storage Software for Cisco MDS 9000 supports MetaLUN for migration purposes in Image Mode only.

---

## Logical unit creation and deletion on the EMC CLARiiON

Binding an LU to a RAID group can take a significant amount of time. The LU must not be added to a storage group until binding is complete. As a safeguard, the SAN Volume Controller Storage Software for Cisco MDS 9000 will not discover the LU if binding is in progress. A subsequent manual discovery is required.

### Related topics:

- “Discovering MDisks using the CLI” on page 158

---

## Configuring settings for the EMC CLARiiON

There are a large number of settings and options which are available through the EMC CLARiiON configuration interface. This topic and its subtopics explains those options and settings that are supported by the SAN Volume Controller Storage Software for Cisco MDS 9000.

These options and settings cover the following:

- Subsystem
- Port
- Logical unit

### Related topics:

- “Global settings for the EMC CLARiiON”
- “Port settings for the EMC CLARiiON” on page 227
- “LU settings for the EMC CLARiiON” on page 228

## Global settings for the EMC CLARiiON

Global settings apply across an EMC CLARiiON subsystem.

*Table 18. EMC CLARiiON global settings supported by the SAN Volume Controller Storage Software for Cisco MDS 9000*

| Option                                   | EMC CLARiiON Default Setting                    | SAN Volume Controller Storage Software for Cisco MDS 9000 Required Setting |
|------------------------------------------|-------------------------------------------------|----------------------------------------------------------------------------|
| Access Controls (Access Logix installed) | Not installed                                   | Either Installed or Not Installed                                          |
| Subsystem Package Type                   | 3                                               | 3                                                                          |
| Queue Full Status                        | Disable                                         | Disable                                                                    |
| Recovered Errors                         | Disable                                         | Disable                                                                    |
| Target Negotiate                         | Displays the state of the target negotiate bit. | Displays the state of the target negotiate bit.                            |
| Mode Page 8 Info                         | Disable                                         | Disable                                                                    |
| Base UUID                                | 0                                               | 0                                                                          |
| Write Cache Enabled                      | Enabled                                         | Enabled                                                                    |
| Mirrored Write Cache                     | Enabled                                         | Enabled                                                                    |

Table 18. EMC CLARiiON global settings supported by the SAN Volume Controller Storage Software for Cisco MDS 9000 (continued)

| Option                    | EMC CLARiiON Default Setting | SAN Volume Controller Storage Software for Cisco MDS 9000 Required Setting |
|---------------------------|------------------------------|----------------------------------------------------------------------------|
| Write Cache Size          | 600 MB                       | Default recommended                                                        |
| Enable Watermarks         | Enabled                      | Enabled                                                                    |
| Cache High Watermark      | 96%                          | Default                                                                    |
| Cache Low Watermark       | 80%                          | Default                                                                    |
| Cache Page Size           | 4 Kb                         | 4 Kb                                                                       |
| RAID3 Write Buffer Enable | Enable                       | Default recommended                                                        |
| RAID3 Write Buffer        | 0 MB                         | Default recommended                                                        |

**Related topics:**

- “Configuring a balanced storage subsystem” on page 217

## Controller settings for the EMC CLARiiON

Controller settings are settings that apply across one EMC CLARiiON subsystem.

Table 19 describes the options that can be set by the EMC CLARiiON.

Table 19. EMC CLARiiON controller settings supported by the SAN Volume Controller Storage Software for Cisco MDS 9000

| Option             | EMC CLARiiON Default Setting | SAN Volume Controller Required Setting |
|--------------------|------------------------------|----------------------------------------|
| Read Cache Enabled | Enable                       | Enable                                 |
| Read Cache Size    | 200 MB                       | Enable                                 |
| Statistics Logging | Disable                      | Either Enable or Disable               |

**Note:** The SAN Volume Controller Storage Software for Cisco MDS 9000 cannot obtain or change the configuration options listed above. It is therefore your responsibility to configure the options as recommended.

**Related topics:**

- “Configuring a balanced storage subsystem” on page 217

## Port settings for the EMC CLARiiON

Port settings are configurable at the port level.

Table 20. EMC CLARiiON port settings supported by the SAN Volume Controller Storage Software for Cisco MDS 9000

| Option     | EMC CLARiiON Default Setting | SAN Volume Controller Required Setting |
|------------|------------------------------|----------------------------------------|
| Port speed | 2 GB                         | Either 1 or 2 GB                       |

**Note:** The SAN Volume Controller Storage Software for Cisco MDS 9000 cannot obtain or change the configuration options listed above. It is therefore your responsibility to configure the options as recommended.

**Related topics:**

- “Configuring a balanced storage subsystem” on page 217

## LU settings for the EMC CLARiiON

LU settings are configurable at the LU level.

Table 21 describes the options that must be set for each logical unit that is accessed by the SAN Volume Controller Storage Software for Cisco MDS 9000. LUs that are accessed by hosts can be configured differently.

*Table 21. EMC CLARiiON LU settings supported by the SAN Volume Controller Storage Software for Cisco MDS 9000*

| Option                | EMC CLARiiON Default Setting | SAN Volume Controller Storage Software for Cisco MDS 9000 Required Setting |
|-----------------------|------------------------------|----------------------------------------------------------------------------|
| LU ID                 | Auto                         | N/A                                                                        |
| RAID Type             | 5                            | Any RAID Group                                                             |
| RAID Group            | Any available RAID Group     | Any available RAID Group                                                   |
| Offset                | 0                            | Any setting                                                                |
| LU Size               | ALL LBAs in RAID Group       | Any setting                                                                |
| Placement             | Best Fit                     | Either Best Fit or First Fit                                               |
| UID                   | N/A                          | N/A                                                                        |
| Default Owner         | Auto                         | N/A                                                                        |
| Auto Assignment       | Disabled                     | Disabled                                                                   |
| Verify Priority       | ASAP                         | N/A                                                                        |
| Rebuild Priority      | ASAP                         | N/A                                                                        |
| Strip Element Size    | 128                          | N/A                                                                        |
| Read Cache Enabled    | Enabled                      | Enabled                                                                    |
| Write Cache Enabled   | Enabled                      | Enabled                                                                    |
| Idle Threshold        | 0–254                        | 0–254                                                                      |
| Max Prefetch Blocks   | 0–2048                       | 0–2048                                                                     |
| Maximum Prefetch IO   | 0–100                        | 0–100                                                                      |
| Minimum Prefetch Size | 0–65534                      | 0–65534                                                                    |
| Prefetch Type         | 0, 1, or 2                   | 0, 1, or 2                                                                 |
| Prefetch Multiplier   | 0 to 2048 or 0 to 324        | 0 to 2048 or 0 to 324                                                      |
| Retain prefetch       | Enabled or Disabled          | Enabled or Disabled                                                        |
| Prefetch Segment Size | 0 to 2048 or 0 to 32         | 0 to 2048 or 0 to 32                                                       |
| Idle Delay Time       | 0 to 254                     | 0 to 254                                                                   |
| Verify Priority       | ASAP, High, Medium, or Low   | Low                                                                        |
| Write Aside           | 16 to 65534                  | 16 to 65534                                                                |

| **Note:** The SAN Volume Controller Storage Software for Cisco MDS 9000 cannot  
| obtain or change the configuration options listed above. It is therefore your  
| responsibility to configure the options as recommended.

| **Related topics:**

- | • “Configuring a balanced storage subsystem” on page 217





---

## Chapter 31. Configuring the Enterprise Storage Server

---

### Configuring the Enterprise Storage Server (ESS)

#### Steps:

Perform the following steps to configure the ESS:

1. Use a Web browser to access the ESS Specialist by entering the IP address of the ESS.
2. Login using the user name and password.
3. Click **ESS Specialist**.
4. Click **Storage Allocation**.
5. Click **Open System Storage**.
6. Click **Modify Host Systems**.
7. Create a host entry for every initiator port on every SAN Volume Controller Storage Software for Cisco MDS 9000 node in your cluster. Complete the following fields:

#### **Nickname**

Type a unique name for each port (for example, knode or 1node).

#### **Host Type**

Select **IBM SAN Volume Controller** or **RS/6000** if that is not available.

#### **Host Attachment**

Select **Fibre Channel attached**.

#### **Hostname/IP address**

Leave this field blank.

#### **WWPN**

Either select the WWPN from the list, or type it manually. A configuration command will fail if you use WWPN 0 in the command string.

8. After you are finished adding all of the ports, click **Perform Configuration Update**.
9. Click **Add Volumes** to add the volumes on which you want the SAN Volume Controller Storage Software for Cisco MDS 9000 to run.
10. From the Add Volumes window, perform the following actions:
  - a. Select any of the SAN Volume Controller Storage Software for Cisco MDS 9000 host ports that you created earlier.
  - b. Select the necessary ESS adapter to create the volumes.
  - c. Click **Next**.
  - d. Create volumes using your desired size, placement, and RAID level.
  - e. After you are done creating all the volumes, click **Perform Configuration Update**.
11. Map the volumes to all of your SAN Volume Controller Storage Software for Cisco MDS 9000 ports by performing the following steps:
  - a. Click **Modify Volume Assignments**.
  - b. Select all of the volumes that you created earlier.
  - c. Click **Assigning selected volumes to target hosts**.

- d. Select all of the remaining SAN Volume Controller Storage Software for Cisco MDS 9000 host ports that you created earlier.
- e. Select the **Use same ID/LUN in source and target** check box.
- f. Click **Perform Configuration Update**.

**Related topics:**

- Chapter 29, “Configuring and servicing storage subsystems,” on page 217
- “Configuring a balanced storage subsystem” on page 217
- “Configuring FAStT disk controllers for the storage server” on page 235
- “Support actions for the FAStT controller” on page 236

---

## Supported models of the ESS

The following table lists the supported models of the ESS.

*Table 22. Supported models of the Enterprise Storage Server*

| Model    |
|----------|
| 2105-F20 |
| 2105-800 |

---

## Supported firmware levels for the ESS

See the following Web site for specific firmware levels and the latest supported hardware: <http://www.ibm.com/storage/support/2145/>

---

## Concurrent maintenance on the ESS

Concurrent maintenance is the capability to perform I/O operations to an ESS while simultaneously performing maintenance operations on it. IBM supports all ESS concurrent maintenance procedures.

---

## Sharing the ESS between a host and the SAN Volume Controller Storage Software for Cisco MDS 9000

The ESS can be shared between a host and a SAN Volume Controller Storage Software for Cisco MDS 9000. This topic briefly discusses the restrictions.

IBM supports sharing an ESS between a SAN Volume Controller Storage Software for Cisco MDS 9000 and other hosts. However, if an ESS port is in the same zone as a SAN Volume Controller Storage Software for Cisco MDS 9000 port, that same ESS port should not be in the same zone as another host.

A single host can have both ESS direct-attached and SAN Volume Controller Storage Software for Cisco MDS 9000 virtualized disks configured to it. If a LUN is managed by the SAN Volume Controller Storage Software for Cisco MDS 9000, it should never be mapped to another host.

See the following Web site for the latest supported configurations: <http://www.ibm.com/storage/support/2145/>

---

## Switch zoning limitations for the ESS

The minimum number of cables recommended for redundancy is 2 cables from 2 separate adapter bays. Up to 16 cables can be used to connect to an ESS. Only 1 or 2 Gb fibre-channel attached is supported.

**Note:** ESCON, FICON, and Ultra SCSI attachment is not supported with the SAN Volume Controller Storage Software for Cisco MDS 9000.

---

## Quorum disks on the ESS

Managed disks presented by the ESS will be chosen by the SAN Volume Controller Storage Software for Cisco MDS 9000 as quorum disks.

**Related topics:**

- Chapter 29, “Configuring and servicing storage subsystems,” on page 217

---

## Advanced functions for the ESS

Only open system storage is supported by the SAN Volume Controller Storage Software for Cisco MDS 9000.

**FlashCopy and Concurrent copy:**

FlashCopy and Concurrent copy are not supported on any LUN that is managed by the SAN Volume Controller Storage Software for Cisco MDS 9000.

**Remote Copy or extended distance Remote Copy:**

Remote Copy or extended distance Remote Copy are not supported on any LUN that is managed by the SAN Volume Controller Storage Software for Cisco MDS 9000.

---

## Logical unit creation and deletion on the ESS

Certain ESS types are supported for use with the SAN Volume Controller Storage Software for Cisco MDS 9000.

Before deleting or un-mapping a LUN from the SAN Volume Controller Storage Software for Cisco MDS 9000, it must first be removed from the MDisk group it was apart of. The following is supported:

- The supported LUN size is 1GB - 2TB.
- RAID 5 and RAID 10 LUNs are supported.
- LUNs can be added dynamically. When adding a new LUN, the “Use same ID/LUN in source and target” check-box **MUST** be checked.

**Note:** Failure to do this could cause loss in redundancy, or data integrity issues. The detect MDisks action in the Console for SAN Volume Controller for Cisco MDS 9000 or the **linfo detectmdisks** command must be run in order for the SAN Volume Controller Storage Software for Cisco MDS 9000 to detect the new disks.

**Related topics:**

- “Configuring a balanced storage subsystem” on page 217



---

## Chapter 32. Configuring the FAStT disk controller system

---

### Configuring FAStT disk controllers for the storage server

**Attention:** The SAN Volume Controller Storage Software for Cisco MDS 9000 does not concurrently support I/O operations with the download of ESM (Environmental Services Monitor) firmware. You must quiesce all I/O operations from the hosts that are using storage provided by the FAStT controllers you wish to update before installing new ESM firmware.

The FAStT storage server has many options and actions. The following list the supported actions and its impact on the SAN Volume Controller Storage Software for Cisco MDS 9000 and its configuration.

1. host type:

- a. You must set either the default host type of your FAStT or the host type of the chosen array to:

IBM TS SAN VCE

You can set the host type in 2 ways:

- 1) Click **Storage Subsystem** → **Change** → **Default Host Type**, or
- 2) For each host port you can specify the host type of that port or modify existing ports.

2. WWNN:

- a. Set the subsystem so that both controllers have the same WWNN. Scripts are available from the FAStT support Web site to change the set up of the FAStT if required.

[www.storage.ibm.com](http://www.storage.ibm.com)

3. auto volume transfer (AVT):

- a. Make sure the auto volume transfer is enabled. The host type selection should have enabled this function already.
- b. View the storage subsystem profile data to confirm that you have the AVT function enabled. This storage profile is presented as a text view in a separate window.
- c. Scripts are available from the FAStT Web site to enable AVT if required.

[www.storage.ibm.com](http://www.storage.ibm.com)

4. limitations:

- a. Only one FAStT storage array can be created that contains any of the ports of any of the nodes in a single SAN Volume Controller Storage Software for Cisco MDS 9000 cluster.
- b. You must not map more than one array to any of the ports on any of the nodes in the same SAN Volume Controller Storage Software for Cisco MDS 9000 cluster. Otherwise, unexpected behavior might result. For example, there will not be any warning messages, however, there will be errors logged in the SAN Volume Controller Storage Software for Cisco MDS 9000 error log and access to storage may be lost.

5. access LUN:

- a. The access LUN, also known as the Universal Transport Mechanism (UTM) LUN, might not be in a partition that contains the SAN Volume Controller Storage Software for Cisco MDS 9000 ports. It is not required by the SAN

Volume Controller Storage Software for Cisco MDS 9000. The UTM LUN is a special LUN that allows the SAN Volume Controller Storage Software for Cisco MDS 9000 to be configured through suitable software over the Fibre channel connection. However, the SAN Volume Controller Storage Software for Cisco MDS 9000 does not require the UTM LUN, therefore does not generate errors either way.

- b. The FAStT *must not* have the Access (UTM) LUN presented as Logical Unit Number 0 (zero).
6. logical unit:
    - a. The SAN Volume Controller Storage Software for Cisco MDS 9000 attempts to follow the FAStT specified preferred ownership. You can specify which controller (A or B) is used to do I/O operations to a given Logical Unit. If the SAN Volume Controller Storage Software for Cisco MDS 9000 can see the ports of the preferred controller and no error conditions exist, then it will access that Logical Unit through one of the ports on that controller.
    - b. Under error conditions, the ownership is ignored. Meaning, the SAN Volume Controller Storage Software for Cisco MDS 9000 has found a given path through the fabric to be errant, or there is no connection to a given port.
  7. Copy services (FlashCopy and Remote Copy):
    - a. FAStT copy services must *not* be used when the SAN Volume Controller Storage Software for Cisco MDS 9000 is attached to the FAStT. Partitioning might allow copy services to be used on other host platforms.

**Related topics:**

- Chapter 29, “Configuring and servicing storage subsystems,” on page 217
- “Configuring a balanced storage subsystem” on page 217
- “Support actions for the FAStT controller”
- “Configuring the Enterprise Storage Server (ESS)” on page 231

---

## Support actions for the FAStT controller

The FAStT storage manager has many options and actions. The following shows the supported actions and their impact on the SAN Volume Controller Storage Software for Cisco MDS 9000 and its configuration.

1. controller run diagnostics:
  - a. The diagnostics should be automatically recovered by the SAN Volume Controller Storage Software for Cisco MDS 9000 software.
  - b. Check your MDisks to make sure that they have not been set to degraded mode after this action.
2. controller disable data transfer:
  - a. This option is not supported when a SAN Volume Controller Storage Software for Cisco MDS 9000 is attached to the FAStT. Loss of availability and redundancy may occur if data transfer is disabled.
3. setting an array Offline:
  - a. Do not set an array Offline. If you use this setting, you might lose access to the MDisk group.
4. array increase capacity:
  - a. Increasing capacity is supported but the new capacity is not usable until the MDisk is removed from an MDisk group and then added again. You might have to migrate data to increase the capacity.
5. redistribute logical drives or change ownership of the preferred path:

- a. These actions are supported but might not take effect until a cluster rediscovery is initiated on the SAN Volume Controller Storage Software for Cisco MDS 9000 cluster. This can be achieved using the **svctask detectmdisk** command.
6. controller reset
  - a. Controller reset should only be performed if directed to do so by service personnel, the alternate controller is functional and available to the SAN. The SAN Volume Controller Storage Software for Cisco MDS 9000 reset should be automatically recovered by the SAN Volume Controller Storage Software for Cisco MDS 9000 software.
  - b. Check your MDisks to make sure that they have not been set to degraded state during this operation. You can issue the **svctask includemdisk** to repair degraded MDisks.

**Related topics:**

- Chapter 29, “Configuring and servicing storage subsystems,” on page 217
- “Configuring a balanced storage subsystem” on page 217
- “Configuring FAStT disk controllers for the storage server” on page 235
- “Configuring the Enterprise Storage Server (ESS)” on page 231

---

## Supported models of the IBM FAStT controller

The following table lists the supported models of the IBM FAStT controller.

*Table 23. Supported models of the IBM FAStT controller*

| Model                    |
|--------------------------|
| 3542 FAStT model 200     |
| 3552 FAStT model 500     |
| 1722 FAStT model 600     |
| 1742/1RU FAStT model 700 |
| 1742/90U FAStT model 900 |

---

## Supported firmware levels for the FAStT

See the following Web site for specific firmware levels and the latest supported hardware: <http://www.ibm.com/storage/support/2145/>

---

## Concurrent maintenance on the IBM FAStT

Concurrent maintenance is the capability to perform I/O operations to an IBM FAStT controller while simultaneously performing maintenance operations on it. Refer to your FAStT documentation for information about concurrent maintenance.

---

## Sharing the IBM FAStT controller between a host and the SAN Volume Controller Storage Software for Cisco MDS 9000

**Attention:** The use of the FAStT term partitioning, does not withhold the same meaning as used by IBM.

The FAStT function known as partitioning, must be used to separate groups of logical units that are directly attached to hosts or groups of hosts from the SAN Volume Controller Storage Software for Cisco MDS 9000 accessed logical units.

**Note:** The SAN Volume Controller Storage Software for Cisco MDS 9000 partition must either contain all the ports of the SAN Volume Controller Storage Software for Cisco MDS 9000 cluster that are connected to the SAN, or are zoned to have access to the FAStT ports. At least one port from each FAStT controller must be visible by the SAN Volume Controller Storage Software for Cisco MDS 9000 cluster.

---

## Quorum disks on the IBM FAStT

Managed disks presented by the IBM FAStT controller will be chosen by the SAN Volume Controller Storage Software for Cisco MDS 9000 as quorum disks.

**Note:** The FAStT 200 does not support quorum disks.

In order to avoid a single point of failure where the node pairs are connected to two different switches, there should be a connection from each switch to both of the controllers in the FAStT subsystem.

### Related topics:

- Chapter 29, “Configuring and servicing storage subsystems,” on page 217

---

## Advanced functions for the IBM FAStT

FlashCopy and Remote Copy facilities are provided as advanced functions on the IBM FAStT controller, however, these controller-delivered copy services are not supported by SAN Volume Controller Storage Software for Cisco MDS 9000.

## Data migration on an existing FAStT installation which contains arrays

You can enable the SAN Volume Controller Storage Software for Cisco MDS 9000 to be introduced to an existing SAN environment, so that you have the option of utilizing image mode LUNs to import the existing data into the virtualization environment without requiring a backup and restore cycle. For example, each FAStT array may contain up to 32 LUNs. Each array can only be accessed by a unique set of HBA ports (as defined by WWPNs). That is, for a single host to access multiple arrays, unique host fibre ports (WWPNs) need to be assigned to each array. All LUNs within an array are surfaced to assigned host fibre ports (no sub-array LUN mapping).

Host A is mapped to LUN 0, 1, 2 in Array 0  
Host B is mapped to LUN 0, 1, 2, 3, 4, 5 in Array 1  
Host C is mapped to LUN 0, 1, 2 in Array 2

To allow Host A to access the LUNs in array B, it is necessary to remove one of the HBAs (for example, A1) from the access list for array 0 and add it to array 1 (A1 cannot be on the access list for more than one array).

To add a SAN Volume Controller Storage Software for Cisco MDS 9000 into this configuration without save and restore cycles would require that a set of unique SAN Volume Controller Storage Software for Cisco MDS 9000 HBA port WWPNs for each array. This would allow the FAStT to surface the LUNs (with your data) to the SAN Volume Controller Storage Software for Cisco MDS 9000, which would



then configure these LUNs as image-mode LUNs and surface them to the required hosts. Unfortunately, this violates a requirement that all SAN Volume Controller Storage Software for Cisco MDS 9000 nodes be able to see all backend storage. To work around this problem, change the FAStT to allow more than 32 LUNs in 1 storage array, so that you can move all the LUNs from all the other arrays into 1 array and map to the SAN Volume Controller Storage Software for Cisco MDS 9000 cluster.

For example, lets say the FAStT has 8 arrays with 30 LUNs in each, and all need to be migrated to a 4-node SAN Volume Controller Storage Software for Cisco MDS 9000 cluster with 4 ports on each SAN Volume Controller Storage Software for Cisco MDS 9000. Perform the following:

1. Change the mappings for the first 4 arrays on the FAStT such that each array is mapped to 1 port on each node, this maintains redundancy across the cluster.
2. Create a new array on the FAStT that is mapped to all 4 ports on all the SAN Volume Controller Storage Software for Cisco MDS 9000s (actually not a array at all)
3. Gradually migrate the data into the MDisks in the target array, as storage is freed from the source arrays this can be reused as new storage in the target array. As arrays are deleted new arrays that need to be migrated can be mapped and migrated in the same way. The host side data access and integrity would be maintained throughout this process.

---

## Logical unit creation and deletion on the IBM FAStT

Certain IBM FAStT controller types are supported for use with the SAN Volume Controller Storage Software for Cisco MDS 9000.

To create a logical disk, you must set either the default host type of your FAStT or the host type of the chosen array to:

IBM TS SAN VCE

You can set the host type in 2 ways:

1. Click **Storage Subsystem** → **Change** → **Default Host Type**, or
2. For each host port you can specify the host type of that port or modify existing ports.

### Related topics:

- “Configuring a balanced storage subsystem” on page 217
- “Configuring FAStT disk controllers for the storage server” on page 235

---

## Configuration interface for the IBM FAStT

The access LUN, also known as the Universal Transport Mechanism (UTM) LUN is the configuration interface for the IBM FAStT controller.

The access LUN, might not be in a partition that contains the SAN Volume Controller Storage Software for Cisco MDS 9000 ports. It is not required by the SAN Volume Controller Storage Software for Cisco MDS 9000. The UTM LUN is a special LUN that allows the SAN Volume Controller Storage Software for Cisco MDS 9000 to be configured through suitable software over the Fibre channel connection. However, the SAN Volume Controller Storage Software for Cisco MDS

9000 does not require the UTM LUN, therefore does not generate errors either way. The FAStT *must not* have the Access (UTM) LUN presented as Logical Unit Number 0 (zero).

It is possible to use in-band (over fibre-channel) and out-of-band (Ethernet) to allow the FAStT configuration software to communicate with more than one FAStT. If using in-band configuration, the "Access" logical unit will need to be configured in a partition that does not include any logical units being accessed by the SAN Volume Controller Storage Software for Cisco MDS 9000 cluster.

**Note:** In-band is not supported via access to the LUN while in the SAN Volume Controller Storage Software for Cisco MDS 9000 partition.

**Related topics:**

- "Support actions for the FAStT controller" on page 236
- "Configuring FAStT disk controllers for the storage server" on page 235

---

## Controller settings for the IBM FAStT

Controller settings are settings that apply across one FAStT controller. For restrictions on controller settings, see the following:

- You must set either the default host type of your FAStT or the host type of the chosen array to:

IBM TS SAN VCE

You can set the host type in 2 ways:

1. Click **Storage Subsystem** → **Change** → **Default Host Type**, or
  2. For each host port you can specify the host type of that port or modify existing ports.
- Set the subsystem so that both controllers have the same WWNN. Scripts are available from the FAStT support Web site to change the set up of the FAStT if required.

[www.storage.ibm.com](http://www.storage.ibm.com)

- Make sure the auto volume transfer is enabled. The host type selection should have enabled this function already. View the storage subsystem profile data to confirm that you have the AVT function enabled. This storage profile is presented as a text view in a separate window. Scripts are available from the FAStT Web site to enable AVT if required.

[www.storage.ibm.com](http://www.storage.ibm.com)

- Ensure that you have the following enabled on any logical units mapped to the SAN Volume Controller Storage Software for Cisco MDS 9000:
  - read caching
  - write caching
  - write cache mirroring

Caching without batteries must not be enabled.

**Related topics:**

- "Configuring a balanced storage subsystem" on page 217
- "Configuring FAStT disk controllers for the storage server" on page 235

## Configuring settings for the IBM FAStT

There are a number of settings and options which are available through the IBM FAStT controller configuration interface. This topic and its subtopics discusses those options and settings that are supported with the SAN Volume Controller.

These options and settings can have a scope of a:

- Subsystem
- Logical unit, see the following:
  - The SAN Volume Controller Storage Software for Cisco MDS 9000 attempts to follow the FAStT specified preferred ownership. You can specify which controller (A or B) is used to do I/O operations to a given Logical Unit. If the SAN Volume Controller Storage Software for Cisco MDS 9000 can see the ports of the preferred controller and no error conditions exist, then it will access that Logical Unit through one of the ports on that controller. Under error conditions, the ownership is ignored. Meaning, the SAN Volume Controller Storage Software for Cisco MDS 9000 has found a given path through the fabric to be errant, or there is no connection to a given port.
  - Ensure that you have the following enabled on any logical units mapped to the SAN Volume Controller Storage Software for Cisco MDS 9000:
    - read caching
    - write caching
    - write cache mirroring

Caching without batteries must not be enabled.

### Related topics:

- “Configuring a balanced storage subsystem” on page 217
- “Configuring FAStT disk controllers for the storage server” on page 235

## Global settings for the IBM FAStT

Global settings apply across an IBM FAStT controller. This topic lists any global settings.

*Table 24. IBM FAStT controller global settings supported by the SAN Volume Controller Storage Software for Cisco MDS 9000*

| Option           | FAStT Default Setting |
|------------------|-----------------------|
| Start flushing   | 80%                   |
| Stop flushing    | 80%                   |
| Cache block size | 4 Kb                  |

These settings can be adjusted depending on the performance requirements. Modification of these settings is not recommended unless directed to do so by support personnel.

If partitioning is not being used, meaning all the FAStT logical units are visible to the SAN Volume Controller Storage Software for Cisco MDS 9000, then the default host type for the FAStT subsystem can be set. See step 1 on page 235. If partitioning is being used to group SAN Volume Controller Storage Software for Cisco MDS 9000 ports and host ports apart then the host type for each partition, or group of SAN Volume Controller Storage Software for Cisco MDS 9000 ports, must be defined. When defining host ports, the host type must be set to:IBM TS SAN VCE

| **Related topics:**

- | • “Configuring a balanced storage subsystem” on page 217

| **LU settings for the IBM FAStT**

| LU settings are configurable at the LU level. LUs that are accessed by hosts can be  
| configured differently. This topic lists those settings, what the IBM FAStT controller  
| defaults are, and what the required settings are for the SAN Volume Controller  
| Storage Software for Cisco MDS 9000.

| Read ahead cache multiplier is typically set to 0 or 1. Modification of these settings  
| is not recommended unless directed to do so by support personnel.

| Ensure that you have the following enabled on any logical units mapped to the SAN  
| Volume Controller Storage Software for Cisco MDS 9000:

- | • read caching  
| • write caching  
| • write cache mirroring

| Caching without batteries must not be enabled.

| When creating a new logical unit set the host type for that logical unit to the host  
| type IBM TS SAN VCE.

| **Note:** IBM TS SAN VCE is set as the default if the default type was already  
| displayed.

| **Related topics:**

- | • “Configuring a balanced storage subsystem” on page 217  
| • “Global settings for the IBM FAStT” on page 241

| **Miscellaneous settings for the IBM FAStT**

| There are options of a miscellaneous nature that must be set properly in order for  
| the IBM FAStT controller to work with the SAN Volume Controller Storage Software  
| for Cisco MDS 9000. Refer to your FAStT documentation for information about  
| other settings.

| **Related topics:**

- | • “Configuring a balanced storage subsystem” on page 217

| **Mapping and virtualization settings for IBM FAStT**

| This topic discusses LUN mapping or masking and virtualization from the point of  
| view of the IBM FAStT controller and their use in a SAN Volume Controller Storage  
| Software for Cisco MDS 9000 environment.

---

## Chapter 33. Configuring the HDS Lightning disk controller system

---

### Supported models of the HDS Lightning

The following table lists the supported models of the HDS Lightning.

*Table 25. Supported MDS 99xxV models*

| Model           |
|-----------------|
| Lightning 9970V |
| Lightning 9980V |

---

### Supported firmware levels for HDS Lightning

See the following Web site for specific firmware levels and the latest supported hardware: <http://www.ibm.com/storage/support/2145/>

**Note:** Concurrent upgrade of the controller firmware is **not** supported with the SAN Volume Controller Storage Software for Cisco MDS 9000.

---

### Concurrent maintenance on the HDS Lightning 99xxV

Concurrent upgrade of the controller firmware is supported with the SAN Volume Controller Storage Software for Cisco MDS 9000.

---

### Sharing the HDS Lightning 99xxV between host and the SAN Volume Controller

#### Sharing ports:

The HDS Lightning 99xxV can be shared between a host and a SAN Volume Controller Storage Software for Cisco MDS 9000 according to certain restrictions. The following restriction applies:

- The same host cannot be connected to both a SAN Volume Controller Storage Software for Cisco MDS 9000 and a Lightning at the same time because HDLM and the Subsystem Device Driver do not coexist.
- A controller port cannot be shared between a host and a SAN Volume Controller Storage Software for Cisco MDS 9000. In other words, if a controller port is used by a SAN Volume Controller Storage Software for Cisco MDS 9000 it must not be present in a switch zone which allows a host to access the port.
- Logical units (LUs) cannot be shared between a host and a SAN Volume Controller Storage Software for Cisco MDS 9000.

#### Supported Topologies:

SAN Volume Controller Storage Software for Cisco MDS 9000 supports connection to the Lightning according to the following restrictions:

- The SAN Volume controller resolves up to four WWNNs per subsystem and allows up to 512 LUs per WWNN. Lightning assigns a WWNN per port, therefore the SAN Volume Controller could be a limitation to both capacity (2048 LUs) and

bandwidth (4 ports). You can use the following procedure for Lightning subsystems with 8 ports if more capacity or bandwidth is required:

1. Divide the set of ports into groups of between 2 and 4.
2. Assign a discreet set of logical units to each group.

SAN Volume Controller then interprets each group as a separate subsystem.

- If a logical unit is mapped to the SAN Volume Controller port as LUNx, it must appear as LUNx to all the SAN Volume Controller ports in the cluster and must also appear as LUNx through all of the controller ports that it is mapped to.
- Command LUNs must not be mapped to the SAN Volume Controller.
- LUN Expansion (LUSE) and Virtual LVI/LUN operations cannot be run on a disk managed by the SAN Volume Controller. LUNs created using LUSE and Virtual LVI/LUN can be mapped to the SAN Volume Controller after they have been created.
- Only disks with open emulation can be mapped to the SAN Volume Controller. S/390 disks cannot be used with the SAN Volume Controller. Only fibre channel connections can be used to connect the SAN Volume Controller to the Lightning.

---

## Quorum disks on HDS Lightning 99xxV

Lightning is not an approved host for quorum disks. Therefore, configurations with only Lightning is not possible.

---

## Advanced functions for HDS Lightning

### ShadowImage:

ShadowImage is functionally similar to FlashCopy. ShadowImage is not supported when the disk controller system is being used with the SAN Volume Controller Storage Software for Cisco MDS 9000. Even when a Lightning 99xxV is shared between a host and a SAN Volume Controller Storage Software for Cisco MDS 9000, ShadowImage is not supported on the ports that are zoned directly with the host.

### LU Expansion:

The Lightning 99xxV supports Logical Unit expansion (LUSE). LUSE is a non-concurrent operation. LUSE is accomplished by concatenating between 2 and 26 existing logical units together. Before LUSE can be performed on a logical unit, it must be removed from an mdisk group and unmapped from the SAN Volume Controller Storage Software for Cisco MDS 9000.

**Attention:** This procedures will destroy all data that exists on the logical unit, except on a Windows system.

### TrueCopy:

TrueCopy is functionally similar to Remote Copy. TrueCopy is not supported when the disk controller system is being used with the SAN Volume Controller Storage Software for Cisco MDS 9000. Even when a Lightning 99xxV is shared between a host and a SAN Volume Controller Storage Software for Cisco MDS 9000, TrueCopy is not supported on the ports that are zoned directly with the host.

### Virtual LVI:

| The Lightning 99xxV supports Virtual LVI/LUNs. You can use this method to modify  
| a LUN size that the Lightning uses by dividing it into several smaller virtual LUNs.  
| This is a non-concurrent procedure that requires you to first create existing LUNs  
| into free space, then you must define their own LUNs using that free space. Virtual  
| LVI/LUNs must not be managed or mapped to a SAN Volume Controller Storage  
| Software for Cisco MDS 9000.

| LUNs that are set up using either LUSE or Virtual LVI/LUNs appear as normal  
| LUNs after they have been created. Therefore, LUNs set up using LUSE or Virtual  
| LVI/LUNs can be used by the SAN Volume Controller after they have been created.

| **Write protect:**

| Logical units (LUs) cannot be explicitly set to be write-protected. However, some of  
| the advanced features, such as Remote Copy, can be used to write-protect a LU as  
| part of the function. Remote Copy must not be used for LUs in use by a SAN  
| Volume Controller.





## Chapter 34. Configuring the HP StorageWorks subsystem

### Managed Disk (MDisk) Groups and MDisks:

An MDisk group should contain either no HSG80 LUNs or LUNs that are only from a single HSG80 subsystem. **No other configuration is supported.**

### Related topics:

- Chapter 29, “Configuring and servicing storage subsystems,” on page 217

## HP StorageWorks definitions

The following terms are used in both IBM and HP documentation, however, contain two different definitions.

| IBM term                            | IBM definition                                                                                                                                         | HP term                             | HP definition                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>container</b>                    | A visual user-interface component that holds objects.                                                                                                  | <b>container</b>                    | (1) Any entity that is capable of storing data, whether it is a physical device or a group of physical devices. (2) A virtual, internal controller structure representing either a single disk or a group of disk drives linked as a storageset. Stripesets and mirrorsets are examples of storageset containers that the controller uses to create units.   |
| <b>device</b>                       | A piece of equipment that is used with the computer. A device does not generally interact directly with the system, but is controlled by a controller. | <b>device</b>                       | 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, once the devices have been made known to the controller. |
| <b>just a bunch of disks (JBOD)</b> | See <i>non-RAID</i> .                                                                                                                                  | <b>just a bunch of disks (JBOD)</b> | A group of single-device logical units not configured into any other container type.                                                                                                                                                                                                                                                                         |

|                  |                                                                                                                                                                                                     |                  |                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>mirrorset</b> | See <i>RAID 1</i> .                                                                                                                                                                                 | <b>mirrorset</b> | A RAID storage set of two or more physical disks that maintains a complete and independent copy of the entire virtual disk's data. This type of storage set has the advantage of being highly reliable and extremely tolerant of device failure. Raid level 1 storage sets are referred to as mirrorsets.                                                                                 |
| <b>non-RAID</b>  | Disks that are not in a redundant array of independent disks (RAID).                                                                                                                                | <b>non-RAID</b>  | See <i>just a bunch of disks</i> .                                                                                                                                                                                                                                                                                                                                                        |
| <b>RAID 0</b>    | 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.                            | <b>RAID 0</b>    | 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 stripesets. |
| <b>RAID 1</b>    | A form of storage array in which two or more identical copies of data are maintained on separate media. Also known as mirrorset.                                                                    | <b>RAID 1</b>    | See <i>mirrorset</i> .                                                                                                                                                                                                                                                                                                                                                                    |
| <b>RAID 5</b>    | 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 disks in the array. | <b>RAID 5</b>    | See <i>RAIDset</i> .                                                                                                                                                                                                                                                                                                                                                                      |

|                  |                                                |                  |                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------|------------------------------------------------|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>RAIDset</b>   | See <i>RAID 5</i> .                            | <b>RAIDset</b>   | A specially developed RAID storage set that stripes data and parity across three or more members in a disk array. A RAIDset combines the best characteristics of RAID level 3 and RAID level 5. A RAIDset is the best choice for most applications with small to medium I/O requests, unless the application is write intensive. A RAIDset is sometimes called parity RAID. RAID level 3/5 storage sets are referred to as RAIDsets. |
| <b>partition</b> | A logical division of storage on a fixed disk. | <b>partition</b> | A logical division of a container represented to the host as a logical unit.                                                                                                                                                                                                                                                                                                                                                         |
| <b>stripeset</b> | See <i>RAID 0</i> .                            | <b>stripeset</b> | See <i>RAID 0</i> .                                                                                                                                                                                                                                                                                                                                                                                                                  |

---

## Configuring the HP StorageWorks controller

### Prerequisites:

Assume that the HP StorageWorks subsystem is not in use.

### Steps:

Perform the following steps to configure the HSG80 controller:

1. Review the SAN Volume Controller Storage Software for Cisco MDS 9000 error log and ensure there are no errors.
2. Ensure that the HP StorageWorks Operator Control Panel (OCP) on each HSG80 controller is clear of errors. The Operator Control Panel consists of seven green LED's at the rear of each HSG80 controller.
3. Ensure that you can use an HP StorageWorks command-line interface (CLI) to configure the HSG80 controllers.
4. Issue the **SHOW THIS** command and **SHOW OTHER** command to verify the following:
  - That the controller software is at the supported level. See .
  - Ensure that the controllers are configured for MULTIBUS FAILOVER with each other.
  - Ensure that the controllers are running in SCSI-3 mode.
  - Ensure that MIRRORED\_CACHE is enabled.
  - Ensure that the Host Connection Table is *not* locked.
5. Issue the **SHOW DEVICES FULL** command to verify the following:
  - Ensure that none of the LUNs are TRANSPORTABLE.
  - Ensure that all LUNs are configured. For example, the LUNs report their serial numbers and TRANSFER\_RATE\_REQUESTED correctly.

6. Issue the **SHOW FAILEDSET** command to verify that there are no failing disks.

**Note:** To verify, there should be no orange lights on any disks in the subsystem.

7. Issue the **SHOW UNITS FULL** command to verify the following:
  - Ensure that all LUNs are set to RUN and NOWRITEPROTECT
  - Ensure that all LUNs are ONLINE to either THIS or OTHER controller.
  - Ensure that all LUNs that are to be made available to the SAN Volume Controller Storage Software for Cisco MDS 9000 have ALL access.
  - Ensure that all LUNs have Host Based Logging NOT specified.

If you have partitioned LUNs, refer to the HP StorageWorks controllers topic.

8. Issue the **SHOW CONNECTIONS FULL** command to verify that you have enough spare entries for all combinations of SAN Volume Controller Storage Software for Cisco MDS 9000 ports and HP StorageWorks ports.
9. Connect up to four known good fibre-channel cables between your fibre-channel switches and your HP StorageWorks subsystem.
10. Ensure that your fibre-channel switches are zoned such that the SAN Volume Controller Storage Software for Cisco MDS 9000 and the HP StorageWorks subsystem are in a zone. Refer to zoning a switch.
11. Issue the **SHOW THIS** command and **SHOW OTHER** command to verify that each connected port is running. Something similar to the following output should be displayed: `PORT_1_TOPOLOGY=FABRIC` (fabric up).
12. Issue the **SHOW CONNECTIONS FULL** command to verify that the new connections have appeared for each SAN Volume Controller Storage Software for Cisco MDS 9000 port and HP StorageWorks port combination.
13. Verify that No rejected hosts is displayed at the end of the **SHOW CONNECTIONS** output.
14. On the SAN Volume Controller Storage Software for Cisco MDS 9000, issue the **svctask detectmdisk** command to discover the controller.
15. Issue the **svcinfolcontroller** command to verify that the two HSG80 serial numbers appear under the `ctrl s/n`.
16. Issue the **svcinfolmdisk** command to verify that the additional MDisks that correspond to the UNITS shown in the HP StorageWorks subsystem.

### **Result:**

You can now use the SAN Volume Controller Storage Software for Cisco MDS 9000 commands to create an MDisk group. You can also create and map VDIs from these MDisk groups. Ensure the SAN Volume Controller Storage Software for Cisco MDS 9000 error log is free of errors. Once you ensure that your host has reloaded its fibre-channel driver then you should be able to perform I/O to the VDIs. See the *IBM TotalStorage SAN Volume Controller: Host Attachment Guide* for detailed information.

### **Related topics:**

- “Creating managed disk (MDisk) groups using the CLI” on page 159
- “Create virtual disks (VDIs)” on page 162

---

## Supported models of the HP StorageWorks controller

**Attention:** The SAN Volume Controller Storage Software for Cisco MDS 9000 only supports configurations in which the HSG80 cache is enabled in writeback mode; running with only a single controller in the HSG80 subsystem results in a single point of data loss.

**Note:** TRANSPORTABLE disks are not supported for any models.

Table 26. Supported models of the HP StorageWorks HSG80

| Model         | Description                                                                                                                       |
|---------------|-----------------------------------------------------------------------------------------------------------------------------------|
| MA8000        | 1 controller enclosure (one or two HSG80 controllers), 3 dual bus 14 bay drive enclosures, 22U modular storage cabinet            |
| EMA12000 D14  | 3 controller enclosures (each with one or two HSG80 controllers), 9 dual bus 14 bay drive enclosures, 42U modular storage cabinet |
| EMA12000 S14  | 1 controller enclosure (with one or two HSG80 controllers), 6 single bus 14 bay drive enclosures, 42U modular storage cabinet     |
| EMA12000 Blue | 1 controller enclosure (with one or two HSG80 controllers), 3 dual bus 14 bay drive enclosures, 41U modular storage cabinet       |
| EMA16000 S14  | 2 controller enclosures with dual HSG80 controllers, 12 single bus 14 bay drive enclosures, wide 41U storage cabinet              |
| EMA16000 D14  | 4 controller enclosures with dual HSG80 controllers, 12 dual bus 14 bay drive enclosures, wide 41U storage cabinet                |

---

## Supported firmware levels for the HP StorageWorks controller

See the following Web site for specific firmware levels and the latest supported hardware: <http://www.ibm.com/storage/support/2145/>

**Note:** Concurrent upgrade of the controller firmware is also not supported with the SAN Volume Controller Storage Software for Cisco MDS 9000.

---

## Concurrent maintenance on the HP StorageWorks

Concurrent maintenance is the capability to perform I/O operations to a HP StorageWorks controller while simultaneously performing maintenance operations on it. This topic lists the components of the HP StorageWorks that can be maintained in this way.

**Note:** HP StorageWorks maintenance documentation uses the phrase 'rolling upgrade' in place of 'concurrent maintenance'. Please refer to this documentation as, in some instances, the level of I/O must be reduced before performing the maintenance procedure.

The HP StorageWorks Controller allows for concurrent replacement of the following components:

- Drive
- EMU
- Blower
- Dual power supply (one unit can be removed and replaced. The fans' speed increases when only one power supply unit is present.)

The following components are hot-pluggable, but maintenance concurrent with the SAN Volume Controller Storage Software for Cisco MDS 9000 I/O is not supported.

- Controller

The HP StorageWorks Controller does **not** allow for concurrent replacement of the following components:

- Single power supply (in a single power supply configuration, the enclosure is disabled when the power supply fails.)
- SCSI bus cables
- I/O module
- Cache

---

## Sharing the HP StorageWorks controller between a host and the SAN Volume Controller Storage Software for Cisco MDS 9000

---

### Switch zoning limitations for the HP StorageWorks subsystem

**Attention:** HSG80 based subsystems are supported with a single controller or dual controllers in the subsystem. Since the SAN Volume Controller Storage Software for Cisco MDS 9000 only supports configurations in which the HSG80 cache is enabled in write-back mode, running with only a single controller in the HSG80 subsystem results in a single point of data loss.

#### Switch zoning:

All switches should be zoned so that ports on the HSG80 are in the switch zone which contains the SAN Volume Controller Storage Software for Cisco MDS 9000 nodes.

#### Connecting to the SAN:

If you're using partitioned containers then there are restrictions on the way the HSG80 should be connected to the SAN Volume Controller Storage Software for Cisco MDS 9000 nodes. See "Support for HP StorageWorks advanced functions" on page 253.

#### Related topics:

- "Support for HP StorageWorks advanced functions" on page 253

---

## Quorum disks on the HP StorageWorks subsystem

Quorum disks are not supported for the HP StorageWorks subsystem.

---

## Support for HP StorageWorks advanced functions

FlashCopy and Remote Copy facilities are provided as advanced functions on the HSG80 controller, however, these controller-delivered copy services are not supported by SAN Volume Controller Storage Software for Cisco MDS 9000.

### Partitioning:

HSG80 supports Partitioning. A partition is a logical division of a container, represented to the host as a logical unit (LU). A container can be a RAID array or a JBOD. Any container type is a candidate for partitions. Any non-transportable disk or storage set can be divided into a maximum of 8 partitions.

Use of this feature is subject to the following restrictions:

- Partitioned containers are fully supported if the HSG80 subsystem is connected to the SAN by a single port.
- Partitioned containers will not be configured by the SAN Volume Controller Storage Software for Cisco MDS 9000 if the HSG80 subsystem is connected to the SAN by multiple ports.
- Partitioned containers will be removed from the configuration if a single port connection becomes a multi-port connection.
- Partitioned containers will be configured if a multi-port connection becomes a single port connection.

It is recommended that containers are partitioned such that no spare capacity exists because there is no way to detect this 'unused' partition. With a multi-port connection, subsequent attempts to use this capacity will remove all partitions on the container from the configuration.

### Dynamic array expansion (LU expansion):

HSG80 does not provide this feature.

### Write protection of LUNs:

This feature is not supported for use with the SAN Volume Controller Storage Software for Cisco MDS 9000.

---

## HP StorageWorks advanced functions

VDisks which are created from MDisks presented by an HSG80 controller may be used in SAN Volume Controller Storage Software for Cisco MDS 9000 FlashCopy mappings or SAN Volume Controller Storage Software for Cisco MDS 9000 Remote Copy relationships. That is, SAN Volume Controller Storage Software for Cisco MDS 9000 copy services fully supports the use of MDisks presented by an HSG80 controller.

---

## Logical unit creation and deletion on the HP StorageWorks

Table 27 lists the valid container types.

Table 27. HSG80 container types for logical unit configuration

| Container                            | Number of Members | Maximum Size |
|--------------------------------------|-------------------|--------------|
| JBOD - transportable (not supported) | 1                 | disk size    |

Table 27. HSG80 container types for logical unit configuration (continued)

| Container                                                                                                                                                                                                                          | Number of Members | Maximum Size             |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------|--------------------------|
| JBOD - non-transportable<br><br><b>Attention:</b> Provides no redundancy at the physical disk drive level, that is, a single disk failure may result in the loss of an entire managed disk group and its associated virtual disks. | 1                 | disk size minus metadata |
| Mirrorset                                                                                                                                                                                                                          | 2 to 6            | smallest member          |
| RAIDset                                                                                                                                                                                                                            | 3 to 14           | 1.024 terabytes          |
| Stripeset                                                                                                                                                                                                                          | 2 to 24           | 1.024 terabytes          |
| Striped Mirrorset                                                                                                                                                                                                                  | 2 to 48           | 1.024 terabytes          |

**Note:** Logical units can be created and deleted on an HSG80 subsystem while I/O operations are performed to other LUs. You do not need to reboot the HSG80 subsystem.

**Related topics:**

- “Configuring a balanced storage subsystem” on page 217

---

## Configuration interface for the HP StorageWorks

The Command Console configuration and service utility is the configuration interface for the HSG80 subsystem. This topic discusses the ways that this interface can connect to the subsystem and the requirements for doing so.

The configuration and service utility can connect to the subsystem in the following ways:

- RS232
- In band over fibre channel
- Over TCP/IP to a proxy agent which then communicates with the HSG80 in band over fibre channel.

**In band:**

**Attention:** POTENTIAL DATA CORRUPTION

In order for the Command Console to communicate with the HSG80 controllers, the host that runs the service utility must be able to access the HSG80 ports over the SAN. This host could therefore also access LUs that are visible to SVC and cause data corruption. To avoid this, set the UNIT\_OFFSET option to 199 for all connections to this host. This ensures that the host is only able to see the CCL.

**Related topics:**

- “Connection settings for the HP StorageWorks” on page 258



## Configuring settings for the HP StorageWorks

There are a large number of settings and options which are available through the HSG80 configuration interface. This topic and its subtopics discusses those options and settings that are supported with the SAN Volume Controller.

These options and settings can have a scope of a:

- Subsystem (global)
- Controller
- Port
- Logical unit
- Connection

### Related topics:

- “Configuring a balanced storage subsystem” on page 217

## Global settings for the HP StorageWorks

Global settings apply across a HSG80 subsystem. This topic lists any global settings.

Table 28. HSG80 global settings supported by the SAN Volume Controller Storage Software for Cisco MDS 9000

| Option                | HSG80 Default Setting | HSG80 SAN Volume Controller Storage Software for Cisco MDS 9000 Required Setting |
|-----------------------|-----------------------|----------------------------------------------------------------------------------|
| DRIVE_ERROR_THRESHOLD | 800                   | Default                                                                          |
| FAILEDSET             | Not defined           | n/a                                                                              |

### Related topics:

- “Configuring a balanced storage subsystem” on page 217

## Controller settings for the HP StorageWorks

Table 29 describes the options that can be set by HSG80 command-line interface (CLI) commands for each controller.

Table 29. HSG80 controller settings supported by the SAN Volume Controller Storage Software for Cisco MDS 9000

| Option               | HSG80 Default Setting | SAN Volume Controller Required Setting |
|----------------------|-----------------------|----------------------------------------|
| ALLOCATION_CLASS     | 0                     | Any value                              |
| CACHE_FLUSH_TIME     | 10                    | Any value                              |
| COMMMAND_CONSOLE_LUN | Not defined           | Any value                              |
| CONNECTIONS_UNLOCKED | CONNECTIONS_UNLOCKED  | CONNECTIONS_UNLOCKED                   |
| NOIDENTIFIER         | Not defined           | No identifier                          |
| MIRRORED_CACHE       | Not defined           | Mirrored                               |
| MULTIBUS_FAILOVER    | Not defined           | MULTIBUS_FAILOVER                      |

Table 29. HSG80 controller settings supported by the SAN Volume Controller Storage Software for Cisco MDS 9000 (continued)

| Option            | HSG80 Default Setting          | SAN Volume Controller Required Setting |
|-------------------|--------------------------------|----------------------------------------|
| NODE_ID           | Worldwide name as on the label | Default                                |
| PROMPT            | None                           | Any value                              |
| REMOTE_COPY       | Not defined                    | Any value                              |
| SCSI_VERSION      | SCSI-2                         | SCSI-3                                 |
| SMART_ERROR_EJECT | Disabled                       | Any value                              |
| TERMINAL_PARITY   | None                           | Any value                              |
| TERMINAL_SPEED    | 9600                           | Any value                              |
| TIME              | Not defined                    | Any value                              |
| UPS               | Not defined                    | Any value                              |

**Related topics:**

- “Configuring a balanced storage subsystem” on page 217
- “Port settings for the HP StorageWorks”

## Port settings for the HP StorageWorks

**Only one port per HSG80 pair may be used with the SAN Volume Controller Storage Software for Cisco MDS 9000.** Port settings are configurable at the port level. This topic lists those per port settings, what the HSG80 defaults are, and what the required settings are for the SAN Volume Controller Storage Software for Cisco MDS 9000.

**Note:** These port settings are set using the following commands:

- SET THIS PORT\_1\_TOPOLOGY=FABRIC
- SET THIS PORT\_2\_TOPOLOGY=FABRIC
- SET OTHER PORT\_1\_TOPOLOGY=FABRIC
- SET OTHER PORT\_2\_TOPOLOGY=FABRIC

These values can be checked using the following commands:

- SHOW THIS
- SHOW OTHER

Table 30. HSG80 port settings supported by the SAN Volume Controller Storage Software for Cisco MDS 9000

| Option            | HSG80 Default Setting | HSG80 SAN Volume Controller Required Setting |
|-------------------|-----------------------|----------------------------------------------|
| PORT_1/2-AL-PA    | 71 or 72              | n/a                                          |
| PORT_1/2_TOPOLOGY | Not defined           | FABRIC                                       |

**Note:** The HSG80 supports LUN masking using the “SET <unit number> ENABLE\_ACCESS\_PATH” command. When used with a SAN Volume Controller Storage Software for Cisco MDS 9000, ENABLE\_ACCESS\_PATH must be set to all (“SET <unit number> ENABLE\_ACCESS\_PATH=ALL”) and all LUN masking handled exclusively by the SAN Volume Controller Storage Software for Cisco MDS 9000. The access rights can be checked

using the "SHOW CONNECTIONS FULL" and any UNIT\_OFFSETs using the "SHOW CONNECTIONS FULL" command.

**Related topics:**

- "Configuring a balanced storage subsystem" on page 217
- "LU settings for the HP StorageWorks"

## LU settings for the HP StorageWorks

Table 31 describes the options that must be set for each logical unit that is accessed by the SAN Volume Controller Storage Software for Cisco MDS 9000. LUs that are accessed by hosts can be configured differently.

*Table 31. HSG80 LU settings supported by the SAN Volume Controller Storage Software for Cisco MDS 9000*

| Option                                     | HSG80 Default Setting   | SAN Volume Controller Storage Software for Cisco MDS 9000 Required Setting |
|--------------------------------------------|-------------------------|----------------------------------------------------------------------------|
| TRANSFER_RATE_REQUESTED                    | 20MHZ                   | n/a                                                                        |
| TRANSPORTABLE/<br>NOTTRANSPORTABLE         | NOTTRANSPORTABLE        | NOTTRANSPORTABLE                                                           |
| ENABLE_ACCESS_PATH/<br>DISABLE_ACCESS_PATH | ENABLE_ACCESS_PATH=ALL  | ENABLE_ACCESS_PATH=ALL                                                     |
| IDENTIFIER/ NOIDENTIFIER                   | NOIDENTIFIER            | n/a                                                                        |
| MAX_READ_CACHE_SIZE                        | 32                      | n/a                                                                        |
| MAX_WRITE_CACHE_SIZE                       | 32                      | 64 or higher                                                               |
| MAX_CACHED_TRANSFER_SIZE                   | 32                      | n/a                                                                        |
| PREFERRED_PATH/<br>NOPREFERRED_PATH        | NOPREFERRED_PATH is set | n/a                                                                        |
| READ_CACHE/ NOREAD_CACHE                   | READ_CACHE              | n/a                                                                        |
| READAHEAD_CACHE/<br>NOREADAHEAD_CACHE      | READAHEAD_CACHE         | n/a                                                                        |
| RUN/ NORUN                                 | RUN                     | RUN                                                                        |
| WRITE_LOG/NOWRITE_LOG                      | NOWRITE_LOG             | NOWRITE_LOG                                                                |
| WRITE_PROTECT/<br>NOWRITE_PROTECT          | NOWRITE_PROTECT         | NOWRITE_PROTECT                                                            |
| WRITEBACK_CACHE/<br>NOWRITEBACK_CACHE      | WRITEBACK_CACHE         | WRITEBACK_CACHE                                                            |

**Related topics:**

- "Configuring a balanced storage subsystem" on page 217

## Connection settings for the HP StorageWorks

Table 32. HSG80 connection options and their required settings for the SAN Volume Controller Storage Software for Cisco MDS 9000

| Option            | HSG80 Default Setting | HSG80 Required Setting |
|-------------------|-----------------------|------------------------|
| OPERATING_SYSTEM  | Not defined           | WINNT                  |
| RESERVATION_STYLE | CONNECTION_BASED      | n/a                    |
| UNIT_OFFSET       | 0                     | 0 or 199               |

### Related topics:

- “Configuring a balanced storage subsystem” on page 217
- “Configuration interface for the HP StorageWorks” on page 254

## Mapping and virtualization settings for the HP StorageWorks

The HSG80 configuration interface requires you assign a unit number to each logical unit when it is defined. By default the LUN is the unit number. It is therefore possible for gaps to exist in the LUN range if the unit numbers used in the configuration commands are not contiguous. By default, each LUN is visible on all controller ports on both controllers.

### LUN masking:

The HSG80 supports the concept of connection names. It supports a maximum of 96 connection names which contain the following parameters:

- HOST\_ID
- ADAPTER\_ID
- CONTROLLER
- PORT
- REJECTED\_HOST

**Note:** The SAN Volume Controller Storage Software for Cisco MDS 9000 ports should not be in the REJECTED\_HOSTS list. This list can be seen with SHOW CONNECTIONS FULL.

LUN masking must not be used on LUs which are in use by the SAN Volume Controller Storage Software for Cisco MDS 9000 to restrict either the initiator ports or the Target ports which the SAN Volume Controller Storage Software for Cisco MDS 9000 uses to access the LUs. Configurations that use LUN masking in this way are not supported. LUN masking can be used to prevent other initiators on the SAN from accessing LUs which are in use by the SAN Volume Controller Storage Software for Cisco MDS 9000 but the preferred method for this is to use SAN zoning.

### LU Virtualization:

The HSG80 also provides LU virtualization by the port and by the initiator. This is achieved by specifying a UNIT\_OFFSET for the connection. The use of LU Virtualization for connections between the HSG80 target ports and initiator ports on the SAN Volume Controller Storage Software for Cisco MDS 9000 is not supported.

---

## Appendix A. Valid combinations of FlashCopy and Remote Copy functions

Table 33. Valid combinations of FlashCopy and Remote Copy interactions

| FlashCopy        | Remote Copy Primary | Remote Copy Secondary |
|------------------|---------------------|-----------------------|
| FlashCopy source | Supported           | Supported             |
| FlashCopy target | Not supported       | Not supported         |

**Related topics:**

- “Remote Copy” on page 38



## Appendix B. Object types

The following table lists the object codes and its corresponding object type.

Table 34. Object types

| Object code | Object type                 |
|-------------|-----------------------------|
| 0           | IC_TYPE_Unknown             |
| 1           | IC_TYPE_Vlun                |
| 2           | IC_TYPE_Vlungrp             |
| 3           | IC_TYPE_Hlun                |
| 4           | IC_TYPE_Node                |
| 5           | IC_TYPE_Host                |
| 6           | IC_TYPE_Hostgrp             |
| 7           | IC_TYPE_Hws                 |
| 8           | IC_TYPE_Fcgrp               |
| 9           | IC_TYPE_Rcgrp               |
| 10          | IC_TYPE_Fcmap               |
| 11          | IC_TYPE_Rcmap               |
| 12          | IC_TYPE_Wwpn                |
| 13          | IC_TYPE_Cluster             |
| 15          | IC_TYPE_Hba                 |
| 16          | IC_TYPE_Device              |
| 17          | IC_TYPE_SCSILun             |
| 18          | IC_TYPE_Quorum              |
| 19          | IC_TYPE_TimeSeconds         |
| 20          | IC_TYPE_ExtInst             |
| 21          | IC_TYPE_ExtInst             |
| 22          | IC_TYPE_Percentage          |
| 23          | IC_TYPE_VPD_SystemBoard     |
| 24          | IC_TYPE_VPD_Processor       |
| 25          | IC_TYPE_VPD_Processor_Cache |
| 26          | IC_TYPE_VPD_Memory_Module   |
| 27          | IC_TYPE_VPD_Fan             |
| 28          | IC_TYPE_VPD_FC_Card         |
| 29          | IC_TYPE_VPD_FC_Device       |
| 30          | IC_TYPE_VPD_Software        |
| 31          | IC_TYPE_VPD_Front_Panel     |
| 33          | IC_TYPE_VPD_Port            |
| 34          | IC_TYPE_FC_Adapter          |
| 35          | IC_TYPE_Migrate             |





---

## Appendix C. Event codes

Configuration event codes are generated when configuration parameters are set. Configuration event codes are recorded in a separate log and do not generate SNMP traps or e-mails and their error fixed flags are ignored.

### Related topics:

- “Configuration event codes”

---

## Configuration event codes

Configuration event codes are generated when configuration parameters are set. Configuration event codes are recorded in a separate log and do not generate SNMP traps or e-mails and their error fixed flags are ignored.

Table 35. Configuration event codes

| Event code | Description                                                                                  |
|------------|----------------------------------------------------------------------------------------------|
| 990101     | Modify cluster (attributes in the <b>svctask chcluster</b> command)                          |
| 990105     | Delete node from cluster (attributes in the <b>svctask rmnode</b> command)                   |
| 990106     | Create host (attributes in the <b>svctask mkhost</b> command)                                |
| 990112     | Cluster config dumped to file (attributes in the <b>svctask dumpconfig</b> command)          |
| 990117     | Create cluster                                                                               |
| 990118     | Modify node (attributes in the <b>svctask chnode</b> command)                                |
| 990119     | Configure set controller name                                                                |
| 990120     | Shut down node (attributes in the <b>svctask stopcluster</b> command)                        |
| 990128     | Modify host (attributes in the <b>svctask chhost</b> command)                                |
| 990129     | Delete node (attributes in the <b>svctask rmnode</b> command)                                |
| 990138     | Virtual Disk Modify (attributes in the <b>svctask chvdisk</b> command)                       |
| 990140     | Virtual Disk Delete (attributes in the <b>svctask rmvdisk</b> command)                       |
| 990144     | Modify Managed Disk group (attributes in the <b>svctask chmdiskgrp</b> command)              |
| 990145     | Delete Managed Disk group (attributes in the <b>svctask rmdiskgrp</b> command)               |
| 990148     | Create Managed Disk group (attributes in the <b>svctask mkmdiskgrp</b> command)              |
| 990149     | Modify Managed Disk (attributes in the <b>svctask chmdisk</b> command)                       |
| 990158     | VLUN included                                                                                |
| 990159     | Quorum created                                                                               |
| 990160     | Quorum Destroy                                                                               |
| 990168     | Modify the HWS a Virtual Disk is assigned to                                                 |
| 990169     | Create a new Virtual Disk (attributes in the <b>svctask mkvdisk</b> command)                 |
| 990173     | Add a Managed Disk to Managed Disk group (attributes in the <b>svctask addmdisk</b> command) |

Table 35. Configuration event codes (continued)

| Event code | Description                                                                                           |
|------------|-------------------------------------------------------------------------------------------------------|
| 990174     | Delete a Managed Disk from Managed Disk group (attributes in the <b>svctask rmmdisk</b> command)      |
| 990178     | Add a port to a Host (attributes in the <b>svctask addhostport</b> command)                           |
| 990179     | Delete a port from a Host (attributes in the <b>svctask rmhostport</b> command)                       |
| 990182     | Create a Virtual Disk to Host SCSI mapping (attributes in the <b>svctask mkvdiskhostmap</b> command)  |
| 990183     | Delete an Virtual Disk to Host SCSI mapping (attributes in the <b>svctask rmdiskhostmap</b> command)  |
| 990184     | Create a FlashCopy mapping (attributes in the <b>svctask mkfcmap</b> command)                         |
| 990185     | Modify a FlashCopy mapping (attributes in the <b>svctask chfcmap</b> command)                         |
| 990186     | Delete a FlashCopy mapping (attributes in the <b>svctask rmfcmap</b> command)                         |
| 990187     | Prepare a FlashCopy mapping (attributes in the <b>svctask prestartfcmap</b> command)                  |
| 990188     | Prepare a FlashCopy consistency group (attributes in the <b>svctask prestartfcconsistgrp</b> command) |
| 990189     | Trigger a FlashCopy mapping (attributes in the <b>svctask startfcmap</b> command)                     |
| 990190     | Trigger a FlashCopy consistency group (attributes in the <b>svctask startfcconsistgrp</b> command)    |
| 990191     | Stop a FlashCopy mapping (attributes in the <b>svctask stopfcmap</b> command)                         |
| 990192     | Stop a FlashCopy consistency group (attributes in the <b>svctask stopfcconsistgrp</b> command)        |
| 990193     | FlashCopy set name                                                                                    |
| 990194     | Delete a list of ports from a Host (attributes in the <b>svctask rmhostport</b> command)              |
| 990196     | Shrink a Virtual Disk.                                                                                |
| 990197     | Expand a Virtual Disk (attributes in the <b>svctask expandvdisksize</b> command)                      |
| 990198     | Expand single extent a Virtual Disk                                                                   |
| 990199     | Modify govern a Virtual Disk                                                                          |
| 990203     | Initiate manual Managed Disk discovery (attributes in the <b>svctask detectmdisk</b> command)         |
| 990204     | Create FlashCopy consistency group (attributes in the <b>svctask mkfcconsistgrp</b> command)          |
| 990205     | Modify FlashCopy consistency group (attributes in the <b>svctask chfcconsistgrp</b> command)          |
| 990206     | Delete FlashCopy consistency group (attributes in the <b>svctask rmfcconsistgrp</b> command)          |
| 990207     | Delete a list of Hosts (attributes in the <b>svctask rmhost</b> command)                              |

Table 35. Configuration event codes (continued)

| Event code | Description                                                                                          |
|------------|------------------------------------------------------------------------------------------------------|
| 990213     | Change the HWS a node belongs to (attributes in the <b>svctask chiogrp</b> command)                  |
| 990216     | Apply software upgrade (attributes in the <b>svcservicetask applysoftware</b> command)               |
| 990219     | Analyze error log (attributes in the <b>svctask finderr</b> command)                                 |
| 990220     | Dump error log (attributes in the <b>svctask dumperrlog</b> command)                                 |
| 990221     | Clear error log (attributes in the <b>svctask clearerrlog</b> command)                               |
| 990222     | Fix error log entry (attributes in the <b>svctask cherrstate</b> command)                            |
| 990223     | Migrate a single extent (attributes in the <b>svctask migrateexts</b> command)                       |
| 990224     | Migrate a number of extents                                                                          |
| 990225     | Create Remote Copy relationship (attributes in the <b>svctask mkrcrelationship</b> command)          |
| 990226     | Modify Remote Copy relationship (attributes in the <b>svctask chrcrelationship</b> command)          |
| 990227     | Delete Remote Copy relationship (attributes in the <b>svctask rmrcrelationship</b> command)          |
| 990229     | Start Remote Copy relationship (attributes in the <b>svctask startrcrelationship</b> command)        |
| 990230     | Stop Remote Copy relationship (attributes in the <b>svctask stoprcrelationship</b> command)          |
| 990231     | Switch a Remote Copy relationship (attributes in the <b>svctask switchrcrelationship</b> command)    |
| 990232     | Start Remote Copy consistency group (attributes in the <b>svctask startrcconsistgrp</b> command)     |
| 990233     | Stop Remote Copy consistency group (attributes in the <b>svctask stoprcconsistgrp</b> command)       |
| 990234     | Switch a Remote Copy consistency group (attributes in the <b>svctask switchrcconsistgrp</b> command) |
| 990235     | Managed Disk migrated to a Managed Disk group                                                        |
| 990236     | Virtual Disk migrated to a new Managed Disk                                                          |
| 990237     | Create partnership with remote cluster (attributes in the <b>svctask mkpartnership</b> command)      |
| 990238     | Modify partnership with remote cluster (attributes in the <b>svctask chpartnership</b> command)      |
| 990239     | Delete partnership with remote cluster (attributes in the <b>svctask rmpartnership</b> command)      |
| 990240     | Create Remote Copy consistency group (attributes in the <b>svctask mkrcconsistgrp</b> command)       |
| 990241     | Modify Remote Copy consistency group (attributes in <b>svctask chrccconsistgrp</b> )                 |
| 990242     | Delete Remote Copy consistency group (attributes in the <b>svctask rmrcconsistgrp</b> command)       |
| 990245     | Node pend                                                                                            |
| 990246     | Node remove                                                                                          |

Table 35. Configuration event codes (continued)

| Event code | Description                                                                   |
|------------|-------------------------------------------------------------------------------|
| 990247     | Node unpend                                                                   |
| 990380     | Time zone changed (attributes in the <b>svctask settimezone</b> command)      |
| 990383     | Change cluster time (attributes in the <b>svctask setclustertime</b> command) |
| 990385     | System time changed                                                           |
| 990386     | SSH key added                                                                 |
| 990387     | SSH key removed                                                               |
| 990388     | All SSH keys removed                                                          |
| 990390     | Add node to the cluster                                                       |
| 990395     | Shutdown or reset node                                                        |
| 990410     | Software Install started                                                      |
| 990415     | Software Install completed                                                    |
| 990420     | Software Install failed                                                       |
| 990501     | The featurization has changed. See feature log for details.                   |
| 991024     | IO tracing has finished, trigger occurred for given Managed Disk.             |

**Related topics:**

- Appendix C, “Event codes,” on page 263

---

## Appendix D. Accessibility

### Features:

These are the major accessibility features in the SAN Volume Controller Storage Software for Cisco MDS 9000 console workstation:

- You can operate all features using the keyboard instead of the mouse.

### Navigating by keyboard:

You can use keys or key combinations to perform operations and initiate many menu actions that can also be done through mouse actions. You can navigate the Console for SAN Volume Controller for Cisco MDS 9000 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 Storage Software for Cisco MDS 9000 in Adobe Portable Document Format (PDF) using the Adobe Acrobat Reader. The PDFs are provided on a CD that is packaged with the product or you can access them at the following Web site:

<http://www.ibm.com/storage/support/2062-2300>

### Related topics:

- “Related publications” on page xi



---

## Notices

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

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

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

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

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

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATIONS "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

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

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

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

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

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

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

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

**Related topics:**

- "Trademarks"

---

## Trademarks

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

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

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

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

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

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

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



---

# Glossary

This glossary includes terms for the IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000.

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.

## A

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

## 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.** A communication service that links a machine 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. With access to the

machine, service personnel can perform service tasks, such as viewing error and problem logs or initiating trace and dump retrievals.

**cluster.** In SAN Volume Controller Storage Software for Cisco MDS 9000, two or four nodes that provide a single configuration and service interface.

**concurrent maintenance.** Service that is performed on a unit while it is operational.

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

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

**consistent copy.** In a Remote Copy 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.

**container.**

- IBM definition: A visual user-interface component that holds objects.
- HP definition:
  1. Any entity that is capable of storing data, whether it is a physical device or a group of physical devices.
  2. A virtual, internal controller structure representing either a single disk or a group of disk drives linked as a storageset. Stripesets and mirrorsets are examples of storageset containers that the controller uses to create units.

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

**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.

## 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.

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

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

**device.**

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

**directed maintenance procedures.** The set of maintenance procedures that can be run for a cluster. These procedures are documented in the service guide.

**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 zone.** A zone defined in the storage area network (SAN) fabric in which the SAN Volume Controller Storage Software for Cisco MDS 9000 can detect and address the logical units that the disk controllers present.

## E

**error code.** A value that identifies an error condition.

**ESS.** See *IBM TotalStorage Enterprise Storage Server*<sup>®</sup>.

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

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

**excluded.** In SAN Volume Controller Storage Software for Cisco MDS 9000, 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

**failover.** In SAN Volume Controller Storage Software for Cisco MDS 9000, the function that occurs when one redundant part of the system takes over the workload of another part of the system that has failed.

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

**fibre-channel extender.** A long-distance communication device that interconnects storage area network (SAN) fabric components.

**FC.** See *fibre channel*.

**FlashCopy service.** In SAN Volume Controller Storage Software for Cisco MDS 9000, 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*.

**FlashCopy mapping.** A relationship between two virtual disks.

**FlashCopy relationship.** See *FlashCopy mapping*.

## H

**HBA.** See *host bus adapter*.

**host bus adapter (HBA).** In SAN Volume Controller Storage Software for Cisco MDS 9000, an interface card that connects a host bus, such as a peripheral component interconnect (PCI) bus, to the storage area network.

**host.** An open-systems computer that is connected to the SAN Volume Controller Storage Software for Cisco MDS 9000 through a fibre-channel interface.

**host ID.** In SAN Volume Controller Storage Software for Cisco MDS 9000, 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 Controller Storage Software for Cisco MDS 9000s.

## I

**IBM Subsystem Device Driver (SDD).** An IBM pseudo device driver designed to support the multipath configuration environments in IBM products.

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

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

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

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

**inconsistent.** In a Remote Copy relationship, pertaining to a secondary virtual disk (VDisk) that is being synchronized with the primary VDisk.

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

**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.

**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*.

## 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*.

**local fabric.** In SAN Volume Controller Storage Software for Cisco MDS 9000, 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)

**LU.** See *logical unit*.

**LUN.** See *logical unit number*.

## M

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

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

**mapping.** See *FlashCopy mapping*.

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

**MDisk.** See *managed disk*.

**migration.** See *data migration*.

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

## N

**node.** An addressable entity connected to an I/O bus or network. Used primarily to refer to computers, storage devices, and storage subsystems. The component of a node that connects to the bus or network is a port. (S)

**node rescue.** In SAN Volume Controller Storage Software for Cisco MDS 9000, the process by which a SAN Volume Controller Storage Software for Cisco MDS 9000 node that has no valid software installed on its hard disk drive can copy the software from another SAN Volume Controller Storage Software for Cisco MDS 9000 node connected to the same fibre-channel fabric.

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

## O

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

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

## P

**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.

**partnership.** In Remote Copy, 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 Storage Software for Cisco MDS 9000, 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.

**primary virtual disk.** In a Remote Copy relationship, the target of write operations issued by the host application.

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

## Q

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

## R

**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 storageset 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 storagesets are referred to as stripesets.

**RAID 1.** SNIA dictionary definition: A form of storage array in which two or more identical copies of data are maintained on separate media. IBM definition: A form of storage array in which two or more identical copies of data are maintained on separate media. Also known as mirrorset. HP definition: See *mirrorset*.

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

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

**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.

**redundant SAN.** A storage area network (SAN) configuration in which any one single component might fail, but connectivity between the devices within the SAN is maintained, possibly with degraded performance. This configuration is normally achieved by splitting the SAN into two, independent, counterpart SANs. See also *counterpart SAN*.

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

**relationship.** In Remote Copy, the association between a master virtual disk (VDisk) and an auxiliary VDisk. These VDIs 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*.

**Remote Copy.** In SAN Volume Controller Storage Software for Cisco MDS 9000, a copy service that enables host data on a particular source virtual disk (VDisk) to be copied to the target VDisk designated in the relationship.

## S

**SAN.** See *storage area network*.

**SAN Volume Controller Storage Software for Cisco MDS 9000 fibre-channel port fan in.** The number of hosts that can see any one SAN Volume Controller Storage Software for Cisco MDS 9000 port.

**SCSI.** See *Small Computer Systems Interface*.

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

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

**secondary virtual disk.** In Remote Copy, the virtual disk (VDisk) in a relationship that contains a copy of data written by the host application to the primary VDisk.

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

**SNMP.** See *Simple Network Management Protocol*.

**stand-alone relationship.** In FlashCopy and Remote Copy, 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)

**stripeset.** See *RAID 0*.

**superuser authority.** The level of access required to add users.

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

**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 Remote Copy, the status condition that exists when both virtual disks (VDisks) of a pair that has a copy relationship contain the same data.

## T

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

## U

**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 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.

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

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

## V

**valid configuration.** A configuration that is supported.

**VDisk.** See *virtual disk*.

**virtual disk (VDisk).** A set of disk blocks presented to an operating environment as a range of consecutively numbered logical blocks with disk-like storage and I/O semantics. The virtual disk is the disk array object that most closely resembles a physical disk from the operating environment's viewpoint. (S)

**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.

**vital product data (VPD).** Information which uniquely defines each hardware element and firmware element in the subsystem.

## W

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

**WWNN.** See *worldwide node name*.

**WWPN.** See *worldwide port name*.

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

---

# Index

## A

- about this guide xi
- accessibility 267
  - keyboard 267
  - shortcut keys 267
- adding
  - storage controllers
    - using the CLI (command-line interface) 183
    - using the Console for SAN Volume Controller for Cisco MDS 9000 99
- advanced functions
  - overview
    - using CLI (command-line interface) 194
    - using Console for SAN Volume Controller for Cisco MDS 9000 89, 105
  - Remote Copy
    - using the CLI (command-line interface) 194
    - using the Console for SAN Volume Controller for Cisco MDS 9000 104
- analyzing error logs 110, 198
- audience xi

## B

- book
  - about this xi

## C

- Cisco Systems
  - MDS 9000 Caching Services Module 3
  - MDS 9000 switch 3, 12
- CLI (command-line interface)
  - examples 153
  - getting started 147
  - issuing commands from a PuTTY SSH client
    - system 147
  - scenarios 153
  - upgrading software 201
  - using to set cluster features 150
- clusters
  - configuring
    - using the CLI (command-line interface) 194
    - using the Console for SAN Volume Controller for Cisco MDS 9000 105
  - overview 12
  - resetting the SSH fingerprint 107
  - setting
    - features 150
  - shutting down 110, 198
- codes
  - configuration events 263
  - events 263
- command-line interface (CLI)
  - examples 153
  - getting started 147

- command-line interface (CLI) (*continued*)
  - issuing commands from a PuTTY SSH client
    - system 147
  - scenarios 153
  - upgrading software 201
  - using to set cluster features 150
- commands
  - svcinfo caterrlog 211
  - svcinfo caterrlogbyseqnum 211
  - svcinfo lserrlogbyxxxx 211
  - svcinfo lsxxxx 211
  - svcinfo lsxxxxcandidate 211
  - svcinfo lsxxxxdumps 211
  - svcinfo lsxxxxextent 211
  - svcinfo lsxxxxmember 211
  - svcinfo lsxxxxprogress 211
  - svcservicetask rmnode 211
  - svctask rmnode 211
- communications
  - determining between hosts and virtual disks 171
- configuration
  - event codes 263
  - rules 43
- configuring
  - clusters 67, 150
    - using the CLI (command-line interface) 194
    - using the Console for SAN Volume Controller for Cisco MDS 9000 105
  - disk controllers 217
  - Enterprise Storage Server 217, 231
  - FASTT Storage Manager 217, 236
  - FASTT Storage Server 217, 235
  - PuTTY 54
  - secure shell (SSH) 54
  - SSH (secure shell) 54
  - switches 46
- consistency group, Remote Copy 41
- consistency groups, FlashCopy 33
  - creating 83
  - deleting 87
  - starting 86
  - stopping 86
- console
  - master
    - overview 53
    - upgrading software 55
  - SAN Volume Controller Storage Software for Cisco MDS 9000
    - banner area 60
    - layout 60
    - portfolio 60
    - starting 59
    - task bar 60
    - work area 61
- console workstation
  - upgrading software 55

- controllers
  - adding
    - using the CLI (command-line interface) 183
    - using the Console for SAN Volume Controller for Cisco MDS 9000 99
  - advanced functions
    - EMC CLARiiON 225
    - Enterprise Storage Server 233
    - FASTT 238
    - HDS Lightning 244
    - HP StorageWorks 253
  - concurrent maintenance
    - EMC CLARiiON 224
    - Enterprise Storage Server 232
    - FASTT 237
    - HDS Lightning 243
    - HP StorageWorks 251
  - configuration
    - EMC CLARiiON 221, 226
    - Enterprise Storage Server 231
    - FASTT 235
    - HDS Lightning 243
    - HP StorageWorks 247, 249, 255
  - controller settings
    - EMC CLARiiON 227
  - firmware
    - EMC CLARiiON 223
    - Enterprise Storage Server 232
    - FASTT 237
    - HDS Lightning 243
    - HP StorageWorks 251
  - global settings
    - EMC CLARiiON 226
    - FASTT 241
  - interface
    - FASTT 239
    - HP StorageWorks 254
  - logical unit creation and deletion
    - EMC CLARiiON 226
    - Enterprise Storage Server 233
    - FASTT 239
    - HP StorageWorks 253
  - LU settings
    - EMC CLARiiON 228
    - FASTT 242
    - HP StorageWorks 257
  - models
    - EMC CLARiiON 223
    - Enterprise Storage Server 232
    - FASTT 237
    - HDS Lightning 243
    - HP StorageWorks 251
  - port settings
    - EMC CLARiiON 227
    - HP StorageWorks 256
  - quorum disks
    - EMC CLARiiON 225, 252
    - Enterprise Storage Server 233
    - FASTT 238
    - HDS Lightning 244

- controllers (*continued*)
  - registering
    - EMC CLARiiON 221
  - removing
    - using the CLI (command-line interface) 184
    - using the Console for SAN Volume Controller for Cisco MDS 9000 100
  - settings
    - FASTT 240, 241
    - HP StorageWorks 255, 258
  - sharing
    - EMC CLARiiON 224
    - Enterprise Storage Server 232
    - FASTT 237
    - HDS Lightning 243
    - HP StorageWorks 252
  - storage groups
    - EMC CLARiiON 222
  - switch zoning
    - EMC CLARiiON 225
    - Enterprise Storage Server 233
    - HP StorageWorks 252
  - conventions
    - emphasis in text xiv
    - numbering xiv
  - creating
    - clusters
      - from the Console for SAN Volume Controller for Cisco MDS 9000 63
    - FlashCopy
      - mappings 84, 168, 169
    - VDisk-to-host mappings 82
    - virtual disk-to-host mappings 166

## D

- data migration
  - FASTT 238
- deleting
  - FlashCopy
    - mappings 85
  - hosts 98
  - nodes 194
- determining
  - communications between hosts and virtual disks 171
- disability 267
- discovering
  - managed disks 158, 161
- disk controllers
  - overview 16
- disks
  - migrating 191

## E

- emphasis in text xiv
- errors
  - notification settings 107
- events
  - codes 263



- events *(continued)*
  - configuration 263
- examples
  - using the CLI (command-line interface) 153
  - using the Console for SAN Volume Controller for Cisco MDS 9000 73
- expanding
  - virtual disks 185
- extents
  - migrating
    - using the CLI (command-line interface) 189

## F

- features
  - setting
    - using the CLI (command-line interface) 150
  - viewing logs 107, 109
- FlashCopy
  - consistency groups 33
  - mappings 29, 167
  - overview 29

## G

- general cluster properties
  - viewing 70, 151
- getting started
  - using SAN Volume Controller Storage Software for Cisco MDS 9000 59, 115
  - using the CLI (command-line interface) 147
  - using the command-line interface (CLI) 147
- glossary 271
- groups
  - managed disk 20
- guide
  - about this xi
  - who should read xi

## H

- HBAs (host bus adapters)
  - configuration 45
- host bus adapters (HBAs)
  - configuration 45
- hosts
  - creating 82
  - deleting 98
  - overview 24

## I

- I/O groups
  - overview 14
- image-mode VDisks
  - converting to managed mode
    - using CLI (command-line interface) 193
    - using Console for SAN Volume Controller for Cisco MDS 9000 103

- information
  - center xi
- installing
  - SAN Volume Controller Storage Software for Cisco MDS 9000 123, 129
  - software 211
  - verification 135
- IP addresses
  - modifying 108
- issuing
  - CLI commands 147

## K

- keyboard 267
  - shortcut keys 267

## L

- listing
  - dump files 109, 196
  - log files 109, 196

## M

- managed disk (MDisk) groups
  - creating 80
  - overview 20
  - status 20
- managed disks (MDisks)
  - access modes 18
  - description 18
  - extents 18
  - overview 18
  - status 18
- managed mode virtual disks
  - converting from image mode
    - using the CLI (command-line interface) 193
    - using the Console for SAN Volume Controller for Cisco MDS 9000 103
- mappings, FlashCopy
  - creating 84
  - deleting 85
  - starting 85
  - stopping 85
- MDisk (managed disk) groups
  - description 20
  - overview 20
  - status 20
- MDisks (managed disks)
  - access modes 18
  - description 18
  - extents 18
  - overview 18
  - status 18
- measurements xiv
- migrating
  - extents
    - using the CLI (command-line interface) 189
- migration 102, 238

monitoring  
software upgrades, automatic 205, 207

## N

nodes  
adding 74, 154  
configuration 13, 46  
overview 12  
status 13  
viewing  
general details 79, 156

notices  
legal 269

## O

ordering publications xiii

overview  
advanced functions  
using the CLI (command-line interface) 194  
using the Console for SAN Volume Controller for  
Cisco MDS 9000 89, 105  
creating a cluster 63  
managed disk groups 20

## P

plink utility  
running 148

preinstalled software  
recovering from installation failures 213

publications  
ordering xiii

PuTTY 54  
configuring 54  
issuing CLI commands from 147  
running the plink utility 148

## R

related information xi

relationships, Remote Copy  
overview 39

Remote Copy  
overview 38, 41  
using the CLI (command-line interface) 194  
using the Console for SAN Volume Controller for  
Cisco MDS 9000 104

partnerships 39

removing  
nodes 105  
storage controllers  
using the CLI (command-line interface) 184  
using the Console for SAN Volume Controller for  
Cisco MDS 9000 100

requirements 115, 117, 119, 121

resetting  
SSH fingerprint for a cluster 107

running  
PuTTY plink utility 148

## S

SAN Volume Controller Storage Software for Cisco  
MDS 9000  
advanced functions 89

Console  
banner area 60  
examples 73  
layout 60  
portfolio 60  
post installation tasks 137  
scenarios 73  
starting 59  
task bar 60  
using to create a cluster 63  
work area 61

launching the Web application 70

overview 3

removing 141

scenarios  
using the CLI (command-line interface) 153  
using the Console for SAN Volume Controller for  
Cisco MDS 9000 73

secure shell (SSH) 53  
client system  
issuing CLI commands from 147  
configuring 54

security  
overview 53

setting  
cluster features  
using the CLI (command-line interface) 150  
features  
using the CLI (command-line interface) 150

settings  
error notification 196

shortcut keys 267

shrinking  
VDisks 101

shutting down  
clusters 110

software  
description 115, 211  
installing 211  
upgrading 203, 209

software, upgrading  
console workstation 55  
using the CLI (command-line interface) 201

SSH (secure shell) 53  
client system  
issuing CLI commands from 147  
configuring 54  
resetting fingerprint 107

starting  
FlashCopy  
consistency groups 86  
mappings 85

- stopping
  - FlashCopy
    - mappings 85
  - Remote Copy
    - consistency groups 86
- storage controllers
  - adding
    - using the CLI (command-line interface) 183
    - using the Console for SAN Volume Controller for Cisco MDS 9000 99
  - removing
    - using the CLI (command-line interface) 184
    - using the Console for SAN Volume Controller for Cisco MDS 9000 100
- strategy
  - software upgrade
    - using the CLI (command-line interface) 201
- support
  - Web sites xiv
- synchronous copy
  - overview 39

## T

- text emphasis xiv
- trademarks 270

## U

- uninterruptible power supply
  - overview 16
- upgrading software
  - console workstation 55
  - strategy
    - using the CLI (command-line interface) 201
- using
  - object classes and instances 261

## V

- VDisks (virtual disks)
  - converting
    - from image mode to managed mode 103, 193
  - creating 81, 162
  - creating VDisk-to-host mappings 82
  - creating virtual disk-to-host mappings 82
  - expanding 185
  - migrating 192
  - modes
    - image 22
    - sequential 22
    - striped 22
  - shrinking 101
  - status 22
- viewing
  - clusters
    - feature logs 109, 197
- virtual disk-to-host mapping
  - description 25

- virtual disks (VDisks)
  - converting
    - from image mode to managed mode 103, 193
  - migrating 102
  - modes
    - image 22
    - sequential 22
    - shrinking 101
    - striped 22
- virtualization
  - asymmetric 8
  - overview 6
  - symmetric 9

## W

- Web sites xiv
- who should read this guide xi



---

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

IBM TotalStorage SAN Volume Controller  
Storage Software for Cisco MDS 9000  
Configuration Guide  
Version 1.1.8

Publication No. SC26-7554-01

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

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

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

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

Please tell us how we can improve this book:

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

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

---

Name

---

Address

---

Company or Organization

---

Phone No.



Fold and Tape

Please do not staple

Fold and Tape



NO POSTAGE  
NECESSARY  
IF MAILED IN THE  
UNITED STATES

# BUSINESS REPLY MAIL

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

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation  
Information Development  
Department 61C  
9032 South Rita Road  
Tucson, Arizona  
USA 85775-4401



Fold and Tape

Please do not staple

Fold and Tape





Part Number: 64P7905

Printed in USA

SC26-7554-01



(1P) P/N: 64P7905

