

IBM TotalStorage SAN Volume Controller



Configuration Guide

Version 3.1.0

IBM TotalStorage SAN Volume Controller



Configuration Guide

Version 3.1.0

Sixth Edition (October 2005)

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

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

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

Contents

Figures.	xi
Tables.	xiii
About this guide.	xv
Who should use this guide.	xv
Summary of changes.	xv
Summary of changes for SC26-7543-05 SAN Volume Controller Configuration Guide	xv
Summary of changes for SC26-7543-04 SAN Volume Controller Configuration Guide	xvii
Emphasis.	xix
Numbering conventions	xix
SAN Volume Controller library and related publications	xx
Related Web sites	xxi
How to order IBM publications	xxi
How to send your comments	xxii
Chapter 1. SAN Volume Controller overview	1
SAN fabric overview	3
SAN Volume Controller operating environment.	5
Virtualization	6
Asymmetric virtualization.	8
Symmetric virtualization	9
Physical links between SAN Volume Controller nodes and a switch	10
Support for long links between the local and remote fabric	11
Object overview	11
Nodes and clusters	13
I/O groups and uninterruptible power supply	17
Storage subsystems and MDisks	22
MDisk groups and VDIsks	26
Copy Services	35
FlashCopy	35
Metro Mirror	47
Configuration rules and requirements.	52
Configuration rules	54
Switch zoning for the SAN Volume Controller.	65
Limiting queue depth in large SANs	74
Configuration requirements	76
Maximum configuration	77
Supported fibre-channel extenders.	78
Performance of fibre-channel extenders.	78
Chapter 2. Creating a SAN Volume Controller cluster	79
Creating a cluster from the front panel	79
SAN Volume Controller Console layout	81
SAN Volume Controller Console banner.	82
SAN Volume Controller Console task bar	83
SAN Volume Controller Console portfolio	83
SAN Volume Controller Console work area	84
Browser requirements for the SAN Volume Controller.	84
Configuring the Web browser	85
Changing browser settings for password protection	85

Accessing the SAN Volume Controller Console	85
Creating a cluster using the SAN Volume Controller Console	86
Chapter 3. Using the SAN Volume Controller Console	93
Launching the SAN Volume Controller Console application.	93
Setting the cluster time	94
Modifying the cluster IP addresses	95
Maintaining cluster passwords	95
Viewing cluster properties	96
Adding nodes to a cluster	96
Viewing node properties	98
Increasing the size of a cluster	99
Adding a node to increase the size of a cluster	99
Migrating a VDisk to a new I/O group	101
Replacing a faulty node with a spare node	101
Renaming a node	105
Deleting a node from a cluster.	105
Renaming an I/O group	107
Renaming a cluster.	107
Shutting down a cluster	108
Shutting down a node	109
Discovering MDisks.	109
Renaming MDisks	110
Adding excluded MDisks to a cluster	110
Setting quorum disks	110
Determining the relationship between MDisks and VDIsks.	111
Determining the relationship between MDisks and RAID arrays or LUNs	111
Displaying MDisk groups	112
Creating MDisk groups	112
Adding MDisks to MDisk groups	113
Removing MDisks from an MDisk group	113
Viewing the progress of an MDisk removal	113
Renaming MDisk groups	114
Displaying VDIsks	114
Deleting MDisk groups.	114
Creating VDIsks	115
Viewing the progress of VDisk formatting	115
Migrating VDIsks	115
Viewing the progress of VDisk migration	117
Shrinking VDIsks	117
Viewing virtual disk-to-host mappings	118
Determining the relationship between VDIsks and MDisks.	118
Recovering from offline VDIsks	118
Deleting VDIsks	120
Using image mode VDIsks	120
Creating an image mode VDisk	121
Migration methods	122
Viewing the progress of image mode migration	123
Viewing the progress of extent migration	123
Creating hosts.	123
Filtering hosts	124
Viewing general details	124
Viewing port details.	124
Viewing mapped I/O groups.	125
Displaying VDIsks that are mapped to a host	125
Modifying a host	125

I	Adding ports to a host	126
I	Deleting ports from a host	126
	Replacing an HBA in a host.	127
	Deleting hosts.	127
	Creating FlashCopy mappings	128
	Filtering FlashCopy mappings	128
	Starting FlashCopy mappings	128
I	Viewing the progress of a FlashCopy	129
	Stopping FlashCopy mappings	129
	Modifying FlashCopy mappings	129
	Deleting FlashCopy mappings	129
	Creating FlashCopy consistency groups	130
	Filtering FlashCopy consistency groups	130
	Starting FlashCopy consistency groups	131
	Stopping FlashCopy consistency groups	131
	Renaming FlashCopy consistency groups	131
	Deleting FlashCopy consistency groups	132
	Creating Metro Mirror relationships	132
	Filtering Metro Mirror relationships	132
	Starting a Metro Mirror copy	133
I	Viewing the progress of Metro Mirror copy processes	133
	Stopping a Metro Mirror copy	133
	Modifying Metro Mirror relationships.	134
	Switching the copy direction of a Metro Mirror relationship	134
	Deleting Metro Mirror relationships	134
	Creating Metro Mirror consistency groups	135
	Filtering Metro Mirror consistency groups	135
	Renaming a Metro Mirror consistency group	135
	Starting a Metro Mirror consistency group copy	135
	Stopping a Metro Mirror consistency group copy	136
	Deleting Metro Mirror consistency groups.	136
	Creating Metro Mirror partnerships	136
	Modifying Metro Mirror partnerships	137
	Deleting Metro Mirror partnerships	137
	Viewing the feature log	138
	Viewing and updating feature settings	138
	Running the cluster maintenance procedure.	138
	Configuring error notification settings	138
	Displaying and saving log and dump files.	139
	Analyzing the error log	140
	Recovering a node and returning it to the cluster	141
	Managing SSH keys	142
	Adding SSH keys for hosts other than the master console	142
	Adding subsequent SSH public keys to the SAN Volume Controller	143
	Replacing the client SSH private key known to the SAN Volume Controller software	143
	Replacing the SSH key pair.	144
	Resetting a refused SSH key	144
	Resetting the SSH fingerprint	145
	Chapter 4. Using the CLI	147
	Preparing the SSH client system	147
	Preparing the SSH client system to issue CLI commands.	148
	Preparing the SSH client on an AIX host	149
	Issuing CLI commands from a PuTTY SSH client system	150
	Running the PuTTY and plink utilities	150

Setting the cluster time using the CLI	152
Reviewing and setting the cluster features using the CLI	152
Displaying cluster properties using the CLI	153
Maintaining passwords for the front panel using the CLI	153
Adding nodes to a cluster using the CLI	153
Displaying node properties using the CLI	156
Discovering MDisks using the CLI	157
Creating MDisk groups using the CLI	158
Adding MDisks to MDisk groups using the CLI	160
Creating VDIs	161
Creating host objects using the CLI	162
Creating VDisk-to-host mappings using the CLI	164
Creating FlashCopy mappings using the CLI	164
Creating a FlashCopy consistency group and adding mappings using the CLI	165
Preparing and triggering a FlashCopy mapping using the CLI	166
Preparing and triggering a FlashCopy consistency group using the CLI	167
Determining the WWPNs of a node using the CLI	168
Determining the VDisk name from the device identifier on the host	168
Determining the host that a VDisk is mapped to	169
Determining the relationship between VDIs and MDisks using the CLI	169
Determining the relationship between MDisks and RAID arrays or LUNs using the CLI	170
Increasing the size of your cluster using the CLI	170
Adding a node to increase the size of a cluster using the CLI	171
Migrating a VDisk to a new I/O group using the CLI	171
Replacing a faulty node in the cluster using the CLI	172
Recovering from offline VDIs using the CLI	176
Recovering a node and returning it to the cluster using the CLI	177
Moving offline VDIs to the recovery I/O group using the CLI	178
Moving offline VDIs to their original I/O group using the CLI	178
Informing the SAN Volume Controller of changes to host HBAs using the CLI	178
Expanding VDIs	179
Expanding a VDisk that is mapped to an AIX host	180
Expanding a VDisk that is mapped to a Windows 2000 host using the CLI	180
Shrinking a virtual disk using the CLI	181
Migrating extents using the CLI	182
Migrating VDIs between MDisk groups using the CLI	184
Migrating a VDisk between I/O groups using the CLI	185
Creating an image mode VDisk using the CLI	186
Migrating to an image mode virtual disk using the CLI	187
Deleting a node from a cluster using the CLI	187
Performing the cluster maintenance procedure using the CLI	188
Changing the cluster IP address using the CLI	189
Changing the cluster gateway address using the CLI	189
Changing the cluster subnet mask using the CLI	189
Maintaining SSH keys using the CLI	190
Setting up error notifications using the CLI	190
Changing cluster passwords using the CLI	191
Changing the language setting using the CLI	191
Viewing the feature log using the CLI	192
Analyzing the error log using the CLI	192
Shutting down a cluster using the CLI	193
Chapter 5. Backing up and restoring the cluster configuration	195
Backing up the cluster configuration.	195
Backing up the cluster configuration using the CLI	196

Downloading backup configuration data files	198
Restoring the cluster configuration using the CLI	199
Deleting backup configuration files	201
Deleting backup configuration files using the CLI	202
Chapter 6. Upgrading the SAN Volume Controller software	203
Installing or upgrading the SAN Volume Controller software	204
Copying the SAN Volume Controller software upgrade files using PuTTY scp	205
Upgrading the SAN Volume Controller software automatically	206
Upgrading the SAN Volume Controller software using the SAN Volume Controller Console	206
Upgrading the SAN Volume Controller software using the CLI	209
Performing a disruptive software upgrade using the CLI	211
Performing the node rescue.	212
Recovering from software upgrade problems automatically	213
Recovering from software upgrade problems manually	213
Chapter 7. Configuring and servicing storage subsystems	215
Identifying your storage subsystem	215
Configuration guidelines	215
Storage subsystem logical disks	216
RAID array configuration	217
Optimal MDisk group configurations.	217
Considerations for FlashCopy mappings	218
Image mode and migrating existing data	219
Configuring a balanced storage subsystem	222
Discovering logical units	225
Expanding a logical unit using the CLI	225
Modifying a logical unit mapping using the CLI.	226
Accessing controller devices with multiple remote ports	227
Determining a storage subsystem name from its SAN Volume Controller name	228
Determining a storage subsystem name from its SAN Volume Controller name using the CLI	229
Renaming a storage subsystem	229
Changing a configuration for an existing storage subsystem using the CLI	229
Adding a new storage controller to a running configuration	230
Adding a new storage controller to a running configuration using the CLI	231
Removing a storage subsystem	233
Removing a storage subsystem using the CLI	234
Removing MDisks that represent unconfigured LUs using the CLI.	235
Creating a quorum disk	236
Manual discovery	237
Servicing storage subsystems	237
Configuring the EMC CLARiiON subsystem	238
Access Logix	239
Configuring the EMC CLARiiON controller with Access Logix installed	239
Configuring the EMC CLARiiON controller without Access Logix installed	242
Supported models of the EMC CLARiiON	242
Supported firmware levels for the EMC CLARiiON	243
Concurrent maintenance on the EMC CLARiiON	243
User interface on EMC CLARiiON	244
Sharing the EMC CLARiiON between a host and the SAN Volume Controller	244
Switch zoning limitations for the EMC CLARiiON	244
Quorum disks on the EMC CLARiiON	245
Advanced functions for the EMC CLARiiON	246
Logical unit creation and deletion on the EMC CLARiiON	246

Configuring settings for the EMC CLARiiON	246
Configuring the EMC Symmetrix subsystem	250
Supported models of the EMC Symmetrix controller	250
Supported firmware levels for the EMC Symmetrix controller	251
Concurrent maintenance on the EMC Symmetrix	251
User interface on EMC Symmetrix	252
Sharing the EMC Symmetrix controller between a host and a SAN Volume Controller	252
Switch zoning limitations for the EMC Symmetrix	253
Quorum disks on EMC Symmetrix	254
Advanced functions for EMC Symmetrix	254
Logical unit creation and deletion on EMC Symmetrix	254
Configuring settings for the EMC Symmetrix	255
Configuring the IBM TotalStorage Enterprise Storage Server	258
Configuring the IBM ESS	259
Supported models of the IBM ESS	260
Supported firmware levels for the IBM ESS	260
Concurrent maintenance on the IBM ESS	261
User interface on the IBM ESS	261
Target port groups for the IBM ESS	261
Sharing the IBM ESS between a host and the SAN Volume Controller	261
Switch zoning limitations for the IBM ESS	262
Quorum disks on the IBM ESS	262
Advanced functions for the IBM ESS	262
Logical unit creation and deletion on the IBM ESS	263
Configuring the IBM TotalStorage DS4000 (formerly FASTT) series subsystem	263
Configuring IBM DS4000 series disk controllers for the storage server	264
Supported options of the IBM DS4000 series controller	266
Supported models of the IBM DS4000 series of controllers	267
Supported firmware levels for the IBM DS4000 series	268
Concurrent maintenance on the IBM DS4000 series	269
User interface on the IBM DS4000 series	269
Sharing the IBM DS4000 series controller between a host and the SAN Volume Controller	269
Quorum disks on the IBM DS4000 series	270
Advanced functions for the IBM DS4000 series	270
Logical unit creation and deletion on the IBM DS4000 series	272
Configuration interface for the IBM DS4000 series	272
Controller settings for the IBM DS4000 series	273
Configuring the HDS Lightning series subsystem	276
Supported models of the HDS Lightning	277
Supported firmware levels for HDS Lightning	277
Concurrent maintenance on the HDS Lightning	278
User interface on HDS Lightning	278
Sharing the HDS Lightning 99xxV between a host and the SAN Volume Controller	278
Quorum disks on HDS Lightning 99xxV	280
Advanced functions for HDS Lightning	280
Logical unit configuration for HDS Lightning	281
Configuring settings for HDS Lightning	282
Configuring the HDS Thunder subsystem	284
Supported models of the HDS Thunder	285
Supported firmware levels for HDS Thunder	286
Concurrent maintenance on the HDS Thunder	286
User interface on the HDS Thunder	286
Sharing the HDS Thunder between host and the SAN Volume Controller	287

Setting up an HDS Thunder with more than four ports	288
Quorum disks on HDS Thunder	288
Advanced functions for HDS Thunder	289
Logical unit creation and deletion on HDS Thunder	290
Configuring settings for HDS Thunder	291
Configuring the HP StorageWorks subsystem	296
HP StorageWorks definitions	298
Configuring the HP StorageWorks controller.	300
Supported models of the HP StorageWorks controller	303
Supported firmware levels for the HP StorageWorks controller	304
Concurrent maintenance on the HP StorageWorks	304
Configuration interface for the HP StorageWorks	305
Sharing the HP StorageWorks controller between a host and the SAN Volume Controller	305
Switch zoning limitations for the HP StorageWorks subsystem	306
Quorum disks on HP StorageWorks.	306
Advanced functions for HP StorageWorks	307
HP StorageWorks advanced functions	308
Logical unit creation and deletion on the HP StorageWorks	308
Configuring settings for the HP StorageWorks	308
Configuring the HPQ Enterprise Virtual Array subsystem	314
Supported models of the HPQ EVA	314
Supported firmware levels for HPQ EVA	315
Concurrent maintenance on the HPQ EVA	315
User interface on HPQ EVA.	315
Sharing the HPQ EVA controller between a host and the SAN Volume Controller	316
Switch zoning limitations for the HPQ EVA subsystem	316
Quorum disks on HPQ EVA.	316
Advanced functions for HPQ EVA	316
Logical unit configuration on the HPQ EVA	316
Logical unit presentation	317
Configuration interface for the HPQ EVA	318
Configuring settings for the HPQ EVA	318
Chapter 8. Installing or upgrading the SAN Volume Controller Console overview.	321
SAN Volume Controller Console hardware installation requirements	322
SAN Volume Controller Console workstation space requirements	322
SAN Volume Controller Console software installation requirements	323
Installing or upgrading the SAN Volume Controller Console in graphical mode	323
Installing or upgrading the SAN Volume Controller Console in unattended (silent) mode	328
Verifying the Windows services associated with the SAN Volume Controller Console	333
Post installation tasks	334
Removing the SAN Volume Controller Console	336
Chapter 9. IBM TotalStorage support for Microsoft Volume Shadow Copy Service	339
Installation overview	339
System requirements for the IBM TotalStorage hardware provider.	340
Installing the IBM TotalStorage hardware provider	340
Verifying the installation	346
Creating the free and reserved pools of volumes	347
Changing the configuration parameters	348

Adding and removing volumes.	349
Error codes.	350
Uninstalling the IBM TotalStorage hardware provider	352
Appendix A. Configuring IBM Director overview	353
Setting up an event action plan	353
Setting up an e-mail notification to IBM	354
Setting up an e-mail user notification	355
Setting up SNMP traps	356
Appendix B. Event codes	357
Information event codes	357
Configuration event codes	358
Appendix C. Object types	363
Appendix D. Valid combinations of FlashCopy and Metro Mirror functions	365
Appendix E. Moving data between MDisk groups with Copy Services	367
Accessibility	369
Notices	371
Trademarks.	372
Glossary	375
Index	385

I

Figures

1.	A SAN Volume Controller 2145-4F2 node	2
2.	A SAN Volume Controller 2145-8F2 node	2
3.	Example of a SAN Volume Controller in a fabric	4
4.	Levels of virtualization	8
5.	Asymmetrical virtualization	9
6.	Symmetrical virtualization	10
7.	Virtualization	12
8.	Configuration node	17
9.	I/O group and UPS	18
10.	2145 UPS-1U	20
11.	2145 UPS	20
12.	Controllers and MDisks	25
13.	MDisk group	27
14.	MDisk groups and VDIs	30
15.	Hosts, WWPNs, and VDIs	34
16.	Hosts, WWPNs, VDIs and SCSI mappings	34
17.	Disk controller system shared between SAN Volume Controller and a host	57
18.	ESS LUs accessed directly with a SAN Volume Controller	58
19.	DS4000 direct connection with a SAN Volume Controller on one host	59
20.	Fabric with ISL between nodes in a cluster	64
21.	Fabric with ISL in a redundant configuration	64
22.	Zoning a 1024 host configuration	71
23.	Basic frame layout	82
24.	Banner	82
25.	Task bar	83
26.	Software upgrade panel	207
27.	Software upgrade - file upload panel	208
28.	Node-rescue-request display	212
29.	PUTTY Configuration panel	326
30.	Updating Embedded WAS Ports panel	327

Tables

1.	Node state	16
2.	MDisk status	25
3.	MDisk group status	28
4.	Capacities of the cluster given extent size	29
5.	VDisk status	31
6.	VDisk cache modes	32
7.	FlashCopy mapping events	40
8.	Background copy	46
9.	Configuration terms and definitions	53
10.	Four hosts and their ports	67
11.	Six hosts and their ports	67
12.	Extent size	159
13.	Calculate the I/O rate	222
14.	Calculate the impact of FlashCopy relationships	223
15.	Determine if the storage subsystem is overloaded	224
16.	Controller device port selection algorithm	228
17.	Supported models of the EMC CLARiiON	242
18.	EMC CLARiiON global settings supported by the SAN Volume Controller	247
19.	EMC CLARiiON controller settings supported by the SAN Volume Controller	248
20.	EMC CLARiiON port settings supported by the SAN Volume Controller	248
21.	EMC CLARiiON LU settings supported by the SAN Volume Controller	249
22.	Supported models of the EMC Symmetrix	250
23.	EMC Symmetrix global settings supported by the SAN Volume Controller	256
24.	EMC Symmetrix port settings supported by the SAN Volume Controller	256
25.	EMC Symmetrix LU settings supported by the SAN Volume Controller	257
26.	Supported models of the IBM ESS	260
27.	Supported models of the IBM DS4000 series, StorageTek FlexLine series and StorageTek D series of controllers	268
28.	IBM DS4000 series controller global settings supported by the SAN Volume Controller	275
29.	Supported models of the HDS Lightning, Sun StorEdge and HPQ XP	277
30.	HDS Lightning global settings supported by the SAN Volume Controller	283
31.	HDS Lightning controller settings supported by the SAN Volume Controller	283
32.	HDS Lightning port settings supported by the SAN Volume Controller	284
33.	HDS Lightning LU settings for the SAN Volume Controller	284
34.	Supported Thunder 9200 models	285
35.	Supported Thunder 95xxV models	285
36.	Thunder global settings supported by the SAN Volume Controller	292
37.	Thunder port settings supported by the SAN Volume Controller	293
38.	Thunder LU settings for the SAN Volume Controller	294
39.	Determining partition usage	302
40.	Supported models of the HP StorageWorks HSG80	304
41.	HSG80 container types for logical unit configuration	308
42.	HSG80 global settings supported by the SAN Volume Controller	309
43.	HSG80 controller settings supported by the SAN Volume Controller	310
44.	HSG80 port settings supported by the	311
45.	HSG80 LU settings supported by the SAN Volume Controller	311
46.	HSG80 connection options and their required settings for the	312
47.	Supported HPQ EVA models	315
48.	HPQ EVA global settings supported by the SAN Volume Controller	318
49.	HPQ EVA LU settings supported by the SAN Volume Controller	319
50.	HPQ EVA host settings supported by the SAN Volume Controller	319
51.	Configuration commands	348
52.	Pool management commands	349

53.	Error messages for the IBM TotalStorage hardware provider	350
54.	Information event codes	357
55.	Configuration event codes	358
I 56.	Object types	363
57.	Valid combinations of FlashCopy and Metro Mirror interactions	365

About this guide

The IBM® TotalStorage® SAN Volume Controller Configuration guide provides information that helps you configure and use the IBM® TotalStorage® SAN Volume Controller™.

The IBM TotalStorage SAN Volume Controller 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.

Who should use this guide

The IBM TotalStorage SAN Volume Controller Configuration guide is intended for system administrators or others who install and use the SAN Volume Controller.

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

Summary of changes

This document contains terminology, maintenance, and editorial changes.

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

Summary of changes for SC26-7543-05 SAN Volume Controller Configuration Guide

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

New information

This topic describes the changes to this guide since the previous edition, SC26-7543-04. The following sections summarize the changes that have since been implemented from the previous version.

This version includes the following new information:

- Adding the following new SAN Volume Controller Console topics:
 - Creating Metro Mirror relationships
 - Filtering Metro Mirror relationships
 - Starting a Metro Mirror copy
 - Stopping a Metro Mirror copy
 - Modifying Metro Mirror relationships
 - Switching the copy direction of a Metro Mirror relationship
 - Deleting Metro Mirror relationships
 - Creating Metro Mirror consistency groups
 - Filtering Metro Mirror consistency groups
 - Renaming a Metro Mirror consistency group

- Starting a Metro Mirror consistency group copy
- Stopping a Metro Mirror consistency group copy
- Deleting Metro Mirror consistency groups
- Backing up the cluster configuration
- Downloading backup configuration data files
- Deleting backup configuration files
- Determining the relationship between MDisks and VDisks
- Discovering logical units
- Concurrent maintenance on the HPQ EVA
- Viewing the progress of a FlashCopy
- Viewing the progress of an MDisk removal
- Viewing the progress of VDisk formatting
- Viewing the progress of image mode migration
- Viewing the progress of extent migration
- Viewing the progress of VDisk migration
- Viewing the progress of Metro Mirror copy processes
- Modifying a host
- Adding ports to a host
- Deleting ports from a host
- Viewing mapped I/O groups
- Adding the following new topics:
 - Flushing data from the host volumes
 - Concurrent maintenance on the HPQ EVA
 - Changing the cluster gateway address using the CLI
 - Changing the cluster subnet mask using the CLI
 - Cluster IP failover
 - Accessing controller devices with multiple remote ports

Changed information

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

- The previous release referred to the uninterruptible power supply (UPS) as UPS 5115 and UPS 5125, by model number. For this release, the UPS is referred to by machine type. For example, this publication states 2145 uninterruptible power supply-1U (2145 UPS-1U) and uninterruptible power supply (2145 UPS). 2145 UPS-1U refers to UPS 5115 and 2145 UPS refers to UPS 5125.

Note: If text is referring to the UPS or to the uninterruptible power supply, then it is referring to a generic UPS and can be referring to either UPS. When the UPS is referred to as the 2145 UPS-1U or the 2145 UPS, then the specific UPS is designated.

- There is a new SAN Volume Controller supported model. The SAN Volume Controller is now documented by model number. For example, this publication states two SAN Volume Controller models types: SAN Volume Controller 2145-4F2 and the new SAN Volume Controller 2145-8F2.

Note: If text is referring to the SAN Volume Controller then it is referring to a generic SAN Volume Controller and can be referring to either SAN Volume Controller model. When the SAN Volume Controller is referred to as the

SAN Volume Controller 2145-4F2 or the SAN Volume Controller 2145-8F2, then the specific SAN Volume Controller is designated.

- The IBM TotalStorage FAStT series is now called the IBM TotalStorage DS4000 series
- Support for HDS Thunder 9520V
- Support for HP XP48
- Support for HP XP128
- Support for HP XP512
- Support for HP XP1024
- Support for Sun StorEdge 9910
- Support for Sun StorEdge 9960
- Support for Sun StorEdge 9970
- Support for IBM ESS 2105-E10
- Support for IBM ESS 2105-E20
- Support for IBM ESS 2105-F10
- Support for StorageTek D-series
- Support for StorageTek FlexLine series
- Support for 1024 host objects
- Support for cache disabled VDisks
- Support for Cisco MDS 9000 Interoperability modes
- Increased the number of supported SAN fabrics to 4
- Updated FlashCopy mapping states
- Updated zoning guidelines for SAN Volume Controller zones
- Updated zoning guidelines for host zones

Deleted information

This section lists information that was deleted from this document.

- The SAN Volume Controller no longer arrives with a CD set. All publication and product upgrades are available from the following Web site:
<http://www.ibm.com/servers/storage/support/virtual/2145.html>
- The maximum configuration table has been removed from this publication. See the following Web site for the latest supported maximum configurations:
<http://www.ibm.com/servers/storage/support/virtual/2145.html>

Summary of changes for SC26-7543-04 SAN Volume Controller Configuration Guide

This summary of changes provides a list of new, modified, and deleted information since the previous edition of this guide.

New information

This edition includes the following new information:

- The SAN Volume Controller is paired with a new uninterruptible power supply (UPS). The UPS 5115 is the preferred UPS for the SAN Volume Controller. Previous releases of the SAN Volume Controller were paired with the UPS 5125, which is still supported.

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

- Support for Microsoft® Volume Shadow Copy Service has been added.
- Instructions for copying a VDisk between managed disk (MDisk) groups that have different extent sizes using Copy Services have been added.
- Added the following new topics:
 - Extents
 - Input/Output governing
 - Metro Mirror
 - Limiting queue depth in large SANs
 - Queue depth
 - Calculating a queue depth limit
 - Homogeneous queue depth calculation
 - Non-homogeneous queue depth calculation
 - Limiting the queue depth
 - Creating an image mode virtual disk
 - Migration to image mode
 - Deleting an image mode virtual disk
 - Deleting a managed mode virtual disk
 - Creating an image mode VDisk using the CLI
 - Migrating to an image mode virtual disk using the CLI
 - Quorum disk creation and extent allocation
 - Access Logix
 - User interface on Symmetrix
 - User interface on ESS
 - TrueCopy
 - Write protect
 - Configuring settings for Lightning
 - Global settings for Lightning
 - Controller settings for Lightning
 - Port settings for Lightning
 - LU settings for Lightning
 - Zoning guidelines
 - IBM TotalStorage support for Microsoft Volume Shadow Copy service
 - Installation overview
 - System requirements for the IBM TotalStorage support for Microsoft Volume Shadow Copy service
 - Installing the IBM TotalStorage hardware provider
 - Verifying the installation
 - Creating the free and reserved pools of volumes
 - Configuration commands
 - Adding and removing VDIs
 - Error codes
 - Uninstalling the IBM TotalStorage hardware provider

Changed information

This section lists the updates that were made in this edition of the guide.

- Remote Copy is now known as IBM TotalStorage Metro Mirror for SAN Volume Controller.
- The maximum number of Metro Mirror relationships per cluster was increased to 4096.
- Information about image mode virtual disks has been enhanced.
- Information about managed disk (MDisk) extents has been enhanced.
- The SAN Volume Controller now supports connection to SANs that contain switches from different manufacturers.
- Instructions for calculating queue depth have been expanded.
- Complete information about the master console hardware and software have been moved to *IBM TotalStorage Master Console Installation and User's Guide*.

Deleted information

This section lists the information that was removed from this edition of the guide.

- Information about the SAN Volume Controller master console was removed. It now appears in the *IBM TotalStorage Master Console Installation and User's Guide*.

Emphasis

Different typefaces are used in this guide to show emphasis.

The following typefaces are used to show emphasis:

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

Numbering conventions

A specific numbering convention is used in this guide and product.

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

SAN Volume Controller library and related publications

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

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

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

SAN Volume Controller library

The following table lists and describes the publications that make up the SAN Volume Controller library. Unless otherwise noted, these publications are available in Adobe portable document format (PDF) from the following Web site:

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

Title	Description	Order number
<i>IBM TotalStorage SAN Volume Controller: CIM agent Developer's Reference</i>	This reference guide describes the objects and classes in a Common Information Model (CIM) environment.	SC26-7590
<i>IBM TotalStorage SAN Volume Controller: Command-Line Interface User's Guide</i>	This guide describes the commands that you can use from the SAN Volume Controller command-line interface (CLI).	SC26-7544
<i>IBM TotalStorage SAN Volume Controller: Configuration Guide</i>	This guide provides guidelines for configuring your SAN Volume Controller.	SC26-7543
<i>IBM TotalStorage SAN Volume Controller: Host Attachment Guide</i>	This guide provides guidelines for attaching the SAN Volume Controller to your host system.	SC26-7575
<i>IBM TotalStorage SAN Volume Controller: Installation Guide</i>	This guide includes the instructions the service representative uses to install the SAN Volume Controller.	SC26-7541
<i>IBM TotalStorage SAN Volume Controller: Planning Guide</i>	This guide introduces the SAN Volume Controller and lists the features you can order. It also provides guidelines for planning the installation and configuration of the SAN Volume Controller.	GA22-1052
<i>IBM TotalStorage SAN Volume Controller: Service Guide</i>	This guide includes the instructions the service representative uses to service the SAN Volume Controller.	SC26-7542

Title	Description	Order number
<i>IBM TotalStorage SAN Volume Controller: Translated Safety Notices</i>	This guide contains the danger and caution notices for the SAN Volume Controller. The notices are shown in English and in numerous other languages.	SC26-7577
<i>IBM TotalStorage Master Console Installation and User's Guide</i>	This guide includes the instructions on how to install and use the SAN Volume Controller Console	GC30-4090

Other IBM publications

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

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

Related Web sites

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

Type of information	Web site
SAN Volume Controller support	http://www.ibm.com/servers/storage/support/virtual/2145.html
Technical support for IBM storage products	http://www.ibm.com/storage/support/

How to order IBM publications

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

The IBM publications center

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

<http://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:

<http://www.ibm.com/shop/publications/order/>

How to send your comments

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

- e-mail

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

starpubs@us.ibm.com

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

- Mail

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

International Business Machines Corporation
RCF Processing Department
Department 61C
9032 South Rita Road
Tucson, Arizona 85775-4401
U.S.A.

Chapter 1. SAN Volume Controller overview

The *SAN Volume Controller* is a SAN (storage area network) appliance that attaches open-systems storage devices to supported open-systems hosts.

The SAN Volume Controller is a rack-mounted unit that you can install in a standard Electrical Industries Association (EIA) 19-inch rack. It provides symmetric virtualization by creating a pool of managed disks (MDisks) from the attached storage subsystems. Those storage systems are then mapped to a set of virtual disks (VDisks) for use by attached host systems. System administrators can view and access a common pool of storage on the SAN. This lets the administrators use storage resources more efficiently and provides a common base for advanced functions.

A SAN is a high-speed fibre-channel network that connects host systems and storage devices. It allows a host system to be connected to a storage device across the network. The connections are made through units such as routers, gateways, hubs, and switches. The area of the network that contains these units is known as the *fabric* of the network. For more information about SANs, see *IBM Storage Networking Virtualization: What's it all about?* and *IBM TotalStorage SAN Volume Controller: What is it and how to use it*.

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

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

Each SAN Volume Controller is a *node*; that is, it is either an end point of a link or a junction that is common to two or more links of the SAN. There are two types of SAN Volume Controller nodes: the SAN Volume Controller 2145-4F2 and the SAN Volume Controller 2145-8F2. Figure 1 on page 2 and Figure 2 on page 2 provide illustrations of the two types of SAN Volume Controller nodes. The nodes are always installed in pairs with one-to-four pairs of nodes constituting a *cluster*. Each node in a pair is configured to backup the other. Each pair of nodes is known as an I/O group.

All I/O operations that are managed by the nodes in an I/O group are cached on both nodes. Each virtual volume is defined to an I/O group. To avoid any single point of failure, the nodes of an I/O group are protected by independent uninterruptible power supplies (UPSs). There are two different UPSs. The UPSs are

called the 2145 uninterruptible power supply-1U (2145 UPS-1U) or 2145 uninterruptible power supply (2145 UPS) units.




Figure 1. A SAN Volume Controller 2145-4F2 node




Figure 2. A SAN Volume Controller 2145-8F2 node

A SAN Volume Controller I/O group takes the storage that is presented to the SAN by the back-end controllers as MDisks and translates that storage into logical disks, known as VDisks, that are used by applications on the hosts. Each node must reside in only one I/O group and provide access to the VDisks in that I/O group.

The SAN Volume Controller provides continuous operations and can also optimize the data path to ensure that performance levels are maintained. Ensure that you use the IBM TotalStorage Productivity Center for Disk and Replication tool to analyze the performance statistics. See the following IBM TotalStorage Productivity Center for Disk and Replication publications for more information:

- *IBM TotalStorage Productivity Center for Disk and Replication Version 2.1 Installation and Configuration Guide*
- *IBM TotalStorage Productivity Center for Disk and Replication Version 2.1 User's Guide*
- *IBM TotalStorage Productivity Center for Disk and Replication Version 2.1 Command-Line Interface User's Guide*

You can remove and replace any field replaceable unit (FRU) on one node while the other node of the pair continues to run. This allows the attached hosts to continue to access the attached storage while a node is repaired.

Related concepts

“Virtualization” on page 6

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

“Asymmetric virtualization” on page 8

With asymmetric virtualization, the virtualization engine is outside the data path and performs a metadata-style service. The metadata server contains all the mapping and the locking tables while the storage devices contain only data.

“Symmetric virtualization” on page 9

The SAN Volume Controller provides symmetric virtualization.

SAN fabric overview

The SAN fabric is an area of the network that contains routers, gateways, hubs, and switches. A single cluster SAN contains two distinct types of zones: a host zone and a disk zone.

In the host zone, the host systems can identify and address the SAN Volume Controller nodes. You can have more than one host zone. Generally, you create one host zone per operating system type. In the disk zone, the SAN Volume Controller nodes identify the disk drives. Host systems cannot operate on the disk drives directly; all data transfer occurs through the SAN Volume Controller nodes. Figure 3 on page 4 shows several host systems that are connected in a SAN fabric.




Figure 3. Example of a SAN Volume Controller in a fabric

A cluster of SAN Volume Controller nodes is connected to the same fabric and presents virtual disks (VDisks) to the host systems. You create these VDisks from units of space within a managed disk (MDisk) group. An MDisk group is a collection of MDisks that are presented by the back-end RAID controllers. The MDisk group provides a storage pool. You choose how each group is made up, and you can combine MDisks from different manufacturers' controllers in the same MDisk group.

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

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

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

Cluster configuration information is stored on every SAN Volume Controller node that is in the cluster to allow concurrent replacement of FRUs. When a new FRU is installed and when the SAN Volume Controller node is added back into the cluster,

configuration information that is required by that SAN Volume Controller node is read from other SAN Volume Controller nodes in the cluster.

SAN Volume Controller operating environment

You must set up your SAN Volume Controller operating environment using the supported multipathing software and hosts.

Minimum requirements

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

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

Note: Depending on how you ordered your SAN Volume Controller, the master console can be preconfigured on your platform or delivered as a software-only package.

Features of a SAN Volume Controller 2145-4F2 node

The following features are available with the SAN Volume Controller 2145-4F2 node:

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

Features of a SAN Volume Controller 2145-8F2 node

The following features are available with the SAN Volume Controller 2145-8F2 node:

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

Supported hosts

See the following Web site for a list of the supported operating systems:

<http://www.ibm.com/servers/storage/software/virtualization/svc>

Multipathing software

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

<http://www.ibm.com/servers/storage/software/virtualization/svc>

User interfaces

The SAN Volume Controller provides the following user interfaces:

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

Application programming interfaces

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

Virtualization

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

For data storage, virtualization includes the creation of a pool of storage that contains several disk subsystems. These subsystems can be supplied from various vendors. The pool can be split into virtual disks (VDisks) that are visible to the host systems that use them. Therefore, VDisks 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 the following four main levels:

At the server level

Manages 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.

At the storage device level

Uses striping, mirroring and RAID's to create 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.

At the fabric level

Enables storage pools to be independent of the servers and the physical components that make up the storage pools. One management interface can be used to manage different storage systems without affecting the servers. The SAN Volume Controller performs virtualization at the fabric level.

At the file system level

Provides the highest benefit because data is shared, allocated, and protected at the data level rather than the volume level.

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 provides a central point of control for disk creation and maintenance.

One problem area that virtualization addresses is unused capacity. Before virtualization, individual host systems each had their own storage, which wasted unused storage capacity. Using virtualization, storage is pooled so that jobs from any attached system that need large amounts of storage capacity can use it as needed. Virtualization makes it easier to regulate the amount of available storage without having to use host system resources or to turn storage devices off and on to add or remove capacity.

Types of virtualization

Virtualization can be performed either asymmetrically or symmetrically. Figure 4 on page 8 provides a diagram of the levels of virtualization.

Asymmetric

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

Symmetric

A virtualization engine sits in the data path and presents disks to the hosts, but hides 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, the benefits of those levels can also be combined. For example, you can gain the most benefits if you attach a low-cost RAID controller to a virtualization engine that provides virtual volumes for use by a virtual file system.

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




Figure 4. Levels of virtualization

Related concepts

“Asymmetric virtualization”

With asymmetric virtualization, the virtualization engine is outside the data path and performs a metadata-style service. The metadata server contains all the mapping and the locking tables while the storage devices contain only data.

“Symmetric virtualization” on page 9

The SAN Volume Controller provides symmetric virtualization.

Chapter 7, “Configuring and servicing storage subsystems,” on page 215

You must ensure that your disk controllers and switches are correctly configured to work with the SAN Volume Controller to avoid performance issues.

“VDisks” on page 30

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

Asymmetric virtualization

With asymmetric virtualization, the virtualization engine is outside the data path and performs a metadata-style service. The metadata server contains all the mapping and the locking tables while the storage devices contain only data.

In asymmetric virtual storage networks, the data flow, (2) in the Figure 5 on page 9, 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. However, there are disadvantages to asymmetric virtualization.

Asymmetric virtualization can have the following disadvantages:

- 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 5. Asymmetrical virtualization

Related concepts

“Symmetric virtualization”

The SAN Volume Controller provides symmetric virtualization.

Symmetric virtualization

The SAN Volume Controller provides symmetric virtualization.

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

Figure 6 on page 10 shows that the storage is pooled under the control of the virtualization engine, because the separation of the control from the data occurs in the data path. The virtualization engine performs the logical-to-physical mapping.




Figure 6. 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. Therefore, the virtualization engine is a central point of control for device and advanced function storage management. Symmetric virtualization allows you to build a firewall in the storage network. Only the virtualization engine can grant access through the firewall.

Symmetric virtualization can cause some problems. The main problem that is associated with symmetric virtualization is scalability. Scalability can cause poor performance because all input/output (I/O) must flow through the virtualization engine. To solve this problem, you can use an *n-way* cluster of virtualization engines that has failover capacity. You can scale the additional processor power, cache memory, and adapter bandwidth to achieve the level of performance that you want. Additional memory and processing power are needed to run advanced services such as Copy Services and caching.

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

Related concepts

“Asymmetric virtualization” on page 8

With asymmetric virtualization, the virtualization engine is outside the data path and performs a metadata-style service. The metadata server contains all the mapping and the locking tables while the storage devices contain only data.

“Virtualization” on page 6

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

Physical links between SAN Volume Controller nodes and a switch

Between the SAN Volume Controller nodes and the switch to which they are connected, the SAN Volume Controller supports shortwave small form factor pluggable (SFP) transceivers (850 nm with 50 μm or 62.5 μm multimode cables).

The transceivers can run at up to 500 m, limited by the pulse spreading that is caused by the multimode nature of the transmission. *Multimode* means that each ray of light can take slightly different paths from one part of the cable to another.

The paths are different lengths, which means a very short pulse of light put in at one end can spread out when it reaches the far end. Not all the light that is associated with that pulse travels exactly the same distance to get out of the fibre, and the longer the fibre, the more the pulse is dispersed. For longer distances between the nodes and the local switch, the SAN Volume Controller supports longwave SFP transceivers (1300 nm with 9 µm unimode2 cables). These transceivers are supported up to 10 km.

Support for long links between the local and remote fabric

Ensure that you are familiar with the support for the inter-switch link (ISL) between the local and remote fabric.

See the following IBM Web site for the supported ISLs:

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

Object overview

The SAN Volume Controller is based on a number of virtualization concepts.

A SAN Volume Controller consists of a single node. Nodes are deployed in pairs to make up a cluster. A cluster can have one to four node pairs in it. Each pair of nodes is known as an input/output (I/O) group. Each node must be in only one I/O group.

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

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

The nodes in a cluster recognize the storage that is presented by SAN-attached storage subsystems as a number of disks, known as *managed disks (MDisks)*. Because the SAN Volume Controller does not attempt to provide recovery from physical disk failures within the back-end disk controllers, an MDisk is usually, but not necessarily, a RAID array.

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

MDisks are collected into groups, known as MDisk groups. VDisks are created from the extents that are contained by an MDisk group. The MDisks that constitute a particular VDisk must all come from the same MDisk 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 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 allows you to create logical host objects that group together WWPNs belonging to a single application server or multiple application servers.

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

Figure 7 provides an overview of the virtualization concepts that are defined in this section.

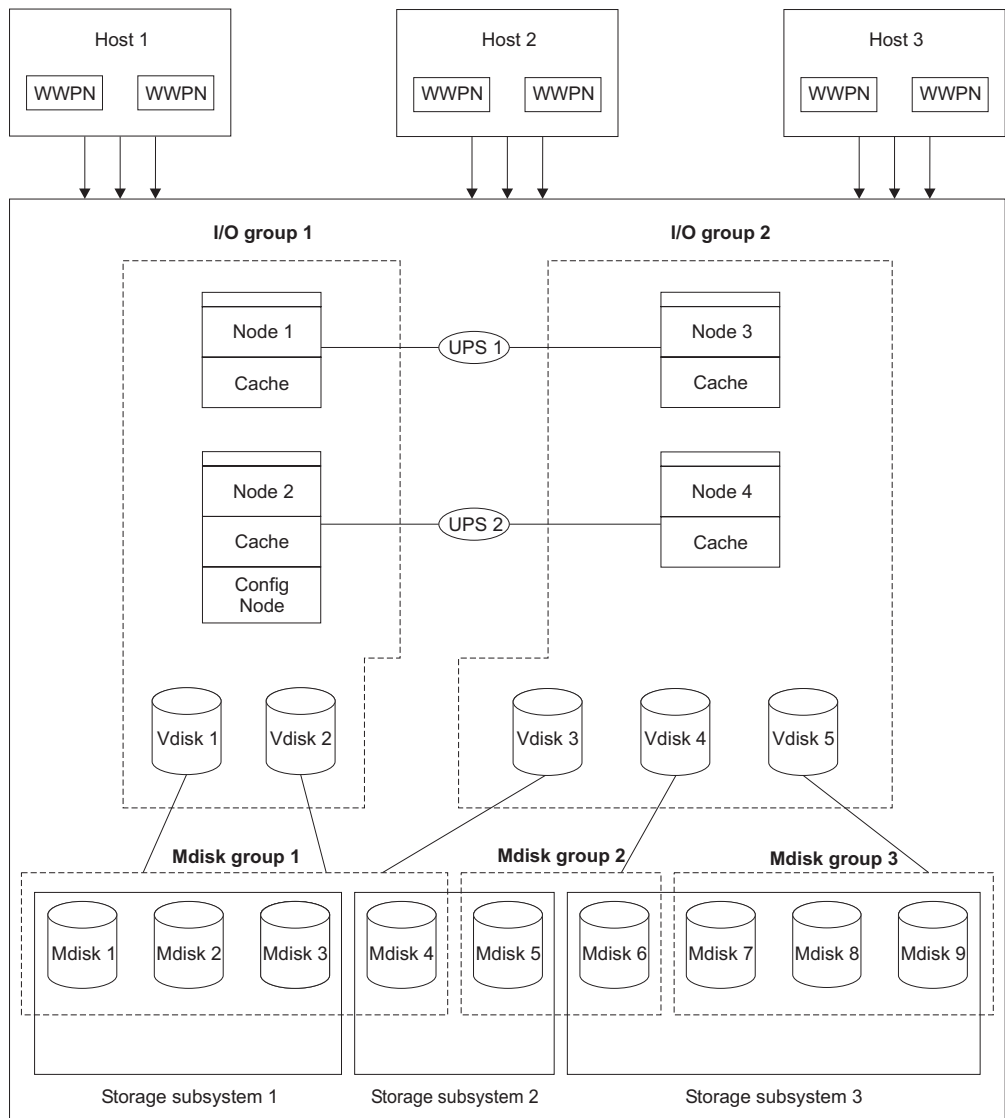


Figure 7. Virtualization

Related concepts

“Nodes and clusters”

A SAN Volume Controller node is a single processing unit, which provides virtualization, cache, and copy services for the SAN.

“I/O groups and uninterruptible power supply” on page 17

Nodes are deployed in pairs to make up a cluster. Each pair of nodes is known as an *I/O group*. Each node can only be in one I/O group.

“Storage subsystems and MDisks” on page 22

The nodes in a cluster see the storage exported by SAN-attached storage subsystems as a number of disks, known as managed disks (MDisks). The SAN Volume Controller does not attempt to provide recovery from physical disk failures within the storage subsystem. An MDisk is usually, but not necessarily, a RAID array.

“MDisk groups and VDIs” on page 26

Managed disks (MDisks) are collected into groups known as *managed disk groups*. Virtual disks (VDIs) are logical disks that are presented to the SAN by SAN Volume Controller nodes. The maximum number of supported VDIs per I/O group is 1024. The maximum number of supported VDIs per cluster is 4096. VDIs, like nodes, are associated with an I/O group.

“VDI-to-host mapping” on page 33

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

Nodes and clusters

A SAN Volume Controller node is a single processing unit, which provides virtualization, cache, and copy services for the SAN.

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 concepts

“Clusters”

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

Related reference

“Nodes” on page 16

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

Clusters

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

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

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

Cluster configuration backup:

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

Backing up the cluster configuration enables you to restore your cluster configuration in the event that it is lost. Only the data that describes the cluster configuration is backed up. You must backup your application data using the appropriate backup methods.

Objects included in the backup

Configuration data is information about a cluster and the objects that are defined in it. Information about the following objects is included in the cluster configuration data:

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

Related concepts

“Clusters” on page 13

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

“Configuration restore”

Configuration restore is the process of using a backup cluster configuration data file or files to restore a specific cluster configuration.

Configuration restore:

Configuration restore is the process of using a backup cluster configuration data file or files to restore a specific cluster configuration.

Restoring the cluster configuration is an important part of a complete backup and disaster recovery solution. You must also regularly back up your application data using appropriate backup methods because you must restore your application data after you have restored your cluster configuration.

The process to restore your cluster configuration consists of two phases:

- Preparing
- Executing

Before issuing the preparation command, the new cluster must be reset to a default state. During the preparation phase, the backup cluster configuration data and the new cluster are analyzed for compatibility and a sequence of commands are prepared to be run.

During the executing phase, the command sequence is run.

Related concepts

“Clusters” on page 13

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

Chapter 5, “Backing up and restoring the cluster configuration,” on page 195

You can back up and restore the cluster configuration data after preliminary tasks are completed.

“Cluster configuration backup” on page 13

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

Cluster IP failover:

If the configuration node fails, the cluster IP address is transferred to a new node. The cluster services are used to manage the IP address transfer from the failed configuration node to the new configuration node.

The following changes are performed by the cluster service:

- If software on the failed configuration node is still operational, the software shuts down the IP interface. If the software cannot shut down the IP interface, the hardware service forces a shut down.
- When the IP interface shuts down, all remaining nodes choose a new node to host the configuration interface.
- The new configuration node initializes the configuration daemons, sshd and httpd, and then binds the configuration IP interface to its Ethernet port.
- The router is configured as the default gateway for the new configuration node.
- The new configuration node sends five unsolicited address resolution protocol (ARP) packets to the local subnet broadcast address. The ARP packets contain the cluster IP and the media access control (MAC) address for the new configuration node. All systems that receive ARP packets are forced to update their ARP tables. Once the ARP tables are updated, these systems can connect to the new configuration node.

Note: Some Ethernet devices might not forward ARP packets. If the ARP packets are not forwarded, connectivity to the new configuration node cannot be established automatically. To avoid this problem, configure all Ethernet devices to pass unsolicited ARP packets. You can restore lost connectivity by logging into the SAN Volume Controller and starting a secure copy to the affected system. Starting a secure copy forces an update to the ARP cache for all systems connected to the same switch as the affected system.

Ethernet link failures

If the Ethernet link to the SAN Volume Controller cluster fails because of an event unrelated to the SAN Volume Controller itself, such as a cable being disconnected or an Ethernet router failure, the SAN Volume Controller does not attempt to failover

the configuration node to restore IP access to the cluster.

Nodes

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

For redundancy, nodes are deployed in pairs to make up a cluster. A cluster can have one to four pairs of nodes. Each pair of nodes is known as an I/O group. Each node can be in *only* one I/O group. A SAN Volume Controller 2145-4F2 node and a SAN Volume Controller 2145-8F2 node cannot be in the same I/O group. A maximum of four I/O groups each containing two nodes is supported.

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

Table 1 describes the operational states of a node.

Table 1. Node state

State	Description
Adding	The node was added to the cluster but is not yet synchronized with the cluster state (see Note). The node state changes to Online after synchronization is complete.
Deleting	The node is in the process of being deleted from the cluster.
Online	The node is operational, assigned to a cluster, and has access to the fibre-channel SAN fabric.
Offline	The node is not operational. The node was assigned to a cluster but is not available on the fibre-channel SAN fabric. Run the Directed Maintenance Procedures to determine the problem.
Pending	The node is transitioning between states and, in a few seconds, will move to one of the other states.
Note: A node can stay in the Adding state for a long time. You should wait for at least 30 minutes before taking further action, but if after 30 minutes, the node state is still Adding, you can delete the node and add it again. If the node that has been added is at a lower code level than the rest of the cluster, the node is upgraded to the cluster code level, which can take up to 20 minutes. During this time, the node is shown as adding.	

Related concepts

“Configuration node”

A *configuration node* is a single node that manages configuration activity of the cluster.

Configuration node:

A *configuration node* is a single node that manages configuration activity of the cluster.

The configuration node is the main source for configuration commands. The configuration node 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 SAN Volume Controller Console.

Figure 8 shows an example cluster containing four nodes. Node 1 has been designated the configuration node. User requests (1) are targeted at Node 1. This can cause requests that are targeted at the other nodes in the cluster to have their data returned to Node 1.




Figure 8. Configuration node

Related reference

“Nodes” on page 16

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

I/O groups and uninterruptible power supply

Nodes are deployed in pairs to make up a cluster. Each pair of nodes is known as an *I/O group*. Each node can only be in one *I/O group*.

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

I/O groups

An *I/O group* is a group that is defined during the cluster configuration process.

| Each node can only be in one I/O group. A SAN Volume Controller 2145-4F2 node
 | and a SAN Volume Controller 2145-8F2 node cannot be in the same I/O group. The
 | I/O groups are connected to the SAN so that all back-end storage and all
 | application servers are visible to all of the I/O groups. Each pair of nodes has the
 | responsibility to serve I/O operations on a particular virtual disk (VDisk).

VDisks are logical disks that are presented to the SAN by SAN Volume Controller nodes. VDisks are also associated with an *I/O group*. The SAN Volume Controller does not contain any internal battery backup units and therefore must be connected to an uninterruptible power supply (UPS) to provide data integrity in the event of a cluster-wide power failure. The UPS only provides power long enough to enable the SAN Volume Controller cluster to shutdown and save cache data. The UPS is not intended to maintain power and keep the SAN Volume Controller running during an outage.

When an application server performs I/O to a VDisk, it has the choice of accessing the VDisk using either of the nodes in the I/O group. When the VDisk is created, you can specify a preferred node. After a preferred node is specified, you should only access the VDisk through the preferred node. Because each I/O group only has two nodes, the distributed cache in the SAN Volume Controller is 2-way. When I/O is performed to a VDisk, the node that processes the I/O duplicates the data onto the partner node that is in the I/O group.

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

Figure 9 shows a write operation from a host (1), that is targeted for VDisk A. This write is targeted at the preferred node, Node 1 (2). The write is cached and a copy of the data is made in the partner node, Node 2's cache (3). The host views the write as complete. At some later time, the data is written, or de-staged, to storage (4). Figure 9 also shows two uninterruptible power supplies (UPSs) correctly configured so that each node is in a different power domain.




Figure 9. I/O group and UPS

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

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

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

The cluster also provides a recovery I/O group, which is used when both nodes in the I/O group have experienced multiple failures. You can move the VDIsks to the recovery I/O group and then into a working I/O group. I/O access is not possible when VDIsks are assigned to the recovery I/O group.

Related concepts

“I/O groups” on page 17

An *I/O group* is a group that is defined during the cluster configuration process.

“UPS”

The uninterruptible power supply (UPS) provides the SAN Volume Controller with a secondary power source if you lose power from your primary power source due to power failures, power sags, power surges, or line noise. The SAN Volume Controller 2145-8F2 supports the 2145 uninterruptible power supply-1U (2145 UPS-1U), but does not support the 2145 uninterruptible power supply (2145 UPS). The SAN Volume Controller 2145-4F2 supports both the 2145 UPS-1U and the 2145 UPS.

“I/O groups and uninterruptible power supply” on page 17

Nodes are deployed in pairs to make up a cluster. Each pair of nodes is known as an *I/O group*. Each node can only be in one I/O group.

I/O governing

You can set the maximum amount of I/O activity that a host sends to a virtual disk (VDisk). This amount is known as the *I/O governing rate*. The governing rate can be expressed in I/Os per second or MB per second.

Read, write, and verify commands that access the physical medium are subject to I/O governing.

I/O governing does not effect FlashCopy and data migration I/O rates.

Governing is applied to Metro Mirror primary and secondary VDIsks as follows:

- If an I/O governing rate is set on a secondary VDisk, then the same governing rate is applied to the primary.
- If you set an I/O governing on the primary and the secondary VDisk, then the I/O governing rate for the pair is the lowest rate that is set.


UPS

The uninterruptible power supply (UPS) provides the SAN Volume Controller with a secondary power source if you lose power from your primary power source due to power failures, power sags, power surges, or line noise. The SAN Volume Controller 2145-8F2 supports the 2145 uninterruptible power supply-1U (2145 UPS-1U), but does not support the 2145 uninterruptible power supply (2145 UPS). The SAN Volume Controller 2145-4F2 supports both the 2145 UPS-1U and the 2145 UPS.

Important: The SAN Volume Controller 2145-8F2 can only operate with the 2145 UPS-1U.

Unlike the traditional UPS that enables continued operation of the devices that they supply when power is lost, these UPS units are used exclusively to maintain data that is held in the SAN Volume Controller dynamic random access memory (DRAM)

in the event of an unexpected loss of external power. Data is saved to the SAN Volume Controller internal disk. The UPS units are required to power the SAN Volume Controllers even if the input power source is itself uninterruptible. Figure 11 and Figure 10 provide illustrations of the two types of UPS units.



svc00001

Figure 10. 2145 UPS-1U




Figure 11. 2145 UPS

Note: The UPS maintains continuous SAN Volume Controller-specific communications with its attached SAN Volume Controller nodes. The SAN Volume Controller cannot operate without the UPS. The SAN Volume Controller UPS must be used in accordance with documented guidelines and procedures and must not power any equipment other than SAN Volume Controller nodes.

Related concepts

“I/O groups” on page 17

An *I/O group* is a group that is defined during the cluster configuration process.

“I/O groups and uninterruptible power supply” on page 17

Nodes are deployed in pairs to make up a cluster. Each pair of nodes is known as an *I/O group*. Each node can only be in one *I/O group*.

UPS configuration:

To provide full redundancy and concurrent maintenance, the SAN Volume Controllers must be installed in pairs.

Each SAN Volume Controller of a pair must be connected to a different uninterruptible power supply (UPS). Each 2145 UPS can support up to two SAN Volume Controller 2145-4F2 nodes. The 2145 UPS-1U can only support one SAN Volume Controller 2145-8F2 or SAN Volume Controller 2145-4F2 node. You can connect the two UPS units for the pair to different independent electrical power sources. This reduces the chance of an input power failure at both UPS units.

The UPS must be in the same rack as the nodes.

The following table provides the UPS guidelines for the SAN Volume Controller 2145-4F2:

Number of SAN Volume Controller 2145-4F2 models	Number of 2145 UPS units	Number of 2145 UPS-1U units
2	2	2
4	2	4
6	4	6
8	4	8

The following table provides the UPS guidelines for the SAN Volume Controller 2145-8F2:

Number of SAN Volume Controller 2145-8F2 models	Number of 2145 UPS units	Number of 2145 UPS-1U units
2	Not supported	2
4	Not supported	4
6	Not supported	6
8	Not supported	8

Attention:

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

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

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

UPS operation:

Each SAN Volume Controller monitors the operational state of the uninterruptible power supply (UPS) to which it is attached.

If the UPS reports a loss of input power, the SAN Volume Controller stops all I/O operations and dumps the contents of its dynamic random access memory (DRAM) to the internal disk drive. When input power to the UPS is restored, the SAN Volume Controllers restart and restore the original contents of the DRAM from the data saved on the disk drive.

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

Note: If input power is disconnected from the UPS, a fully-operational SAN Volume Controller that is connected to that UPS performs a power-down sequence. This operation, which saves the configuration and cache data to an internal disk in the SAN Volume Controller, typically takes about three minutes, at which time power is removed from the output of the UPS. In the event of a delay in the completion of the power-down sequence, the UPS output power is removed five minutes after the power is disconnected from the UPS. Because this operation is controlled by the SAN Volume Controller, a UPS that is not connected to an active SAN Volume Controller does not shut off within the five-minute required period.

Attention: Data integrity can be compromised by pushing the 2145 UPS power-off button or the 2145 UPS-1U on/off button. However, in the case of an emergency, you can manually shut down the UPS by pushing the 2145 UPS power-off button or the 2145 UPS-1U on/off button. Never shut down a UPS without first shutting down the SAN Volume Controller nodes that it supports.

Important: The two SAN Volume Controller 2145-4F2 nodes in the I/O group must be connected to different 2145 UPSs. This configuration ensures that cache and cluster state information is protected in the event of a failure of the UPS or mainline power source.

When nodes are added to the cluster, you must specify the I/O group that they are joining. The configuration interfaces check the UPS units and ensure that the two nodes in the I/O group are not connected to the same UPS units.

Storage subsystems and MDisks

The nodes in a cluster see the storage exported by SAN-attached storage subsystems as a number of disks, known as managed disks (MDisks). The SAN

Volume Controller does not attempt to provide recovery from physical disk failures within the storage subsystem. An MDisk is usually, but not necessarily, a RAID array.

Related concepts

“Storage subsystems”

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

“MDisks” on page 24

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

Storage subsystems

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

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

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).

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

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

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

Many storage subsystems allow the storage that is provided by a RAID to be divided up into many SCSI logical units (LUs) that are presented on the SAN. With the SAN Volume Controller, ensure that the storage subsystems are configured to present each RAID as a single SCSI LU that are recognized by the SAN Volume

Controller as a single MDisk. The virtualization features of the SAN Volume Controller can then be used to divide up the storage into VDIs.

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

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

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

Note: The SAN Volume Controller Console supports storage that is internally configured as a RAID. 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 VDIs that were created from the MDisk group.

Related concepts

“MDisks”

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

“Storage subsystems and MDisks” on page 22

The nodes in a cluster see the storage exported by SAN-attached storage subsystems as a number of disks, known as managed disks (MDisks). The SAN Volume Controller does not attempt to provide recovery from physical disk failures within the storage subsystem. An MDisk is usually, but not necessarily, a RAID array.

MDisks

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

An MDisk might, therefore, consist of multiple physical disks that are presented as a single logical disk to the SAN. An MDisk 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 MDisk is divided into a number of extents, which are numbered, from 0, sequentially from the start to the end of the MDisk. The extent size is a property of MDisk groups. When an MDisk is added to an MDisk group, the size of the extents that the MDisk is divided into depends on the attribute of the MDisk group to which it has been added.

Access modes

The access mode determines how the cluster uses the MDisk. The following list provides the three types of possible access modes:

Unmanaged

The MDisk is not used by the cluster.

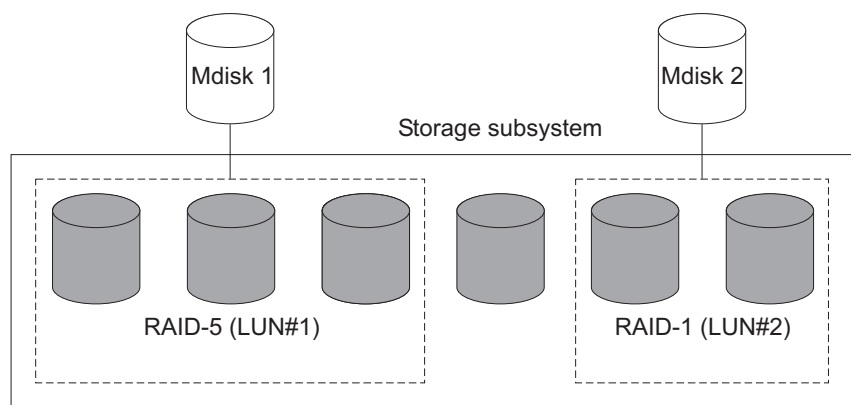
Managed

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

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

Attention: If you add an MDisk that contains existing data to an MDisk group while the MDisk is in unmanaged or managed mode, you lose the data that it contains. The *image mode* is the only mode that preserves this data.

Figure 12 shows physical disks and MDisks.





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

Figure 12. Controllers and MDisks

Table 2 describes the operational states of an MDisk.

Table 2. MDisk status

Status	Description
Online	The MDisk can be accessed by all online nodes. That is, all the nodes that are currently working members of the cluster can access this MDisk. The MDisk is online when the following conditions are met: <ul style="list-style-type: none">• All timeout error recovery procedures complete and report the disk as online.• Logical unit number (LUN) inventory of the target ports correctly reported the MDisk.• Discovery of this LUN completed successfully.• All of the MDisk 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 can be partially excluded; that is, some of the paths to the MDisk (but not all) have been excluded.

Table 2. MDisk status (continued)

Excluded	The MDisk has been excluded from use by the cluster after repeated access errors. Run the Directed Maintenance Procedures to determine the problem.
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 is 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 that provides the logical connection between MDisks and VDIs.

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

```
svctask includemdisk mdiskname/id
```

Where *mdiskname/id* is the name or ID of your MDisk.

MDisk path

Each MDisk has an online path count, which is the number of nodes that have access to that MDisk; this represents a summary of the I/O path status between the cluster nodes and the 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. If the current path count is not equal to the maximum path count, the MDisk might be degraded. That is, one or more nodes might not see the MDisk on the fabric.

Related concepts

“Storage subsystems” on page 23

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

“Storage subsystems and MDisks” on page 22

The nodes in a cluster see the storage exported by SAN-attached storage subsystems as a number of disks, known as managed disks (MDisks). The SAN Volume Controller does not attempt to provide recovery from physical disk failures within the storage subsystem. An MDisk is usually, but not necessarily, a RAID array.

MDisk groups and VDIs

Managed disks (MDisks) are collected into groups known as *managed disk groups*. Virtual disks (VDIs) are logical disks that are presented to the SAN by SAN Volume Controller nodes. The maximum number of supported VDIs per I/O group

is 1024. The maximum number of supported VDisks per cluster is 4096. VDisks, like nodes, are associated with an I/O group.

VDisks are created from the extents of MDisks. Only MDisks that are in the same MDisk group can contribute extents to a VDisk.

Related concepts

“MDisk groups”

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

“VDisks” on page 30

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

MDisk groups

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

Figure 13 shows an MDisk group containing four MDisks.




Figure 13. MDisk group

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

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

You can add 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:

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

Attention: If you delete an MDisk group, you destroy all the VDIs that are made from the extents that are in the group. If the group is deleted, you cannot recover the mapping that existed between extents that are in the group or the extents that the VDIs use. The MDIs 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 3 describes the operational states of an MDisk group.

Table 3. MDisk group status

Status	Description
Online	The MDisk group is online and available. All the MDIs in the group are available.
Degraded	The MDisk group is available; however, one or more nodes cannot access all the MDIs in the group.
Offline	The MDisk group is offline and unavailable. No nodes in the cluster can access the MDIs. The most likely cause is that one or more MDIs 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. Take care when you create MDisk groups to ensure an optimal configuration.

Consider the following guidelines when you create MDisk groups:

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

Extents

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

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

Ensure that your MDisk groups *do not* have different extent sizes. Different extent sizes place restrictions on the use of data migration. You cannot use the SAN Volume Controller data migration function to move a VDisk between MDisk groups that have different extent sizes.

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

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

The choice of extent size affects the total amount of storage that is managed by the cluster. Table 4 shows the maximum amount of storage that can be managed by a cluster for each extent size.

Table 4. Capacities of the cluster given extent size

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

A cluster can manage 4 million extents (4 x 1024 x 1024). For example, with a 16 MB extent size, the cluster can manage up to 16 MB x 4 MB = 64 TB of storage.

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

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

Related concepts

“VDisks” on page 30

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

“MDisk groups and VDIsks” on page 26

Managed disks (MDisks) are collected into groups known as *managed disk groups*. Virtual disks (VDIsks) are logical disks that are presented to the SAN by SAN Volume Controller nodes. The maximum number of supported VDIsks per I/O group is 1024. The maximum number of supported VDIsks per cluster is 4096. VDIsks, like nodes, are associated with an I/O group.

VDisks

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

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

Types

You can create the following types of VDisks:

Striped

A VDisk that has been striped is at the extent level. One extent is allocated, in turn, from each MDisk that is in the group. For example, an MDisk group that has 10 MDisks takes one extent from each MDisk. The 11th extent is taken from the first MDisk, 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 MDisk group. The round-robin procedure is used across the specified stripe set.

Attention: By default, striped VDisks are striped across all MDisks in the group. If some of the MDisks are smaller than others, the extents on the smaller MDisks are 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 if there is sufficient free space to create a striped VDisk, select one of the following options:

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

Figure 14 shows an example of an MDisk group that contains three MDisks. This figure also shows a striped VDisk that is created from the extents that are available in the group.

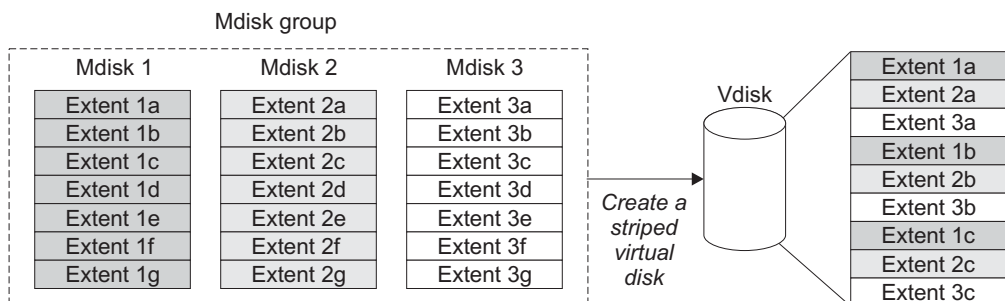


Figure 14. MDisk groups and VDisks

Sequential

When extents are selected, they are allocated sequentially on one MDisk to create the VDisk if enough consecutive free extents are available on the chosen MDisk.

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

When you create an image-mode VDisk, you must assign it to an MDisk group. An image-mode VDisk must be at least one extent in size. 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 VDisk becomes a virtualized disk, and the mode of the MDisk changes from image to managed.

Attention: If you add a managed mode MDisk to an MDisk group, any data on the MDisk is 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 if it contains partitions or data.

A VDisk can be in one of three states: online, offline, and degraded. Table 5 describes the different states of a VDisk.

Table 5. VDisk status

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

You can use more sophisticated extent allocation policies to create VDIs. When you create a striped VDisk, you can specify the same MDisk more than once in the list of MDisks that are used as the stripe set. This is useful if you have an MDisk group in which not all the MDisks are of the same capacity. For example, if you have an MDisk group that has two 18 GB MDisks and two 36 GB MDisks, you can create a striped VDisk 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 VDisk, you destroy access to the data that is on the VDisk. The extents that were used in the VDisk are returned to the pool of free extents that is in the MDisk group. The deletion might fail if the VDisk is still mapped to hosts. The

deletion might also fail if the VDisk is still part of a FlashCopy or a Metro Mirror mapping. If the deletion fails, you can specify the force-delete flag to delete both the VDisk and the associated mappings to hosts. Forcing the deletion deletes the Copy Services relationship and mappings.

Cache modes

You can select if you want read and write operations to be stored in cache by specifying a cache mode. You must specify the cache mode when you create the VDisk. After the VDisk is created, you cannot change the cache mode.

Table 6 describes the two types of cache modes for a VDisk.

Table 6. VDisk cache modes

Cache mode	Description
readwrite	All read and write I/O operations that are performed by the VDisk are stored in cache. This is the default cache mode for all VDIs.
none	All read and write I/O operations that are performed by the VDisk are not stored in cache.

Related concepts

“MDisk groups” on page 27

A *managed disk (MDisk) group* is a collection of MDisks that jointly contain all the data for a specified set of virtual disks (VDIs).

“MDisk groups and VDIs” on page 26

Managed disks (MDisks) are collected into groups known as *managed disk groups*. Virtual disks (VDIs) are logical disks that are presented to the SAN by SAN Volume Controller nodes. The maximum number of supported VDIs per I/O group is 1024. The maximum number of supported VDIs per cluster is 4096. VDIs, like nodes, are associated with an I/O group.

Host objects

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

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 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 virtual disks (VDIs) on the fibre. You must map each VDisk to a particular set of ports to enable the VDisk to be accessed through those ports. The mapping is made between a host object and a VDisk.

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 type the port names under this condition.

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

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

Node login counts

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

Related concepts

“VDisk-to-host mapping”

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

VDisk-to-host mapping

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

VDisk-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. VDisk-to-host mapping is done at the SAN Volume Controller level.

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

VDisks and host mappings

LUN masking usually requires device driver software on each host. The device driver software masks the LUNs. After the masking is complete, only some disks are visible to the operating system. The SAN Volume Controller performs a similar function, but, by default, it presents to the host only those VDisks that are mapped to that host. Therefore, you must map the VDisks to the hosts that you want to access those disks.

Each host mapping associates a VDisk with a host object and allows all HBA ports in the host object to access the VDisk. You can map a VDisk to multiple host objects. When a mapping is created, multiple paths might exist across the SAN fabric from the hosts to the SAN Volume Controllers that are presenting the VDisk. Most operating systems present each path to a VDisk as a separate storage device. The SAN Volume Controller, therefore, requires that multipathing software be running on the host. The multipathing software manages the many paths that are available to the VDisk and presents a single storage device to the operating system.

When you map a VDisk to a host, you can optionally specify a SCSI ID for the VDisk. This ID controls the sequence in which the VDIs are presented to the host. For example, if you present three VDIs to the host, and those VDIs have SCSI IDs of 0, 1, and 3, the VDisk that has an ID of 3 might not be found because no disk is mapped with an ID of 2. The cluster automatically assigns the next available SCSI ID if none is entered.

Figure 15 and Figure 16 show two VDIs, and the mappings that exist between the host objects and these VDIs.

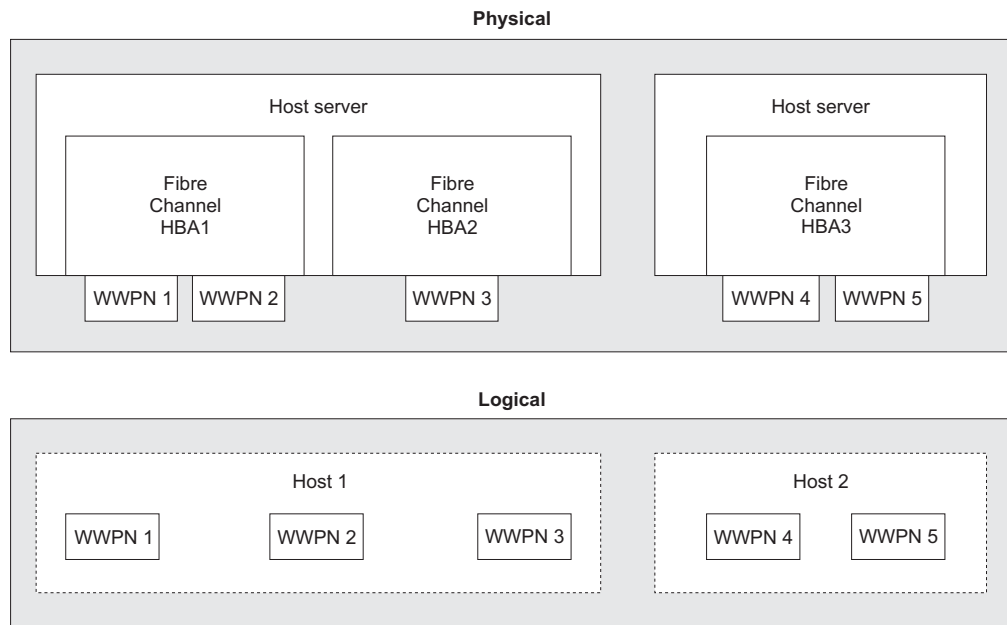


Figure 15. Hosts, WWPNs, and VDIs

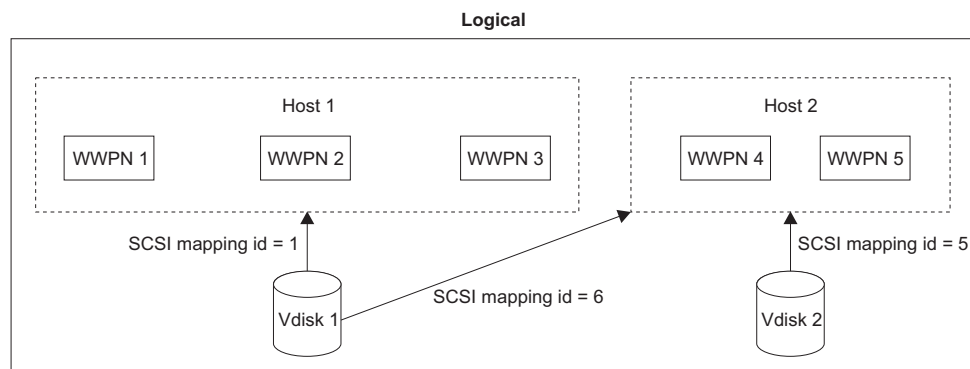


Figure 16. Hosts, WWPNs, VDIs and SCSI mappings

Related concepts

“Host objects” on page 32

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

“VDisk-to-host mapping” on page 33

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

Copy Services

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

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

FlashCopy

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

Metro Mirror

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

Related concepts

“FlashCopy”

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

“Metro Mirror” on page 47

Metro Mirror enables you to set up a relationship between two virtual disks (VDisks), so that updates that are made by an application to one VDisk are mirrored on the other VDisk.

FlashCopy

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

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

It is difficult to make a consistent copy of a data set that is constantly updated. Point-in-time copy techniques help solve this problem. If a copy of a data set is created using a technology that does not provide point-in-time techniques and the data set changes during the copy operation, the resulting copy might contain data 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 copied, the copy contains the referenced object at its new location but the copied reference still points to the old location.

Source VDisks and target VDisks must meet the following requirements:

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

Related concepts

“FlashCopy mappings” on page 38

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

“FlashCopy consistency groups” on page 42

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

Related reference

“FlashCopy applications”

You can use FlashCopy to include taking consistent backups of changing data, testing of applications, and creating copies for auditing purposes and for data mining.

“FlashCopy indirection layer” on page 44

FlashCopy provides the semantics of a point in time copy by using an indirection layer which intercepts I/Os targeted at both the source and target virtual disks.

“Background copy” on page 46

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

Related information

“Host considerations for FlashCopy integrity”

The SAN Volume Controller FlashCopy function transfers a point-in-time copy of a source virtual disk (VDisk) onto a designated target VDisk.

FlashCopy applications

You can use FlashCopy to include taking consistent backups of changing data, testing of applications, and creating copies for auditing purposes and for data mining.

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 become redundant and can now be discarded. Usually in this backup condition, the target data can be managed as read only.

Another use of FlashCopy data is the testing of applications. 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.

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 might be possible to restart the job from a saved copy of its data rather than rerunning the entire multi-day job.

Related concepts

“FlashCopy” on page 35

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

Host considerations for FlashCopy integrity

The SAN Volume Controller FlashCopy function transfers a point-in-time copy of a source virtual disk (VDisk) onto a designated target VDisk.

You must create or already have an existing target VDisk before you can transfer the copy. You must also ensure that the target VDisk has enough space to support the amount of data that is being transferred.

All of the data on the source VDisk is copied to the target VDisk. Therefore, operating system control information, application data and meta-data are included in the data that is copied to the target VDisk. Because all of the data is copied, some operating systems do not allow a source VDisk and a target VDisk to reside on the same host. In order to ensure the integrity of the copy that is made, it is necessary to completely flush the host cache of any outstanding reads or writes before you proceed with the FlashCopy. Host cache flushing is ensured by unmounting the source VDIsks from the source host before you start the FlashCopy.

Because the target VDIsks are overwritten with a complete image of the source VDIsks, it is important that any data held on the host operating system (or application) caches for the target VDIsks 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 VDIsks prior to starting the 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, they can be used to prepare and start a FlashCopy. Refer to your host and application documentation for details.

Some operating systems are unable to use a copy of a VDisk without *synthesis*. Synthesis performs a transformation of the operating system meta-data on the target VDisk to allow the operating system to use the disk. Refer to your host documentation on how to detect and mount the copied VDIsks.

Related concepts

“FlashCopy” on page 35

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

Flushing data from the host volumes:

All outstanding read and write operations must be flushed from the host cache before you use FlashCopy.

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

1. If you are using UNIX[®] or Linux[®] operating systems, perform the following steps:
 - a. Quiesce all applications to the source volumes that you want to FlashCopy.
 - b. Use the **umount** command to unmount the designated drives.
 - c. Prepare and start the FlashCopy for those unmounted drives.
 - d. Remount 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 that you want to FlashCopy.
 - b. Go into your disk management window and remove the drive letter on each drive to you want to copy. This unmounts the drive.
 - c. Prepare and start the FlashCopy for those unmounted drives.

- d. Remount 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 that you want to FlashCopy.
- b. Issue the **chkdsk /x** command on each drive you want to copy. The /x option unmounts, scans, and remounts 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.

FlashCopy mappings

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

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

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

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

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

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

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

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

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

If the background copy rate is zero (or NOCOPY), only the data that changes on the source is copied to the target. The target never contains a copy of the whole source unless every extent is overwritten at the source. You can use this copy rate when you require 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 the target VDisk is taken offline. You must restart the mapping to correct the target.

FlashCopy mapping states

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

Idle or copied

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

Copying

The copy is in progress.

Prepared

The mapping is ready to start. The target VDisk is online, but not accessible. The target VDisk cannot perform read or write caching. Read and write caching is failed by the SCSI front-end as a hardware error.

Preparing

The target VDisk is online, but not accessible. The target VDisk cannot perform read or write caching. Read and write caching is failed by the SCSI front-end as a hardware error. 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 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. Preparing the mapping ensures that the data in the cache is de-staged to disk and a consistent copy of the source exists on disk. At this time, the cache goes into write-through mode. Data that is written to the source is not cached in the SAN Volume Controllers; it passes straight through to the MDisks. The prepare operation for the mapping might take you a few minutes; the actual length of time depends on the size of the source VDisk. You must coordinate the prepare operation with the operating system. Depending on the type of data that is on the source VDisk, the operating system or application

software might also cache data write operations. You must flush, or synchronize, the file system and application program before you prepare for and start the mapping.

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

If you do not want to use consistency groups, the SAN Volume Controller allows a FlashCopy mapping to be treated as an independent entity. In this case the FlashCopy mapping is known as a stand-alone mapping. For FlashCopy mappings which have been configured in this way, the **svctask prestartfcmap** and **svctask startfcmap** commands are directed at the FlashCopy mapping name rather than the consistency group ID.

Veritas Volume Manager

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

Related reference

“FlashCopy mapping events”

FlashCopy mapping events detail the events that modify the state of a FlashCopy mapping.

FlashCopy mapping events:

FlashCopy mapping events detail the events that modify the state of a FlashCopy mapping.

Table 7 provides a description of each FlashCopy mapping event.

Table 7. FlashCopy mapping events

Create	A new FlashCopy mapping is created between the specified source virtual disk (VDisk) and the specified target VDisk. The operation fails if either the source or target VDisk are already a member of a FlashCopy mapping. The operation fails if the SAN Volume Controller has insufficient bitmap memory. The operation also fails if the source and target VDisks 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 into the preparing state. Attention: The prepare command can corrupt any data which previously resided on the target VDisk because cached writes are discarded. Even if the FlashCopy mapping is never started, the data from the target might have logically changed during the act of preparing to start the FlashCopy mapping.
Flush done	The FlashCopy relationship automatically moves from the preparing state to the prepared state after all cached data for the source is flushed and all cached data for the target is invalidated.

Table 7. FlashCopy mapping events (continued)

Start	<p>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.</p> <p>To preserve the cross volume consistency group, the start of all of the FlashCopy mappings in the consistency group must be synchronized correctly with respect to I/Os that are directed at the VDIs. This is achieved during the start command.</p> <p>The following occurs during the start command:</p> <ul style="list-style-type: none"> • New reads and writes to all source VDIs in the consistency group are paused in the cache layer until all ongoing reads and writes below the cache layer are completed. • After all FlashCopy mappings in the consistency group are paused, the internal cluster state is set to allow FlashCopy operations. • After the cluster state is set for all FlashCopy mappings in the consistency group, read and write operations are unpaused on the source VDIs. • The target VDIs are brought online. <p>As part of the start command, read and write caching is enabled for both the source and target VDIs.</p>
Modify	<p>A FlashCopy mapping has two properties that can be modified. These properties are the background copy rate and the consistency group. The background copy rate can be modified during any state. The consistency group can only be modified during the idling, copied or stopped state.</p>
Stop	<p>There are two separate mechanisms by which a FlashCopy mapping can be stopped:</p> <ul style="list-style-type: none"> • You have issued a command • An I/O error has occurred
Delete	<p>This command requests that the specified FlashCopy mapping is deleted. If the FlashCopy mapping is in the stopped state, the force flag must be used.</p> <p>Deleting a FlashCopy mapping in the stopped state can allow unflushed write data from the cache to be destaged to what was the target VDisk.</p> <p>The destaging of old data to what was the target VDisk does not affect the future use of the VDisk because any new data is written over this old data, in the cache or on disk.</p>
Flush failed	<p>If the flush of data from the cache cannot be completed, the FlashCopy mapping enters the stopped state.</p>
Copy complete	<p>After all of the source data has been copied to the target, the source and target are made independent and the state is set to copied. The FlashCopy mapping is not automatically deleted at this time and can be reactivated by preparing and starting again.</p>
Bitmap Online/Offline	<p>The node has failed.</p>

Related concepts

“FlashCopy mappings” on page 38

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

FlashCopy consistency groups

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

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

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

When you copy data from one VDisk to another, the data might not include all that you need to enable you to use the copy. Many applications have data that spans multiple VDIs and requires that data integrity is preserved across VDIs. 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 VDIs. In this situation, FlashCopy must be performed in a way that preserves data integrity across the multiple VDIs. 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 VDIs act independently even if a FlashCopy consistency group exists. Read and write caching is enabled for the source VDIs and target VDIs.

Copying

The copy is in progress.

Prepared

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

Preparing

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

Stopped

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

Suspended

The consistency group was started, but it did not complete. The source VDIs 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 reference

“Dependent writes”

One requirement for preserving the integrity of data that is being written is to ensure that dependent writes are run in the intended sequence of the application.

“Operations on consistency groups”

You can create, change, and delete consistency groups.

“FlashCopy limits for consistency groups” on page 44

The SAN Volume Controller supports up to 2048 FlashCopy mappings. The SAN Volume Controller supports up to 512 FlashCopy mappings per consistency group.

Dependent writes:

One requirement for preserving the integrity of data that is being written is to ensure that dependent writes are run in the intended sequence of the application.

The following list is a typical sequence of write operations for a database update transaction:

1. A write operation updates the database log so that it indicates that a database update is about to take place.
2. A second write operation updates the database.
3. A third write operation updates the database log so that it indicates that the database 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 (VDisks) 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 database 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 database would be in question.

You can perform a FlashCopy operation on multiple VDIs as an atomic operation to create a consistent image of user data. To use FlashCopy this way, the SAN Volume Controller 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 cluster. The SAN Volume Controller allows the **svctask startfcconsistgrp** 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 supports 128 consistency groups per cluster.

Related concepts

“FlashCopy consistency groups” on page 42

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

Operations on consistency groups:

You can create, change, and delete consistency groups.

You can create, change, and delete 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 SAN Volume Controller Console.

Related concepts

“FlashCopy consistency groups” on page 42

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

FlashCopy limits for consistency groups:

The SAN Volume Controller supports up to 2048 FlashCopy mappings. The SAN Volume Controller supports up to 512 FlashCopy mappings per consistency group.

Related concepts

“FlashCopy consistency groups” on page 42

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

FlashCopy indirection layer

FlashCopy provides the semantics of a point in time copy by using an indirection layer which intercepts I/Os targeted at both the source and target virtual disks.

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 determination about each I/O. This determination is based upon the following criteria:

- 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 FlashCopy 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.

Related reference

“Grains and the FlashCopy bitmap” on page 45

When data is copied from the source virtual disk to the target virtual disk, it is copied in units of address space known as grains.

“Source and target reads” on page 45

The source and target must be of equal size. A 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.

“Writes to the source or target” on page 45

Where writes occur to source or target to an area (or grain) which has not yet been copied, these will usually be delayed while a copy operation is performed to copy data from the source to the target, to maintain the illusion that the target contains its own copy.

“FlashCopy mapping limits” on page 46

The SAN Volume Controller supports up to 2048 FlashCopy mappings in a single cluster.

Grains and the FlashCopy bitmap:

When data is copied from the source virtual disk to the target virtual disk, it is copied in units of address space known as grains.

In the SAN Volume Controller, 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.

Related reference

“FlashCopy indirection layer” on page 44

FlashCopy provides the semantics of a point in time copy by using an indirection layer which intercepts I/Os targeted at both the source and target virtual disks.

Source and target reads:

The source and target must be of equal size. A 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.

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 check its bitmap to see if the data has already been copied. If the data being read has already been copied to the target, the read is sent to the target disk. If it has not, 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 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.

Related reference

“FlashCopy indirection layer” on page 44

FlashCopy provides the semantics of a point in time copy by using an indirection layer which intercepts I/Os targeted at both the source and target virtual disks.

Writes to the source or target:

Where writes occur to source or target to an area (or grain) which has not yet been copied, these will usually be delayed while a copy operation is performed to copy data from the source to the target, to maintain the illusion that the target contains its own copy.

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.

Related reference

“FlashCopy indirection layer” on page 44
 FlashCopy provides the semantics of a point in time copy by using an indirection layer which intercepts I/Os targeted at both the source and target virtual disks.

FlashCopy mapping limits:

The SAN Volume Controller supports up to 2048 FlashCopy mappings in a single cluster.

You can have a maximum of 16 TB of virtual disk (VDisk) space for both the source and target VDIsks participate in FlashCopy mappings in any one I/O group in a single cluster.

Related reference

“FlashCopy indirection layer” on page 44
 FlashCopy provides the semantics of a point in time copy by using an indirection layer which intercepts I/Os targeted at both the source and target virtual disks.

Background copy

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

If NOCOPY is specified, background copy is disabled. You can specify NOCOPY for short-lived FlashCopy mappings that are only used for backups. Since the source data set is not expected to significantly change during the lifetime of the FlashCopy mapping, it is more efficient in terms of managed disk (MDisk) I/Os to not perform a background copy.

Table 8 provides the relationship of the background copy rate value to the attempted number of grains to be split per second. A grain is the unit of data represented by a single bit, which is 256K.

Table 8. Background copy

User-specified value	KB/sec	Grains/sec
1 - 10	128	0.5
11 - 20	256	1
21 - 30	512	2
31 - 40	1024	4
41 - 50	2048	8
51 - 60	4096	16
61 - 70	8192	32
71 - 80	16384	64
81 - 90	32768	128
91 - 100	65536	256

The grains/sec numbers represent standards that the SAN Volume Controller tries to achieve. The SAN Volume Controller is unable to achieve these standards if insufficient bandwidth is available from the nodes to the physical disks that make up the managed disks (MDisks) after taking into account the requirements of foreground I/O. If this situation occurs, background copy I/O contends for resources

on an equal basis with I/O that arrives from hosts. Both tend to see an increase in latency and consequential reduction in throughput with respect to the situation had the bandwidth not been limited.

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

The background copy is performed by one of the nodes that belongs to the I/O group in which the source VDisk resides. This responsibility is failed over to the other node in the I/O group in the event of the failure of the node that performs the background copy.

The background copy starts with the grain that contains the highest logical block numbers (LBAs) and works in reverse towards the grain that contains LBA 0. The background copy is performed in reverse to avoid any unwanted interactions with sequential write streams from the application.

Related concepts

“FlashCopy” on page 35

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

Metro Mirror

Metro Mirror enables you to set up a relationship between two virtual disks (VDisks), so that updates that are made by an application to one VDisk are mirrored on the other VDisk.

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

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

Metro Mirror supports the following features:

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

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

- Intercluster and intra-cluster Metro Mirror can be used concurrently within a cluster.
- The intercluster link is bidirectional. This means that it can copy data from clusterA to clusterB for one pair of VDIsks while copying data from clusterB to clusterA for a different pair of VDIsks.

- The copy direction can be reversed for a consistent relationship by issuing either the **svctask switchrconsistgrp** command for a Metro Mirror consistency group or **svctask switchrrelationship** command for a Metro Mirror relationship. See *IBM TotalStorage SAN Volume Controller: Command-Line Interface User's Guide*.
- Metro Mirror consistency groups are supported for ease of managing a group of relationships that must be kept synchronized for the same application. This also simplifies administration, because a single command that is issued to the consistency group is applied to all the relationships in that group.

Related concepts

“Metro Mirror for SAN Volume Controller”

Metro Mirror for SAN Volume Controller provides a *consistent* copy, which means that the primary virtual disk (VDisk) is always the exact match of the secondary VDisk.

“Metro Mirror partnerships”

With Metro Mirror, you can copy a virtual disk (VDisk) in one cluster to a VDisk in another cluster.

“Metro Mirror relationships” on page 49

A Metro Mirror relationship defines the relationship between two virtual disks (VDisks): a master VDisk and an auxiliary VDisk.

“Metro Mirror consistency groups” on page 51

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

Metro Mirror for SAN Volume Controller

Metro Mirror for SAN Volume Controller provides a *consistent* copy, which means that the primary virtual disk (VDisk) is always the exact match of the secondary VDisk.

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

Related concepts

“Metro Mirror” on page 47

Metro Mirror enables you to set up a relationship between two virtual disks (VDisks), so that updates that are made by an application to one VDisk are mirrored on the other VDisk.

Metro Mirror partnerships

With Metro Mirror, you can copy a virtual disk (VDisk) in one cluster to a VDisk in another cluster.

The SAN Volume Controller must know not only about the relationship between the two VDIs but also about the relationship between the two clusters. A Metro Mirror 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 must first issue the **svctask mkpartnership** command from clusterA, and specifying clusterB as the remote cluster. At this point the partnership is partially configured, and sometimes described as one-way communication. Next, you must issue the **svctask mkpartnership**

command from clusterB and specify 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 concepts

“Metro Mirror” on page 47

Metro Mirror enables you to set up a relationship between two virtual disks (VDisks), so that updates that are made by an application to one VDisk are mirrored on the other VDisk.

Background copy bandwidth impact on foreground I/O latency

The background copy bandwidth determines the rate at which the background copy for the IBM TotalStorage Metro Mirror for SAN Volume Controller will be attempted.

The background copy bandwidth can affect foreground I/O latency in one of three ways:

- If the background copy bandwidth is set too high for the Metro Mirror intercluster link capacity, the following results can occur:
 - The background copy I/Os can back up on the Metro Mirror intercluster link.
 - There is a delay in the synchronous secondary writes of foreground I/Os.
 - The foreground I/O latency will increase as perceived by applications.
- If the background copy bandwidth is set too high for the storage at the *primary* site, background copy read I/Os overload the primary storage and delay foreground I/Os.
- If the background copy bandwidth is set too high for the storage at the *secondary* site, background copy writes at the secondary overload the secondary storage and again delay the synchronous secondary writes of foreground I/Os.

In order to set the background copy bandwidth optimally, make sure that you consider all three resources (the primary storage, the intercluster link bandwidth and the secondary storage). Provision the most restrictive of these three resources between the background copy bandwidth and the peak foreground I/O workload.

This provisioning can be done by calculation as above or alternatively by determining experimentally how much background copy can be allowed before the foreground I/O latency becomes unacceptable and then backing off to allow for peaks in workload and some safety margin.

Related concepts

“FlashCopy” on page 35

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

Metro Mirror relationships

A Metro Mirror relationship defines the relationship between two virtual disks (VDisks): a master VDisk and an auxiliary VDisk.

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 VDIs 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 Metro Mirror 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, you can reverse the copy direction 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 Metro Mirror consistency group, for ease of application management.

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

Metro Mirror states

When a Metro Mirror 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 must 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 **svctask startrcrelationship** command is issued to an consistency group in the InconsistentStopped state. This state is also entered when a **svctask startrcrelationship** 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 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 cannot 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 can accept read I/O operations but not write I/O operations

Related concepts

“Metro Mirror” on page 47

Metro Mirror enables you to set up a relationship between two virtual disks (VDisks), so that updates that are made by an application to one VDisk are mirrored on the other VDisk.

Metro Mirror consistency groups

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

Metro Mirror relationships can be based on “loose” or “tight” associations. A more significant use arises when the relationships contain virtual disks (VDisks) with a tight association. A simple example of a tight association is the spread of data for an application 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. In both examples, specific rules exist as to how the relationships must be manipulated, in unison. This ensures that the set of secondary VDIsks contain usable data. The key property is that these relationships are consistent.

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

|
|
|
|
|
|
|
|

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 primary and secondary clusters, sometimes referred to as master and auxiliary clusters. All relationships in a consistency group must also have the same copy direction and state.

Metro Mirror consistency group states

Inconsistent (stopped)

The primary VDIsks are accessible for read and write input/output (I/O)

operations but the secondary VDisks are not accessible for either. A copy process must 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 the **svctask startrcconsistgrp** command is issued to a consistency group in the InconsistentStopped state. This state is also entered when the **svctask startrcconsistgrp** 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 Both the primary VDisks and the secondary VDisks are operating in the primary role. Consequently the VDisks are accessible for write I/O operations.

Idling (disconnected)

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

Inconsistent (disconnected)

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

Consistent (disconnected)

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

Empty

The consistency group contains no relationships.

Related concepts

“Metro Mirror” on page 47

Metro Mirror enables you to set up a relationship between two virtual disks (VDisks), so that updates that are made by an application to one VDisk are mirrored on the other VDisk.

Configuration rules and requirements

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

Table 9 on page 53 provides terms and definitions that can guide your understanding of the rules and requirements.

Table 9. Configuration terms and definitions

Term	Definition
ISL hop	A hop on an interswitch link (ISL). With reference to all pairs of N-ports or end-nodes that are in a fabric, the number of ISL hops is the number of links that are crossed on the shortest route between the node pair whose nodes are farthest apart from each other. The distance is measured only in terms of the ISL links that are in the fabric.
Oversubscription	The ratio of the sum of the traffic that is on the initiator N-node connections to the traffic that is on the most heavily-loaded ISLs or where more than one ISL is in parallel between these switches. This definition assumes a symmetrical network and a specific workload that is applied equally from all initiators and sent equally to all targets. A symmetrical network means that all initiators are connected at the same level and all the controllers are connected at the same level. Note: The SAN Volume Controller puts its back-end traffic onto the same symmetrical network. The back-end traffic can vary by workload. Therefore, the oversubscription that a 100% read hit gives is different from the oversubscription that 100% write-miss gives. If you have an oversubscription of 1 or less, the network is nonblocking.
Virtual SAN (VSAN)	A VSAN is a virtual storage area network (SAN).
Redundant SAN	A SAN configuration in which if any one component fails, connectivity between the devices that are in the SAN is maintained, possibly with degraded performance. Create a redundant SAN by splitting the SAN into two independent counterpart SANs.
Counterpart SAN	A non-redundant portion of a redundant SAN. A counterpart SAN provides all the connectivity of the redundant SAN, but without the redundancy. The SAN Volume Controller is typically connected to a redundant SAN that is made out of two counterpart SANs.
Local fabric	The fabric that consists of those SAN components (switches and cables) that connect the components (nodes, hosts, and switches) of the local cluster. Because the SAN Volume Controller supports Metro Mirror, significant distances might exist between the components of the local cluster and those of the remote cluster.
Remote fabric	The fabric that consists of those SAN components (switches and cables) that connect the components (nodes, hosts, and switches) of the remote cluster. Because the SAN Volume Controller supports Metro Mirror, significant distances might exist between the components of the local cluster and those of the remote cluster.
Local/remote fabric interconnect	The SAN components that connect the local fabrics to the remote fabrics. There might be significant distances between the components in the local cluster and those in the remote cluster. These components might be single-mode optical fibres that are driven by Gigabit Interface Converters (GBICs), or they might be other, more advanced components, such as channel extenders
SAN Volume Controller fibre-channel port fan in	The number of hosts that can see any one port. Some controllers recommend that the number of hosts using each port be limited to prevent excessive queuing at that port. If the port fails or the path to that port fails, the host might failover to another port, and the fan in requirements might be exceeded in this degraded mode.
Invalid configuration	In an invalid configuration, an attempted operation fails and will generate an error code to indicate what caused it to become invalid.

Table 9. Configuration terms and definitions (continued)

Term	Definition
Unsupported configuration	A configuration that might operate successfully, but for which IBM does not guarantee the solution for problems that might occur. Usually this type of configuration does not create an error log entry.
Valid configuration	A configuration that is neither invalid nor unsupported.
Degraded	A valid configuration that has had a failure, but continues to be neither invalid nor unsupported. Typically, a repair action is required to restore the degraded configuration to a valid configuration.
Fibre channel extender	A device for long distance communication connecting other SAN fabric components. Generally these might involve protocol conversion to ATM, IP or some other long distance communication protocol.
Mesh configuration	A network that contains a number of small SAN switches configured to create a larger switched network. With this configuration, four or more switches are connected together in a loop with some of the paths short circuiting the loop. An example of this configuration is to have four switches connected together in a loop with ISLs for one of the diagonals. The SAN Volume Controller does not support this configuration.

Related tasks

“Configuration requirements” on page 76

Ensure that you are familiar with the configuration requirements for the SAN Volume Controller. You must abide by the configuration requirements for the SAN Volume Controller to ensure you have a valid configuration.

Related reference

“Configuration rules”

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

“Maximum configuration” on page 77

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

Configuration rules

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

A SAN configuration that contains SAN Volume Controllers is valid if it observes *all* of the following rules. Some configurations do not work and are known as *invalid*.

Related concepts

“Storage subsystems” on page 55

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

Related reference

“HBAs” on page 59

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

“Nodes” on page 60

Ensure that you are familiar with the configuration rules for SAN Volume

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

“Power requirements” on page 61

Ensure you are familiar with the configuration rules for power requirements. You must abide by the configuration rules for power requirements to ensure you have a valid configuration.

“Fibre-channel switches” on page 61

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

Storage subsystems

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

All SAN Volume Controller nodes of a cluster must be able to see the same set of storage subsystem ports on each device. Any operation that is in this mode in which two nodes do not see the same set of ports on the same device is degraded, and the system logs errors that request a repair action. This rule can have important effects on a storage subsystem such as DS4000 series, which has exclusion rules that determine to which host bus adapter (HBA) WWNNs a storage partition can be mapped.

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

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

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

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

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

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

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

Split controller configurations

The SAN Volume Controller is configured to manage LUs that are exported only by RAID controllers. Operation with other RAID controllers is illegal. While it is possible to use the SAN Volume Controller to manage just a bunch of disks (JBOD) LUs presented by supported RAID controllers, it should be noted that the SAN Volume Controller itself provides no RAID function, so these LUs are exposed to data loss in the event of a disk failure.

If a single RAID controller presents multiple LUs, either by having multiple RAID configured or by partitioning one or more RAID into multiple LUs, each LU can be owned by either SAN Volume Controller or a directly attached host. Suitable LUN masking must be in place to ensure that LUs are not shared between SAN Volume Controllers and direct attached hosts.

In a split controller configuration, a RAID presents some of its LUs to a SAN Volume Controller (which treats the LU as an MDisk) and the remaining LUs to another host. The SAN Volume Controller presents virtual disks (VDisks) that are 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). Figure 17 on page 57 shows that the RAID controller is a DS4000, with RDAC used for pathing on the directly attached host, and SDD used on the host that is attached with the SAN Volume Controller. Hosts can simultaneously access LUs that are provided by the SAN Volume Controller and directly by the device.




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

It is also possible to split a host so that it accesses some of its LUNs through the SAN Volume Controller and some directly. In this case, the multipathing software that is used by the controller must be compatible with SAN Volume Controllers multipathing software. Figure 18 on page 58 is a supported configuration because the same pathing driver is used for both direct and VDisks.




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

In the case where the RAID controller uses multipathing software that is compatible with SAN Volume Controllers multipathing software (see Figure 19 on page 59), it is possible to configure a system where some LUNs are mapped directly to the host and others are accessed through the SAN Volume Controller. An ESS that uses the same multipathing driver as SAN Volume Controller is one example.




Figure 19. DS4000 direct connection with a SAN Volume Controller on one host

Related reference

“Configuration rules” on page 54

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

HBAs

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

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

HBAs that are in dissimilar hosts or dissimilar HBAs that are in the same host must be in separate zones. *Dissimilar* means that the hosts are running different operating systems or that they are different hardware platforms. For example, if you have an HP/UX host and a Windows 2000 server host, those hosts must be in separate zones. Different levels of the same operating system are considered to be similar. A configuration that breaks this requirement is not supported.

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

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

Operation with other HBAs is not supported.

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

Related reference

“Configuration rules” on page 54

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

Nodes

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

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

Note: Configurations consisting of a single Cisco MDS 9000 with two, four, six, and eight nodes are valid, although the configuration is not fully redundant.

A SAN Volume Controller 2145-8F2 and a SAN Volume Controller 2145-4F2 cannot be in the same I/O group.

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

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

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

- Node to a switch
- Host to a switch
- Back-end to a switch

- Interswitch links (ISLs)

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

- Short wave optical fiber
- Long wave optical fiber up to 10 km

High-power Gigabit Interface Converters (GBICs) and long wave fibre connections beyond 10 km are not supported.

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

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

The SAN Volume Controller nodes must be connected to the uninterruptible power supply (UPS) using the supplied cable that joins the signal and power cables.

Related reference

“Configuration rules” on page 54

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

Power requirements

Ensure you are familiar with the configuration rules for power requirements. You must abide by the configuration rules for power requirements to ensure you have a valid configuration.

The uninterruptible power supply (UPS) must be in the same rack that contains the SAN Volume Controller node (or nodes) that it supplies. The SAN Volume Controller 2145-8F2 must be connected to a 2145 uninterruptible power supply-1U (2145 UPS-1U) because it cannot operate with a 2145 uninterruptible power supply (2145 UPS).

The combination power and signal cable for connection between the SAN Volume Controller and the UPS units is 2 m long. The SAN Volume Controller and UPS must connect with both the power and the signal cable to function correctly.

Related reference

“Configuration rules” on page 54

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

Fibre-channel switches

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

The SAN must contain only supported switches.

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

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

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

Configurations with more than four SANs are not supported.

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

The SAN Volume Controller nodes must always and only be connected to SAN switches. Each node must be connected to each of the counterpart SANs that are in the redundant fabric. Any configuration that uses direct connections between host and node or between controller and node is not supported. SANs that are created from a mesh of switches are not supported.

All 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 DS4000 configuration in which the DS4000 contains two redundant controllers, only two controller minihubs are usually used. Controller A of the DS4000 is connected to counterpart SAN A, and controller B of the DS4000 is connected to counterpart SAN B. Any configuration that uses a direct connection between the host and the controller is not supported.

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

When you attach the SAN Volume Controller to a SAN fabric that contains core directors and edge switches, connect the SAN Volume Controller ports to the core directors and connect the host ports to the edge switches. In this type of fabric, the next priority for connection to the core directors is the storage controllers, leaving the host ports connected to the edge switches.

The switch configuration of a SAN Volume Controller SAN must observe the switch manufacturer's configuration rules. These rules might put restrictions on the switch configuration. Any configuration that runs outside the manufacturers' rules is not supported.

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

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

All SAN Volume Controller nodes in a SAN Volume Controller cluster must be able to see the same set of back-end storage ports on each back-end controller. Operation in a mode where two nodes see a different set of ports on the same controller is degraded and the system logs errors that request a repair action. This can occur if inappropriate zoning was applied to the fabric or if inappropriate LUN masking is used. This rule has important implications for back-end storage, such as a DS4000, which imposes exclusive rules on which HBA worldwide node names (WWNNs) a storage partition can be mapped to.

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

Mixing manufacturer switches in a single SAN fabric

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

- BladeCenter®. See the documentation that is provided with your BladeCenter for more information.
- Where one pair of counterpart fabrics (for example, Fabric A and Fabric B) provide a redundant SAN, different manufacturer's switches can be mixed in a SAN Volume Controller configuration, provided that each fabric contains only switches from a single manufacturer. Thus, the two counterpart SANs can have different manufacturer's switches.
- The SAN Volume Controller supports the Interoperability Modes of the Cisco MDS 9000 family of switch and director products with the following restrictions:
 - The Cisco MDS 9000 must be connected to Brocade and McData switch/director products with the multivendor fabric zones connected using MDS Interoperability Mode 1, 2 or 3.
 - All of the SAN Volume Controller nodes that are in the SAN Volume Controller cluster must be attached to the Cisco part of the counterpart fabric or they must be attached to the McData or Brocade part of the counterpart fabric to avoid having a single fabric with a SAN Volume Controller cluster that has part of the SAN Volume Controller nodes connected to Cisco switch ports and part of the SAN Volume Controller nodes connected to Brocade or McData switch ports.

Fibre-channel switches and interswitch links

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

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

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

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

With ISLs between nodes in the same cluster, the ISLs are considered a single point of failure. This is illustrated in Figure 20.




Figure 20. Fabric with ISL between nodes in a cluster

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

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

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

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

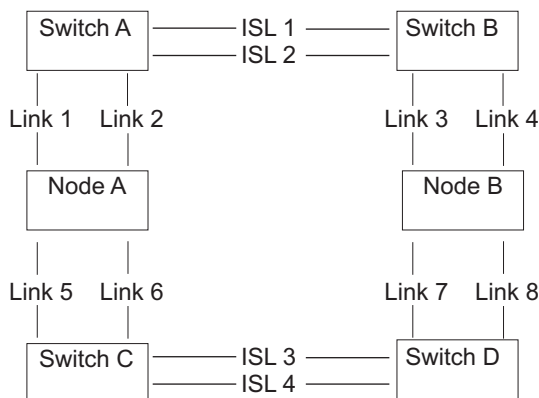


Figure 21. Fabric with ISL in a redundant configuration

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

SAN Volume Controller in a SAN with director class switches

You can use director class switches within the SAN to connect large numbers of RAID controllers and hosts to a SAN Volume Controller cluster. Because director

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

Related concepts

“IBM BladeCenter fabric support”

Ensure that you use fabric and switch modules that are supported by the SAN Volume Controller environment.

Related reference

“Configuration rules” on page 54

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

Switch zoning for the SAN Volume Controller

Ensure that you are familiar with the constraints for zoning a switch.

IBM BladeCenter fabric support

Ensure that you use fabric and switch modules that are supported by the SAN Volume Controller environment.

See the following Web site for the latest support and information:

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

Port switches

The SAN Volume Controller can be used in a SAN with either 16-port or 32-port switches.

Using a 16-port switch SAN

Use a SAN with 16-port switches if your configuration requires only two SAN Volume Controller nodes. A typical configuration has two SAN Volume Controller nodes and up to four RAID controller pairs. With this configuration, the SAN Volume Controller nodes and RAID controllers use eight of the port switches and leave the remaining eight ports on the switch for host connections. You can also adjust the configuration to include more RAID controllers and less host connections. If you want to include an optional service node, you can attach the service node to one fibre-channel port on each of the switches.

Using a 32-port switch SAN

Use a SAN with 32-port switches if your configuration requires more than two SAN Volume Controller nodes. An example configuration for a large SAN Volume Controller based SAN uses a dual switch, redundant SAN fabric and four SAN Volume Controller nodes. The SAN Volume Controller nodes use 16 of the ports and the remaining 16 ports are used for RAID controllers and host connections.

Switch zoning for the SAN Volume Controller

Ensure that you are familiar with the constraints for zoning a switch.

Overview

The number of virtual paths to each virtual disk (VDisk) is limited. The following guidelines can help you achieve the correct number of virtual paths:

- Each host (or partition of a host) can have between one and four fibre-channel ports.
- Brocade and McData switches can be configured in Vendor Interoperability Mode or in Native Mode.
- The SAN Volume Controller supports the Interoperability Modes of the Cisco MDS 9000 family of switch and director products with the following restrictions:
 - The Cisco MDS 9000 must be connected to Brocade and McData switch/director products with the multivendor fabric zones connected using MDS Interoperability Mode 1, 2 or 3.
 - All of the SAN Volume Controller nodes that are in the SAN Volume Controller cluster must be attached to the Cisco part of the counterpart fabric or they must be attached to the McData or Brocade part of the counterpart fabric to avoid having a single fabric with a SAN Volume Controller cluster that has part of the SAN Volume Controller nodes connected to Cisco switch ports and part of the SAN Volume Controller nodes connected to Brocade or McData switch ports.
- The fabric uses the following default timeout values:
 - E_A_TOV=10
 - E_D_TOV=2

Operation with values other than these default timeout values is not supported.

You must manually set the domain IDs prior to building the multiswitch fabric and prior to zoning for the following reasons:

- When two switches are joined while they are active, they can determine if the domain ID is already in use. If there is a conflict, the domain ID cannot be changed in an active switch. This conflict causes the fabric merging process to fail.
- The domain ID identifies switch ports when zoning is implemented using the domain and switch port number. If domain IDs are negotiated at every fabric start up, there is no guarantee that the same switch will have the same ID the next time. Therefore, zoning definitions can become invalid.
- If the domain ID is changed after a SAN is set up, some host systems might have difficulty logging back in with the switch, and it might be necessary to reconfigure the host in order to detect devices on the switch again.

The maximum number of paths from the SAN Volume Controller nodes to a host is eight. The maximum number of host bus adapter (HBA) ports is four. For example, you cannot have more than two 2-port HBAs or four 1-port HBAs).

Consider the SAN environment in the following example:

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

- Each of the four hosts has four ports, as described in Table 10.

Table 10. Four hosts and their ports

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

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

The following is an example configuration:

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

Create the following host zones on switch X:

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

Create the following host zones on switch Y:

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

Create the following storage zone:

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

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

- Two SAN Volume Controller nodes called A and B
- Nodes A and B have four ports each
 1. Node A has ports A0, A1, A2, and A3
 2. Node B has ports B0, B1, B2, and B3
- Six hosts called P, Q, R, S, T and U
- Four hosts have four ports each and the other two hosts have two ports each as described in Table 11.

Table 11. Six hosts and their ports

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

Table 11. Six hosts and their ports (continued)

P3	Q3	R3	S3	—	—
----	----	----	----	---	---

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

The following is an example configuration:

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

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

Create the following host zones on switch X:

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

Create the following host zones on switch Y:

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

Create the following storage zone:

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

Related reference

Switch zoning limitations for the EMC CLARiiON

There are limitations in switch zoning for the SAN Volume Controller and EMC CLARiiON.

Fibre-channel switches

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

Zoning guidelines

Ensure that you are familiar with the following zoning guidelines.

Paths to hosts

- The number of paths through the network from the SAN Volume Controller nodes to a host must not exceed 8. Configurations in which this number is exceeded are unsupported.
 - Each SAN Volume Controller node has four ports and each I/O group has two SAN Volume Controller nodes. Therefore, without any zoning, the number of paths to a VDisk would be $8 \times$ the number of host ports.

- This rule exists to limit the number of paths that must be resolved by the multipathing device driver.

If you want to restrict the number of paths to a host, zone the switches so that each HBA port is zoned with one SAN Volume Controller port for each node in the cluster. If a host has multiple HBA ports, zone each port to a different set of SAN Volume Controller ports to maximize performance and redundancy.

Controller zones

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

SAN Volume Controller zones

The switch fabric must be zoned so that the SAN Volume Controller nodes can see the back-end storage and the front end host HBAs. Usually, the front-end host HBAs and the back-end storage are not in the same zone. The exception to this is where split host and split controller configuration is in use.

It is possible to zone the switches in such a way that a SAN Volume Controller port is used solely for internode communication, or for communication to host, or for communication to back-end storage. This is possible because each SAN Volume Controller node contains 4 ports. Each SAN Volume Controller node must still remain connected to the full SAN fabric. Zoning cannot be used to separate the SAN into two parts.

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

If a SAN Volume Controller node can see another SAN Volume Controller node through multiple paths (some of which use ISLs but some do not), zoning should be used where possible to ensure that the SAN Volume Controller to SAN Volume Controller communication does not travel over an ISL. Likewise if a SAN Volume Controller can see a storage controller through multiple paths (some of which travel over ISLs and but some do not), you should use zoning to restrict communication to those paths which do not travel over ISLs.

| The SAN Volume Controller zones must ensure that every port on each SAN
| Volume Controller node can see at least one port that belongs to every other SAN
| Volume Controller node in the cluster.

| The SAN Volume Controller zones must ensure that the SAN Volume Controller
| nodes in the local cluster can only see SAN Volume Controller nodes that are in the
| remote cluster. You must avoid any configuration that allows two clusters to see
| each other over fibre channel. You can have one or two hot spare SAN Volume
| Controller nodes that are not members of any cluster zoned to see all of the
| clusters.

Host zones

| The configuration rules for host zones are different depending upon the number of
| hosts that will access the SAN Volume Controller cluster. For smaller configurations

of less than 64 hosts per cluster, the SAN Volume Controller supports a simple set of zoning rules which enable a small set of host zones to be created for different environments. For larger configurations of more than 64 hosts, the SAN Volume Controller supports a more restrictive set of host zoning rules.

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

Clusters with less than 64 hosts

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

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

Clusters with 64 to 256 hosts

For clusters with 64 to 256 hosts attached, each HBA port in a host that connects to a SAN Volume Controller *must* be placed into a separate zone. In this separate zone, you must also include exactly one port from each SAN Volume Controller node in the I/O groups which are associated with this host.

The SAN Volume Controller does not specify the number of host fibre-channel ports or HBAs that a host or a partition of a host can have. The number of host fibre-channel ports or HBAs are specified by the host multipathing device driver. The SAN Volume Controller supports this number; however it is subject to the other configuration rules that are specified here.

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

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

Clusters with 256 to 1024 hosts

For clusters with 256 to 1024 hosts attached, the SAN must be zoned so that each HBA port in a host that connects to a SAN Volume Controller node can only see one SAN Volume Controller port for each node in the I/O group that is associated with the host. If you have 1024 hosts, each host must be associated with only one I/O group and each I/O group must only be associated with up to 256 hosts.

Figure 22 provides an example configuration for zoning 1024 hosts. In this example, the hosts are arranged into four groups of 256 hosts and each group of 256 hosts is zoned to one I/O group. You must zone each group of 256 hosts separately so they cannot see other hosts that are in different I/O groups. The controller zone contains all eight of the SAN Volume Controller nodes and all four of the controllers. The intercluster zone contains all of the SAN Volume Controller nodes that are in both clusters to allow you to use Metro Mirror.

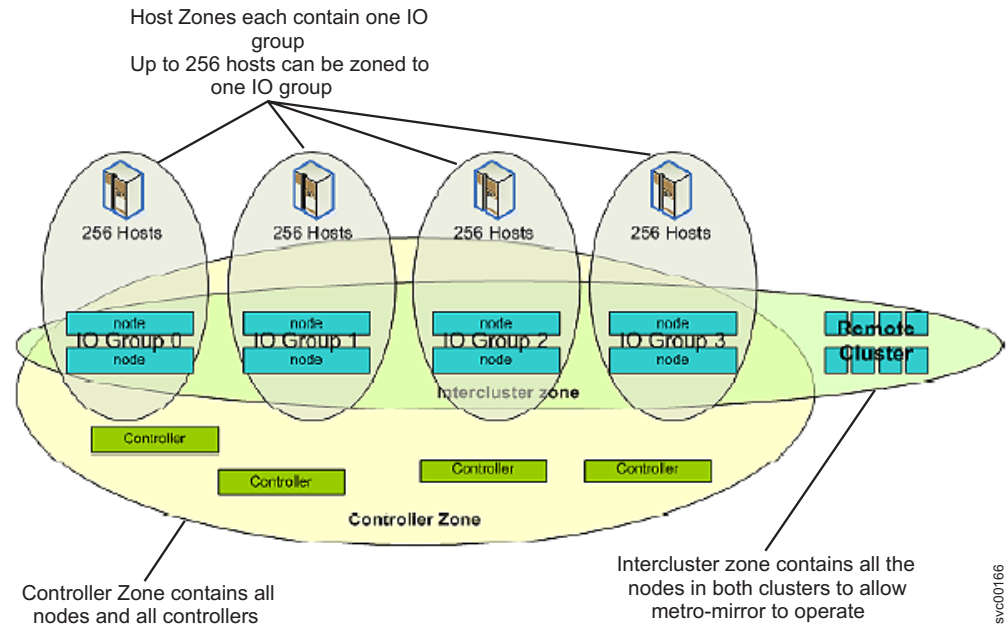


Figure 22. Zoning a 1024 host configuration

You can have up to 512 fibre-channel logins per fibre-channel port. The following logins are counted towards the 512 login maximum:

- Host port logins
- Storage controller logins
- SAN Volume Controller node logins
- Fibre-channel name server logins

If any port has more than 512 logins, the SAN Volume Controller logs an ID 073006 error. You can use the `svcinfolsfabric` command-line interface (CLI) command to list the logins that are seen by each SAN Volume Controller port.

Zoning considerations for Metro Mirror

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

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

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

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

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

1. Form a SAN that contains both clusters that are to be used in the Metro Mirror relationships. If cluster A originally is in SAN A and cluster B is originally in SAN B, there must be at least one fibre-channel connection between SAN A and SAN B. This connection consists of one or more interswitch links. The fibre-channel switch ports that are associated with these interswitch ports must not appear in any zone.
2. Ensure that each switch has a different domain ID before you connect the two SANs. Form a single SAN from the combination of SAN A and SAN B prior to the connection of the two SANs.
3. After the switches in SAN A and SAN B are connected, configure the switches so that they operate as a single group of switches. Each cluster must retain the same set of zones that were required to operate in the original single SAN configuration.

Note: You do not have to perform this step if you are using routing technologies to connect the two SANs.

4. Add a new zone that contains all the switch ports that are connected to SAN Volume Controller ports. This zone contains switch ports that were originally in SAN A and in SAN B.
5. This step is optional because in some cases, this view of both clusters can complicate the way that you operate the overall system. Therefore, unless it is specifically needed, avoid implementing this view. Adjust the switch zoning so that the hosts that were originally in SAN A can recognize cluster B. This allows a host to examine data in both the local and remote cluster, if required.
6. Verify that the switch zoning is such that cluster A cannot recognize any of the back-end storage that is owned by cluster B. Two clusters cannot share the same back-end storage devices.

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

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

Note:

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

Switch operations over long distances

Some SAN switch products provide features that allow the users to tune the performance of I/O traffic in the fabric in a way that can affect Metro Mirror performance. The two most significant features are ISL trunking and extended fabric.

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

Feature	Description
ISL trunking	<p>Trunking enables the switch to use two links in parallel and still maintain frame ordering. It does this by routing all traffic for a given destination over the same route even when there might be more than one route available. Often trunking is limited to certain ports or port groups within a switch. For example, in the IBM 2109-F16 switch, trunking can only be enabled between ports in the same quad (for example, same group of four ports). For more information on trunking with the MDS, refer to "Configuring Trunking" on the Cisco Systems Web site.</p> <p>Some switch types can impose limitations on concurrent use of trunking and extended fabric operation. For example, with the IBM 2109-F16 switch, it is not possible to enable extended fabric for two ports in the same quad. Thus, extended fabric and trunking cannot be used together. Although it is possible to enable extended fabric operation one link of a trunked pair, this does not offer any performance advantages and adds complexity to the configuration setup. Therefore, do not use mixed mode operations.</p>
Extended fabric	<p>Extended fabric operation allocates extra buffer credits to a port. This is important over long links that are usually found in intercluster Metro Mirror operation. Because of the time that it takes for a frame to traverse the link, it is possible to have more frames in transmission at any instant in time than is possible over a short link. The additional buffering is required to allow for the extra frames.</p> <p>For example, the default license for the IBM 2109-F16 switch has two extended fabric options: Normal and Extended Normal.</p> <ul style="list-style-type: none"> • The Normal option is suitable for short links. • The Extended Normal option provides significantly better performance for the links up to 10 km long. <p>Note: With the additional extended fabric license, the user has two extra options: Medium, up to 10 - 50 km and Long, 50 - 100 km. Do not use Medium and Long settings in the intercluster Metro Mirror links that are currently supported.</p>

Limiting queue depth in large SANs

If you are designing a configuration for a large SAN, you must estimate the queue depth for each node in order to avoid application failures.

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

If a SAN Volume Controller node reaches the maximum number of queued commands, many operating systems cannot recover if the situation persists for more than 15 seconds. This can result in one or more servers presenting errors to applications and application failures on the servers.

A large SAN is one in which the number of virtual disks (VDisks) that are mapped to each server and the number of servers is greater than 1 000. For example, 50 servers with each server addressing 20 VDIs.

Queue depth

The queue depth is the number of I/O operations that can be run in parallel on a device. It is usually possible to set a limit on the queue depth on the subsystem device driver (SDD) paths (or equivalent) or the host bus adapter (HBA).

Ensure that you configure the servers to limit the queue depth on all of the paths to the SAN Volume Controller disks in configurations that contain a large number of servers or virtual disks (VDIs).

Note: You might have a number of servers in the configuration that are idle or do not initiate the calculated quantity of I/O operations. If so, you might not need to limit the queue depth.

Calculating a queue depth limit

Several factors are considered in the formula for calculating the queue depth limit.

The formula for queue depth calculation considers the following factors:

- The maximum number of queued commands is per node and there are two nodes in an input/output (I/O) group. The system must continue to function when one of the nodes in an I/O group is not available. Thus, an I/O group is considered to have the same number of queued commands as a node. If a node fails, the number of paths to each disk is cut in half.
- If a virtual disk (VDisk) is mapped so that it can be seen by more than one server, then each of the servers can send commands to it.
- If a device driver times out of a command, it immediately reissues the command. The SAN Volume Controller will have both commands in its command queue.

Homogeneous queue depth calculation

Ensure you are familiar with the homogeneous queue depth calculation.

The homogeneous queues must meet one of the following statements:

- The queued commands are shared among all paths rather than providing servers with additional resources.
- The virtual disks (VDIs) are distributed evenly among the input/output (I/O) groups in the cluster.

You can set the queue depth for each VDisk on the servers using the following calculation:

$$q = ((n \times 7000) / (v \times p \times c))$$

where:

- q is the queue depth per device path
- n is the number of nodes in the cluster
- v is the number of VDIsks configured in the cluster
- p is the number of paths per VDisk per host. A path is a route from a server fibre-channel port to a SAN Volume Controller fibre-channel port that provides the server access to the VDisk.
- c is the number of hosts that can concurrently access each VDisk. Very few applications support concurrent access from multiple hosts to a single VDisk. This number typically is 1.

Example

Consider the following example:

- An eight-node SAN Volume Controller cluster ($n = 8$)
- 4096 VDIsks ($v = 4096$)
- One server with access to each VDisk ($c = 1$)
- Each host has four paths to each VDisk ($p = 4$)

The calculation is $((8 \times 7\ 000) / (4096 \times 4 \times 1)) = 4$.

The queue depth in the operating systems must be set to four concurrent commands per path.

Related reference

“Maximum configuration” on page 77

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

Nonhomogeneous queue depth calculation

For nonhomogeneous queues, use the following calculation.

Nonhomogeneous queues meet one of the following criteria:

- One or more servers are allocated additional resources so that they can queue additional commands.
- VDIsks are not distributed evenly among the I/O groups in the cluster.

Set the queue depth for each VDisk on the servers using the following calculation.

For each VDisk, consider each server to which that VDisk has a mapping. This gives a set of server/VDisk pairs. If the sum of the server and VDisk queue depth for all of the pairs is less than 7 000, the server will not experience problems due to a full queue.

Limiting the queue depth

After you have calculated the queue depth limit, you must apply it.

Each operating system has a way to limit the queue depth on a per virtual disk (VDisk) basis.

An alternative to setting a limit per VDisk is to set a limit on the host bus adapter (HBA). Thus, if the queue depth per path limit is 5, the server has access to 40

VDisks through two adapters (four paths). It might be appropriate to place a queue depth limit of $(40 \times (4 \times 5)) / 2 = 400$ on each adapter. The queue depth limit of $(40 \times (4 \times 5)) / 2 = 400$ on each adapter enables sharing the queue depth allocation between VDisks.

Configuration requirements

Ensure that you are familiar with the configuration requirements for the SAN Volume Controller. You must abide by the configuration requirements for the SAN Volume Controller to ensure you have a valid configuration.

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

1. Your IBM service representative must have installed the SAN Volume Controller.
2. Install and configure your disk controller systems and create the RAID resources that you intend to virtualize. To prevent loss of data, virtualize only those RAID types 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 (VDisks). 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 that you use, the fewer disks that you need to provide availability for the same total capacity (one per array). However, if you use more disks, it takes 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 (stripe 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 that are required to provide availability might be unacceptable if the arrays are too small.

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

If reasonably small RAID types are used, it is easier to extend a managed disk (MDisk) group by adding a new RAID of the same type. Construct multiple RAID devices of the same type, when it is possible.

When you create 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 into more than one SCSI logical unit (LU). When you configure new storage for use with the SAN Volume Controller, you do not have to divide up the array. New storage is presented as one SCSI LU. This gives a one-to-one relationship between MDisks and RAID.

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

3. Install and configure your switches to create the zones that the SAN Volume Controller requires. One zone must contain all the disk controller systems and the SAN Volume Controller nodes. For hosts with more than one port, use switch zoning to ensure that each host fibre-channel port is zoned to exactly one fibre-channel port of each SAN Volume Controller node in the cluster. Set

up a zone on each fibre-channel switch that includes the master console and all of the SAN Volume Controller ports that are connected to that switch.

4. If you want the SAN Volume Controller to export redundant paths to disks, you must install a multipathing device on all of the hosts that are connected to the SAN Volume Controller. Otherwise, you cannot use the redundancy inherent in the configuration. You can install the subsystem device driver (SDD) from the following Web site:

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

5. Install and configure the SAN Volume Controller master console (see the *IBM TotalStorage Master Console Installation and User's Guide*). The communication between the master console and the SAN Volume Controller runs under a client-server network application called Secure Shell (SSH). The SSH server software is already installed on each SAN Volume Controller cluster. The SSH client software called PuTTY is already installed on the master console. You will need to configure the SSH client key pair using PuTTY on the master console.
 - a. You can configure the SAN Volume Controller using the SAN Volume Controller Console Web-based application that is preinstalled on the master console.

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

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

Note:

- AIX comes with an installed SSH client.
- Linux comes with an installed SSH client.
- Use PuTTY for Windows.

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

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

Related reference

“Configuration rules and requirements” on page 52

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

Maximum configuration

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

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

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

Related reference

“Configuration rules and requirements” on page 52

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

Supported fibre-channel extenders

The supported hardware for the SAN Volume Controller frequently changes.

See the following IBM Web site for the latest supported hardware:

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

Related concepts

“Performance of fibre-channel extenders”

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

Performance of fibre-channel extenders

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

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

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

Related reference

“Supported fibre-channel extenders”

The supported hardware for the SAN Volume Controller frequently changes.

Chapter 2. Creating a SAN Volume Controller cluster

You must complete the two phases that are required to create a cluster before you can configure the SAN Volume Controller.

The first phase is performed from the front panel of the SAN Volume Controller. The second phase is performed from the SAN Volume Controller Console which is accessible from a Web server that runs on the master console.

Related concepts

“Browser requirements for the SAN Volume Controller” on page 84
Ensure that you are familiar with the Internet browsers and settings when using the SAN Volume Controller.

Related tasks

“Creating a cluster from the front panel”
After you have created a pair of nodes, you can use the front panel of the SAN Volume Controller to create a cluster.

“Configuring the Web browser” on page 85
You must configure the Web browser to allow new windows to automatically open. These new windows are called popups.

“Changing browser settings for password protection” on page 85
For security reasons, you can configure your Web browser so that a password does not automatically display when you type a user name into the user name field.

“Creating a cluster using the SAN Volume Controller Console” on page 86
After you have created a pair of nodes, you can create and configure a cluster.

Related reference

“SAN Volume Controller Console layout” on page 81
Ensure that you are familiar with the basic frame layout of the SAN Volume Controller Console.

“SAN Volume Controller Console banner” on page 82
The banner of the SAN Volume Controller Console is used for product or customer identification.

“SAN Volume Controller Console task bar” on page 83
The task bar of the SAN Volume Controller Console keeps track of all opened primary tasks and allows you to quickly go back to the previous task or move forward to the next task.

“SAN Volume Controller Console portfolio” on page 83
The portfolio area of the SAN Volume Controller Console contains task-based links that open panels in the work area. Common tasks are grouped under task headings and are expandable and collapsible.

“SAN Volume Controller Console work area” on page 84
The work area of the SAN Volume Controller Console is where you work with a cluster and the objects it contains.

“Accessing the SAN Volume Controller Console” on page 85
The SAN Volume Controller Console is a Web-based application that you can use to manage multiple clusters.

Creating a cluster from the front panel

After you have created a pair of nodes, you can use the front panel of the SAN Volume Controller to create a cluster.

A pair of SAN Volume Controller nodes must be installed before you can create a cluster. You must also check your license to ensure it is correct. The license specifies if you are permitted to use the FlashCopy or Metro Mirror options. It also specifies how much virtualization you are permitted to use.

If you choose to have the customer engineer (CE) initially create the cluster, you must provide the following information to your CE prior to configuring the cluster:

- Cluster IP address
- Subnet mask
- Gateway IP address

Attention: The Cluster IP address must be unique to avoid possible communication problems.

The CE uses the front panel of the SAN Volume Controller to enter the information that you have provided. The SAN Volume Controller generates a random password on the display panel. The CE gives you this password. You must record the password and the IP address. The password and IP address are used to connect to the SAN Volume Controller Console and to create the cluster.

Perform the following steps to create and configure the cluster:

1. Choose a node that you want to make a member of the cluster you are creating.
2. Press and release the up or down navigation button until Node: is displayed on the SAN Volume Controller service panel.
3. Press and release the left or right navigation button until Create Cluster? is displayed.
4. Press the select button.
 - If IP Address: is displayed on line 1 of the screen, go to step 5.
 - If Delete Cluster? is displayed on line 1 of the service display screen, this node is already a member of a cluster. Either you have selected the wrong node, or you have already used this node in an existing cluster. The ID of this existing cluster is displayed on line 2 of the service display screen.
 - If you have selected the wrong node you can exit this task now. The procedure cancels automatically after 60 seconds.
 - If you no longer need the data that is contained on the node and you want to delete the node from the existing cluster, perform the following steps:
 - a. Press and hold the up navigation button.
 - b. Press and release the select button. The node restarts. After the node has restarted, you must restart this task from step 1. IP Address: is displayed on the service display screen.
5. Press and hold the down navigation button to display the menu that shows the current value of the fibre-channel speed setting for the node.
6. Press the select button when the display shows the status of one of the fibre-channel ports. The setting must be either 1 Gb or 2 Gb.

Perform the following steps to change the speed of all the fibre-channel ports on the node:

- a. Press the up or down navigation button to select the speed.
 - b. Press the select button when the speed you want is displayed.
7. Press the select button.

8. Use the up or down navigation button to change the value of the first field of the IP Address to the value that you have chosen.
9. Use the right navigation button to move to the next field and use the up or down navigation button to change the value of this field.
10. Repeat step 9 for each of the remaining fields of the IP Address.
11. Press the select button when you have changed the last field of the IP Address.
12. Press the right navigation button. Subnet Mask: is displayed.
13. Press the select button.
14. Use the up or down navigation button to change the value of the first field of the Subnet Mask to the value that you have chosen.
15. Use the right navigation button to move to the next field and use the up or down navigation buttons to change the value of this field.
16. Repeat step 15 for each of the remaining fields of the Subnet Mask.
17. Press the select button when you have changed the last field of Subnet Mask.
18. Press the right navigation button. Gateway: is displayed.
19. Press the select button.
20. Use the up or down navigation button to change the value of the first field of the Gateway to the value that you have chosen.
21. Use the right navigation button to move to the next field and use the up or down navigation button to change the value of this field.
22. Repeat step 21 for each of the remaining fields of the Gateway.
23. Press the select button when you have changed the last field of Gateway.
24. Press and release the right navigation button until Create Now? is displayed.
25. If you are satisfied with your settings, press the select button.
 - If you want to review your settings before you create the cluster, use the right and left navigation buttons to review those settings. Make any necessary changes, return to Create Now?, then press the select button.
 - If the cluster is created successfully, Password: is displayed on line 1 of the service display screen. Line 2 contains a password that you can use to access the cluster. Record this password now. The password is displayed for 60 seconds, or until the up, down, left or right navigation button is pressed, which deletes it.

Important: If you do not record the password, you must restart the cluster configuration procedure. When the password has been recorded, press the up, down, left, or right navigation button to delete the password from the screen.

After you complete this task, the following information is displayed on the service display screen:

- Cluster: is displayed on line 1.
- The cluster IP address is displayed on line 2.

SAN Volume Controller Console layout

Ensure that you are familiar with the basic frame layout of the SAN Volume Controller Console.

Figure 23 provides, the basic frame layout, which consists of a banner, task bar, portfolio and a work area. An optional frame can be added for embedded task assistance or help.

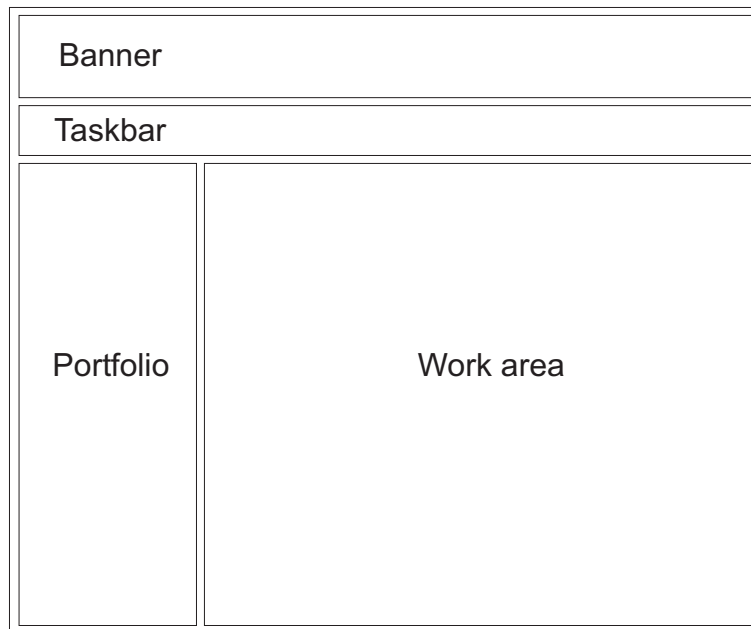


Figure 23. Basic frame layout

Related reference

“SAN Volume Controller Console banner”

The banner of the SAN Volume Controller Console is used for product or customer identification.

“SAN Volume Controller Console task bar” on page 83

The task bar of the SAN Volume Controller Console keeps track of all opened primary tasks and allows you to quickly go back to the previous task or move forward to the next task.

“SAN Volume Controller Console portfolio” on page 83

The portfolio area of the SAN Volume Controller Console contains task-based links that open panels in the work area. Common tasks are grouped under task headings and are expandable and collapsible.

“SAN Volume Controller Console work area” on page 84

The work area of the SAN Volume Controller Console is where you work with a cluster and the objects it contains.

SAN Volume Controller Console banner

The banner of the SAN Volume Controller Console is used for product or customer identification.

Figure 24 shows the banner.



Figure 24. Banner

SAN Volume Controller Console task bar

The task bar of the SAN Volume Controller Console keeps track of all opened primary tasks and allows you to quickly go back to the previous task or move forward to the next task.

Figure 25 shows the task bar. You can click the **question mark (?)** icon on the right side to display the information center in a separate browser window. You can click the **(I)** icon to display a help topic for the panel that is currently displayed in the work area.



Figure 25. Task bar

SAN Volume Controller Console portfolio

The portfolio area of the SAN Volume Controller Console contains task-based links that open panels in the work area. Common tasks are grouped under task headings and are expandable and collapsible.

The following task-based links are available from the Welcome panel of the SAN Volume Controller Console:

- Welcome
- Clusters
- Users
- Truststore Certificate
- Change Password

The following task-based links are available after you have launched the SAN Volume Controller application:

- Welcome
- Manage Cluster
 - View Cluster Properties
 - Maintain Passwords
 - Modify IP Addresses
 - Set Cluster Time
 - Start Statistics Collection
 - Stop Statistics Collection
 - Shut Down Cluster
- Work with Nodes
 - I/O Groups
 - Nodes
- Manage Progress
 - View Progress
- Work with Managed Disks
 - Disk Controller Systems
 - Managed Disks
 - Managed Disk Groups

- Work with Virtual Disks
 - Virtual Disks
 - Virtual Disk-to-Host Mappings
- Work with Hosts
 - Hosts
- Manage Copy Services
 - FlashCopy Mappings
 - FlashCopy Consistency Groups
 - Metro Mirror Relationships
 - Metro Mirror Consistency Groups
 - Metro Mirror Cluster Partnership
- Service and Maintenance
 - Upgrade Software
 - Run Maintenance Procedures
 - Set Error Notification
 - Analyze Error Log
 - Set Features
 - View Feature Log
 - Dump Configuration
 - List Dumps
 - Backup Configuration
 - Delete Backup

SAN Volume Controller Console work area

The work area of the SAN Volume Controller Console is where you work with a cluster and the objects it contains.

The work area is the main area of the application.

Browser requirements for the SAN Volume Controller

Ensure that you are familiar with the Internet browsers and settings when using the SAN Volume Controller.

The following Web browser versions are required.

- Windows and UNIX operating systems
 - Netscape version 6.2
 - 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 displays.
 2. From the left side category, click **Advanced** to expand the suboptions. The suboption Proxies displays.
 3. Click **Proxies**. The Proxies window displays.
 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.

Configuring the Web browser

You must configure the Web browser to allow new windows to automatically open. These new windows are called popups.

- Ensure that the Web browser is not set to block or suppress popup windows.
- Do not install applications on the Web browser that block or suppress popup windows. If such an application is installed with the Web browser, uninstall it or turn it off.

Changing browser settings for password protection

For security reasons, you can configure your Web browser so that a password does not automatically display when you type a user name into the user name field.

1. For Internet Explorer 6.0, complete the following steps to change your browser settings:
 - a. Start an Internet Explorer session.
 - b. Click **Tools** → **Internet Options** from the menu bar. The Internet Options panel is displayed.
 - c. Click the **Content** tab.
 - d. Click **AutoComplete**. The AutoComplete Settings panel is displayed.
 - e. Ensure that the **User names and passwords on forms** box is unchecked.
 - f. Click **OK**.
2. For Netscape, complete the following steps to change your browser settings:
 - a. Start a Netscape session.
 - b. Click **Edit** → **Preferences** from the menu bar.
 - c. Click **Privacy and Security**.
 - d. Click **Web Passwords**.
 - e. Ensure that the **Remember passwords for sites that require me to log in** box is unchecked.
 - f. Click **OK**.

Accessing the SAN Volume Controller Console

The SAN Volume Controller Console 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 because this can prevent the windows in the SAN Volume Controller Console from opening.

You access the SAN Volume Controller Console by pointing a Web browser to the following URL on the master console:

`http://svccconsoleip:9080/ica`

Where `svccconsoleip` is the IP address of your master console.

Log on to the SAN Volume Controller Console using the superuser user name, which is `superuser`, and the superuser password, which is `passwd`. The first time you access the SAN Volume Controller Console you are required to change the superuser password

You can use the SAN Volume Controller Console to configure the SAN Volume Controller clusters in your environment.

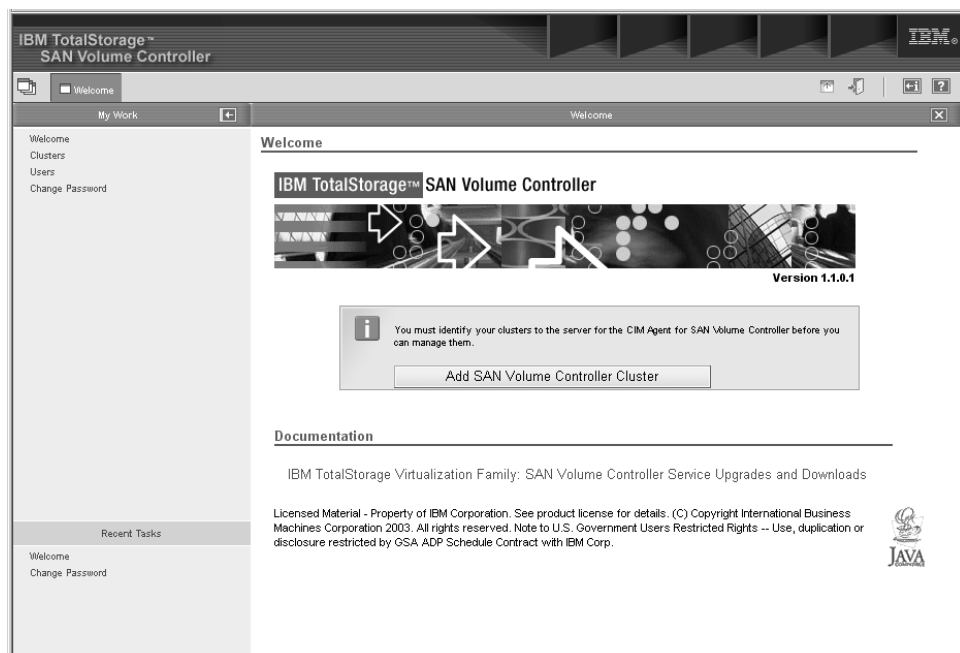
Creating a cluster using the SAN Volume Controller Console

After you have created a pair of nodes, you can create and configure a cluster.

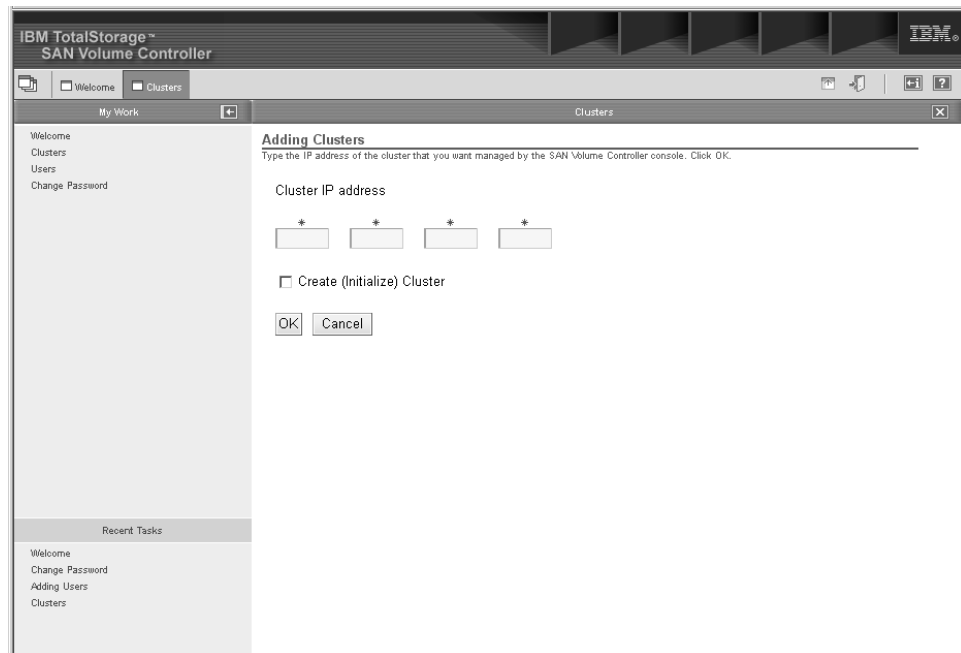
You must generate an SSH key pair before you can use the SAN Volume Controller to create a cluster. If you are adding an SSH public key to enable your system to use the command-line interface (CLI), you must also generate an SSH key pair for the CLI.

Perform the following steps to create a cluster:

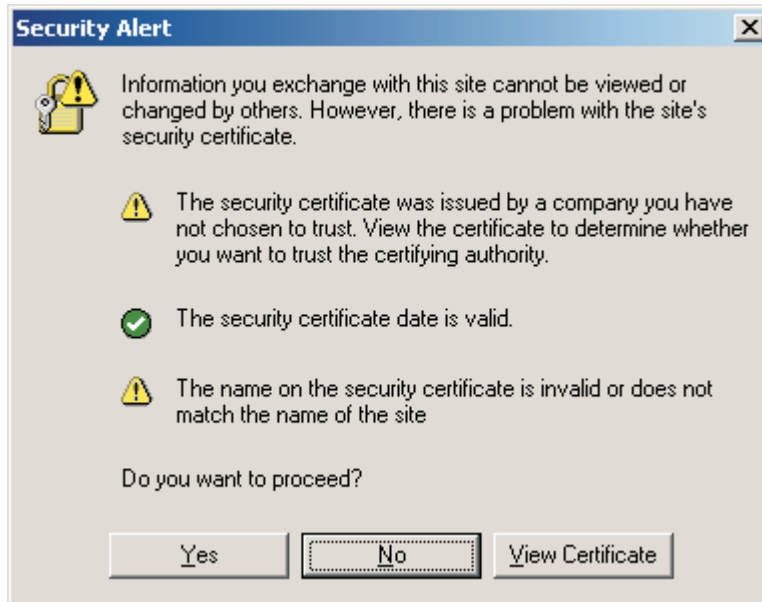
1. Start the SAN Volume Controller Console by clicking on the desktop icon or by pointing your Web browser to `http://<svccconsoleip>:9080/ica`, where `<svccconsoleip>` is the IP address of the master console. The Enter Network Password window is displayed.
2. Type `superuser` for the user ID and `passwd` 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. If this is the first time that you have accessed the SAN Volume Controller Console, go to step 4. Otherwise, go to step 7.
4. Click **Add SAN Volume Controller Cluster** from the Welcome panel.
5. Click **Create new cluster**. The SAN Volume Controller creates the new cluster. When the new administrator password is accepted, the cluster displays the password prompt again.
6. Type the user ID admin and the new administrator password.
7. Select **Clusters** from the portfolio. The Viewing Clusters panel is displayed.
8. From the task list, select **Add a Cluster** and click **Go**. The Adding Clusters panel is displayed.



9. Type the IP address of your cluster.
 - If the cluster has not been fully created (that is, you have created the cluster using the front panel), select the **Create (Initialize) Cluster** check box.
 - If the cluster is already in use (that is you are just adding this cluster to the list of managed clusters for this installation of the SAN Volume Controller Console) do not select the Create (Initialize) Cluster check box. Click **OK**. The Security Alert panel is displayed.



10. Click **View Certificate**. The Certificate panel is displayed.



- a. Click **Install Certificate**.
- b. Click **Next**.
- c. Click **Next**.

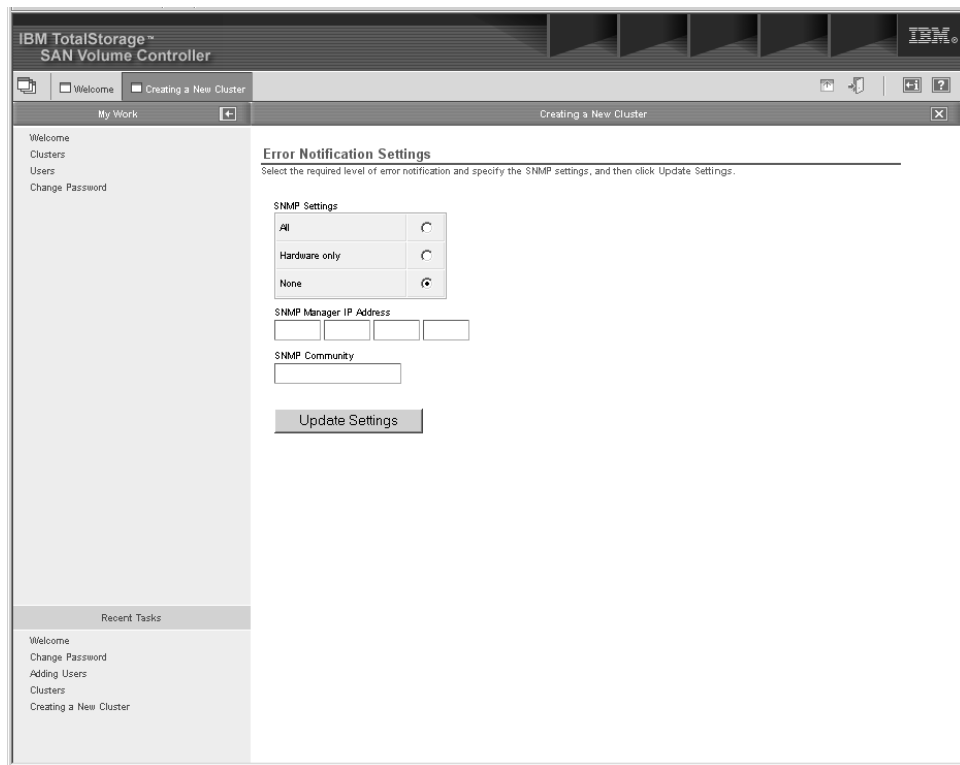
- d. Click **Install**.
 - e. Click **OK** to complete the Install Certificate panel.
 - f. Click **OK** to close the Certificate panel.
 - g. Click **Yes** to close the Security Alert panel.
11. Type the cluster user name admin and the password that was generated when you created the cluster from the front panel.
 12. Click **OK**.
 13. Click **Continue** when the Create a Cluster wizard begins. The Create New Cluster panel is displayed. If the cluster already existed and you did not select the **Initialize Cluster** check box in step 9 on page 87, proceed to step 17 on page 90.

14. Complete the Create New Cluster panel.
 - a. Type a new administrator password.

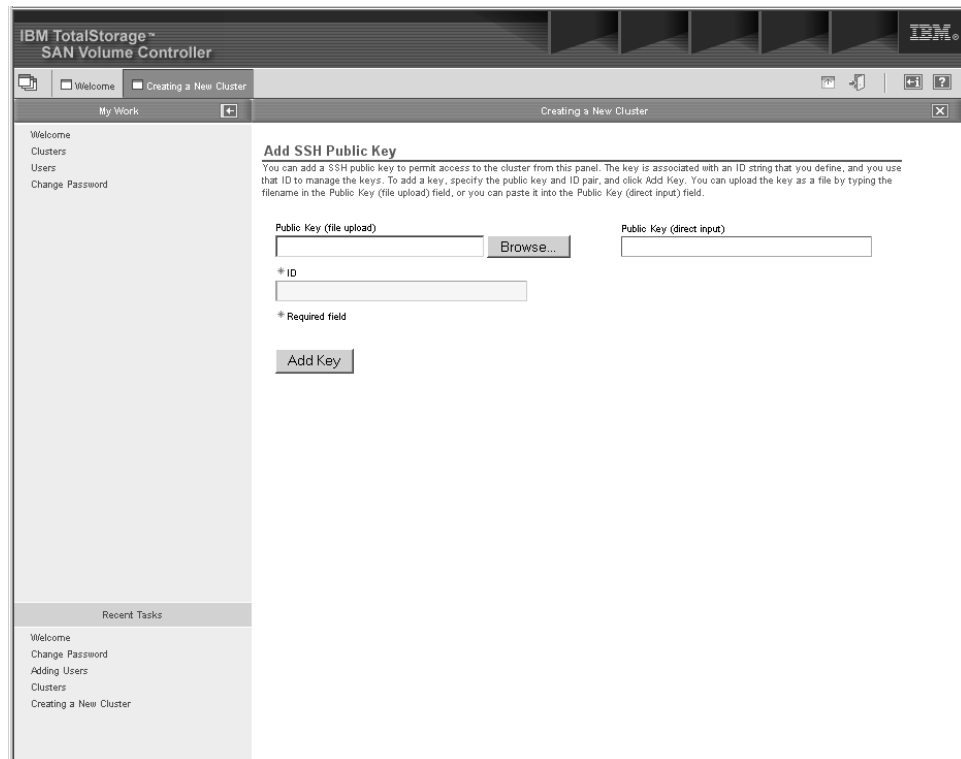
Important: Record this password because you will need it to upload new SSH Keys using the SAN Volume Controller Console.
 - b. Type the service password.

Important: Record this password because you will need it to upload new SSH Keys using the SAN Volume Controller Console.
 - c. Type a name for your cluster.
 - d. Type the service IP address for the cluster. This is the IP address that the system uses if you have to start a single node in service mode.
 - e. Select the speed of your fabric, either 1 or 2 Gbps.
 - f. If you want to reset the administrator password from the front panel, select the **Administrator Password Policy** check box.

- g. Click **Create New Cluster** when you have completed this panel. After a few seconds, the cluster is created.
 - h. Click **Continue** when the Web page returns.
15. Click **Continue** after you are notified that the password has been changed. The Error Notification Settings panel is displayed.



- a. If you want errors forwarded as SNMP traps, select either **Hardware only** or **All**. Selecting *Hardware only* sends SNMP traps for hardware-related errors and selecting *All* sends SNMP traps for both hardware and software errors.
 - b. Type the IP address of the system that is running your SNMP management software.
- Note:** If you are using IBM Director on the master console to collect SNMP traps, type the IP address of the master console here.
- c. Type the SNMP community name.
 - d. Click **Update Settings** to continue.
16. Click **Continue**. The Featurization Settings panel is displayed. The allowed setting for each of the parameters is specified in your user's license.
- a. Enable the FlashCopy or Metro Mirror options if they are licensed.
 - b. Type the virtualization limit as specified in the license. A zero value is not allowed for this field.
 - c. Click **Set features**.
17. Click **Continue**. The Add SSH Public Key panel is displayed.



- a. If prompted, type admin as the user name and type the new password you gave during step 14 on page 89.
 - b. Click **Browse** to locate the public key for the master console.
 - c. Type an ID (label) for this key.
 - d. Click **Add Key**.
18. Click on the **X** that is located in the right corner of the window to close the wizard.

You have successfully connected and configured the cluster. The cluster should be listed on the Viewing Clusters panel.

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

Related concepts

“Browser requirements for the SAN Volume Controller” on page 84
Ensure that you are familiar with the Internet browsers and settings when using the SAN Volume Controller.

Related tasks

“Creating a cluster from the front panel” on page 79
After you have created a pair of nodes, you can use the front panel of the SAN Volume Controller to create a cluster.

“Configuring the Web browser” on page 85
You must configure the Web browser to allow new windows to automatically open. These new windows are called popups.

“Changing browser settings for password protection” on page 85
For security reasons, you can configure your Web browser so that a password does not automatically display when you type a user name into the user name field.

Related reference

“SAN Volume Controller Console layout” on page 81

Ensure that you are familiar with the basic frame layout of the SAN Volume Controller Console.

“SAN Volume Controller Console banner” on page 82

The banner of the SAN Volume Controller Console is used for product or customer identification.

“SAN Volume Controller Console task bar” on page 83

The task bar of the SAN Volume Controller Console keeps track of all opened primary tasks and allows you to quickly go back to the previous task or move forward to the next task.

“SAN Volume Controller Console portfolio” on page 83

The portfolio area of the SAN Volume Controller Console contains task-based links that open panels in the work area. Common tasks are grouped under task headings and are expandable and collapsible.

“SAN Volume Controller Console work area” on page 84

The work area of the SAN Volume Controller Console is where you work with a cluster and the objects it contains.

“Accessing the SAN Volume Controller Console” on page 85

The SAN Volume Controller Console is a Web-based application that you can use to manage multiple clusters.

Chapter 3. Using the SAN Volume Controller Console

The SAN Volume Controller is provided with a console that is Web-browser based. It is known as the SAN Volume Controller Console.

Overview

The SAN Volume Controller Console can be used to create and maintain the configuration of storage associated with the SAN Volume Controller. It also provides user management and access to multiple clusters.

You can use the SAN Volume Controller Console to perform the following functions:

- Initial set up of the cluster, its nodes, and the I/O groups (or node pairs). This function includes diagnostics and error log analysis of the cluster.
- Set up and maintain managed disks and managed disk groups.
- Set up and maintain Secure Shell keys.
- Set up and maintain virtual disks.
- Set up logical host objects.
- Map virtual disks to hosts.
- Navigate from managed hosts to virtual disk and to managed disk groups, and the reverse direction up the chain.
- Set up and start Copy Services:
 - FlashCopy and FlashCopy consistency groups.
 - Synchronous Metro Mirror and Metro Mirror consistency groups.

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

Launching the SAN Volume Controller Console application

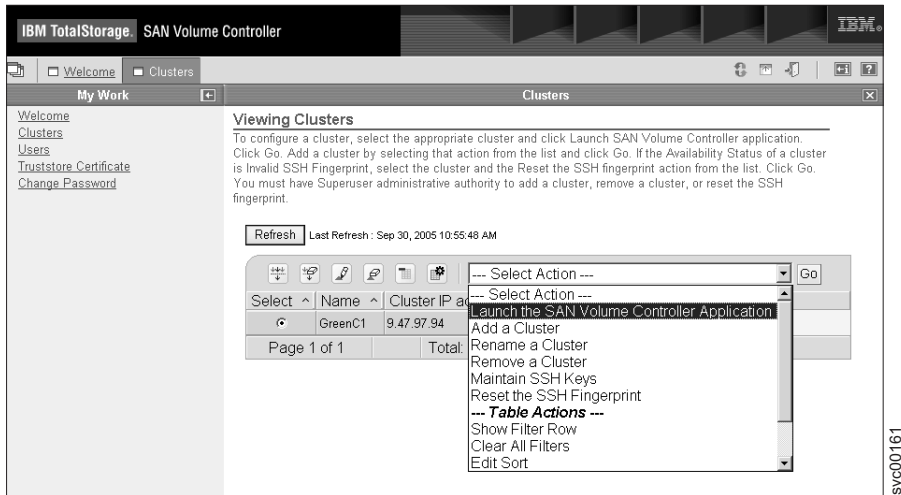
You can launch the SAN Volume Controller from the Viewing Clusters panel.

The SAN Volume Controller Console is the centralized Web application that is used to manage your clusters. It is preinstalled on the master console.

This task assumes that you are at the Welcome panel for the SAN Volume Controller Console.

Perform the following steps to launch the SAN Volume Controller application:

1. Click **Clusters** in 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 task list.



4. Click **Go**. A secondary browser window opens.

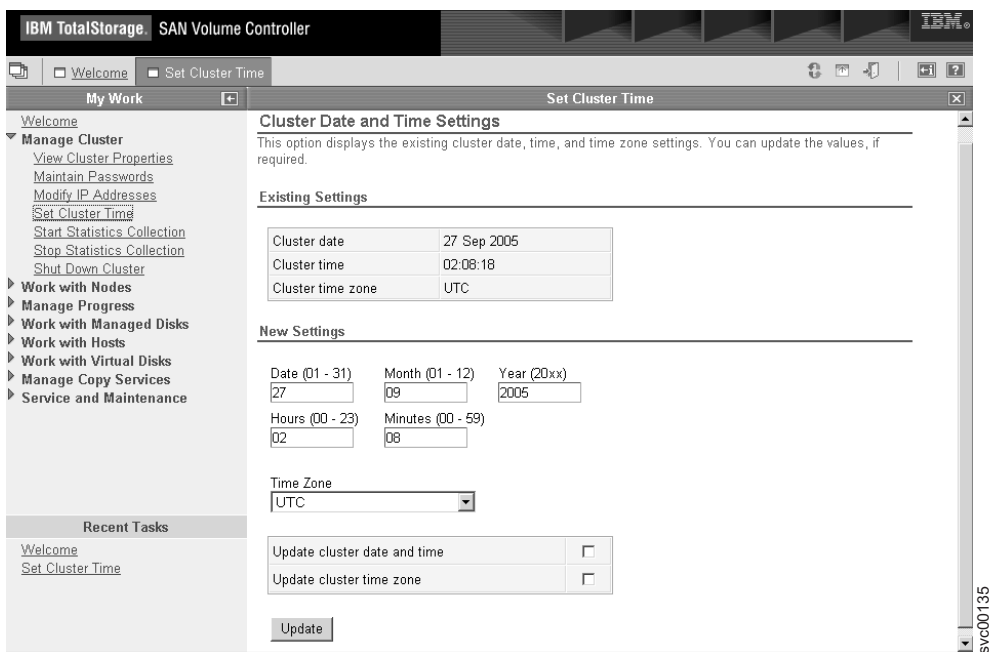
Setting the cluster time

You can set the cluster time for the SAN Volume Controller from the Cluster Date and Time Settings panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to set the cluster time:

1. Click **Manage Clusters** → **Set Cluster Time** in the portfolio. The Cluster Date and Time Settings panel is displayed.



2. Type your changes into the **Date**, **Month**, **Year**, **Hours** and **Minutes** fields and select a new time zone from the **Time Zone** list.

3. Select **Update cluster time and date**, **Update cluster time zone**, or both.
4. Click **Update** to submit the update request to the cluster.

Modifying the cluster IP addresses

You can display and change the IP addresses that are associated with a cluster from the Modify IP Addresses panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to change the IP addresses:

1. Click **Manage Cluster** → **Modify IP Address** in the portfolio. The Modify IP Addresses panel is displayed. The Modify IP Addresses panel displays the existing value for the following IP addresses and enables you to change the settings:
 - Cluster IP Address
 - Service IP Address (used when the node is not part of the cluster)
 - Subnet Mask IP Address
 - Gateway IP Address
 - Master Console IP Address
 - Master Console Port
2. 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.
3. Click **Modify Settings** to update the IP address. 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 Web browser connection.

A new SSL certificate is generated by the cluster to display the new IP address. This new certificate displays when the Web browser first connects to the cluster.

Maintaining cluster passwords

You can use the SAN Volume Controller Console to maintain cluster passwords.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to maintain cluster passwords:

1. Click **Manage Cluster** → **Maintain Passwords** in the portfolio. The Maintain Passwords panel is displayed.
2. Type the new administrator or service password in the appropriate fields and click **Maintain Passwords** to change the password.

Note: Passwords must be typed twice to allow verification. Passwords can consist of A - Z, a - z, 0 - 9, and underscore.

3. If you are changing an administrator password, you must reauthenticate the administrator password by entering the new administrator password in the password prompt.
4. Record the administrator password because you cannot access the cluster through the SAN Volume Controller Console without this password.

Viewing cluster properties

You can view the properties for a cluster from the View Cluster Properties panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to view the properties of a cluster:

1. Click **Manage Cluster** → **View Cluster Properties** in the portfolio. The Viewing General Properties panel is displayed.
2. Click the following tabs:
 - a. **General** to display the general properties.
 - b. **IP Addresses** to view the cluster IP address, service IP address, subnet mask address, default gateway address and master console IP address.
 - c. **Space** to view the space and capacity for managed disks (MDisks), MDisk groups and virtual disks (VDisks).
 - d. **SNMP** to view the SNMP details.
 - e. **Statistics** to view the cluster statistics details.
 - f. **Metro Mirror** to view the Metro Mirror properties of the cluster.
3. Click **Close** to close the panel.

Adding nodes to a cluster

For availability purposes, you must connect the nodes in an I/O group to different uninterruptible power supplies (UPSs).

Before you add a node to a cluster, you must 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. If you are replacing a node and the switch is zoned by worldwide port name (WWPN) rather than by switch port, make sure that the switch is configured such that the node being added is in the same VSAN/zone.

Special procedures when adding a node to a cluster

Applications on the host systems direct I/O operations to file systems or logical volumes that are mapped by the operating system to virtual paths (vpaths), which are pseudo disk objects that are supported by the Subsystem Device Driver (SDD). SDD maintains an association between a vpath and a SAN Volume Controller virtual disk (VDisk). This association uses an identifier (UID) which is unique to the VDisk and is never reused. The UID allows SDD to directly associate vpaths with VDIsks.

SDD operates within a protocol stack that contains disk and fibre channel device drivers that allow it to communicate with the SAN Volume Controller using the SCSI protocol over fibre channel as defined by the ANSI FCS standard. The addressing scheme that is provided by these SCSI and fibre-channel device drivers uses a combination of a SCSI logical unit number (LUN) and the worldwide name (WWN) for the fibre-channel node and ports.

If an error occurs, the 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 that were previously used.

SDD does not check the association of the VDisk with the vpath on every I/O operation that it performs.

Before you add a node to the cluster, you must check to see if any of the following conditions are true:

- The cluster has more than one I/O group.
- The node that is being added to the cluster uses physical node hardware or a slot which has previously been used for a node in the cluster.
- The node that is 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 back-end storage.

If any of the previous conditions are true, the following special procedures apply:

- The node must be added to the same I/O group that it was previously in. You can use the command-line interface (CLI) command **svcinfo lsnode** or the SAN Volume Controller Console to determine the WWN of the cluster nodes.
- Before you add the node back into the cluster, you must shut down all of the hosts that are using the cluster. The node must then be added before the hosts are rebooted. If the I/O group information is unavailable or it is inconvenient to shut down and reboot all of the hosts that are using the cluster, perform the following actions:
 - On all hosts that are connected to the cluster, unconfigure the fibre-channel adapter device driver, the disk device driver and SDD before you add the node to the cluster.
 - Add the node to the cluster and then reconfigure the fibre-channel adapter device driver, the disk device driver, and SDD.

Scenarios where the special procedures can apply

The following two scenarios describe situations where the special procedures can apply:

- Four nodes of an eight-node cluster have been lost because of the failure of a pair of 2145 uninterruptible power supply (2145 UPS) or four 2145 uninterruptible power supply-1U (2145 UPS-1U). In this case, the four nodes must be added back into the cluster using the CLI command **svctask addnode** or the SAN Volume Controller Console.
- You decided to delete four nodes from the cluster and add them back into the cluster using the CLI command **svctask addnode** or the SAN Volume Controller Console.

Adding nodes to a cluster using the SAN Volume Controller Console

This task assumes that you have already launched the SAN Volume Controller Console.

Attention: If you are re-adding a node to the SAN, ensure that you are adding the node to the same I/O group from which it was removed. Failure to do this can result in data corruption. You must 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. The LUNs that are presented to the ports on the new node must be the same as the LUNs that are presented to the nodes that currently exist in the cluster. You must make sure that the LUNs are the same before you add the new node to the cluster.

Each node in an I/O group must be connected to a different uninterruptible power supply. If you do not provide 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.

This task assumes that you have already launched the SAN Volume Controller application.

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

1. Click **Work with Nodes** → **Nodes** in the portfolio. The Viewing Nodes panel is displayed.
2. Select **Add a node** from the task list and click **Go**. The Adding a Node to a Cluster panel is displayed.
3. If you are adding a node into the cluster for the first time, record the following information:
 - Node serial number
 - All WWPNs
 - I/O groups that the node belongs to

Important: You need this information to avoid possible data corruption if you have to remove and re-add the node to the cluster.

4. Select the node that you want to add to the cluster from the **Available Candidate Nodes** list.
5. Select the I/O group from the **I/O Groups** list.
6. In the **Node Name** field, type the name that you want to assign to the node.
7. Click **OK**.

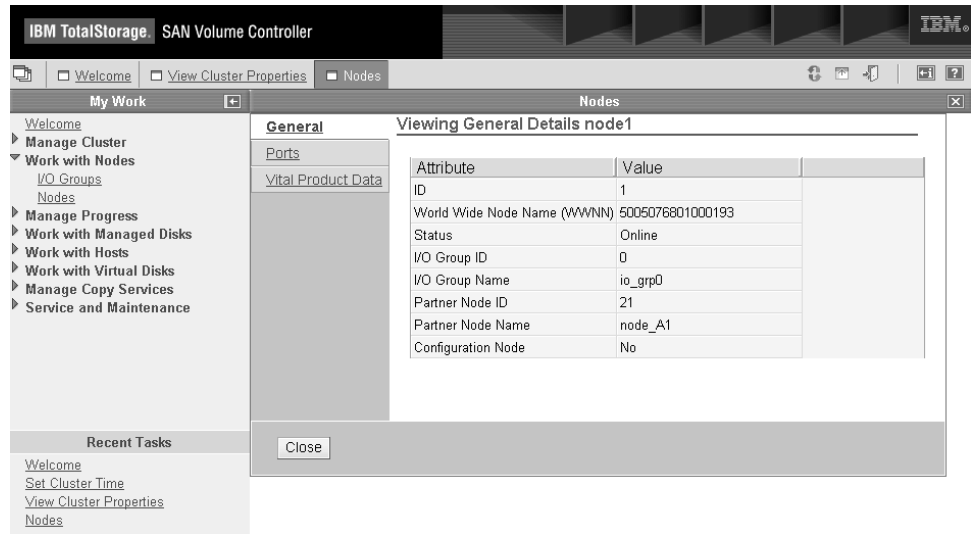
Viewing node properties

You can view the properties for a node using the Viewing General Details panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to view the node properties:

1. Click **Work with Nodes** → **Nodes** in the portfolio. The Viewing Nodes panel is displayed.
2. Click the name of the node for which you want to view detailed information. The Viewing General Details panel is displayed.



svc00138

3. Click **Ports** to view the worldwide port name (WWPN) details. The Viewing Port Details panel is displayed.
4. Click **Vital Product Data** to view the node hardware details. The Viewing Vital Product Data panel is displayed.
5. Click **Close** to close the panel.

Increasing the size of a cluster

You can use the SAN Volume Controller Console to increase the size of a cluster.

You can increase throughput by adding more nodes to the cluster. The nodes must be added in pairs and assigned to a new I/O group.

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

1. Add a node to your cluster and repeat this step for the second node.
2. If you want to balance the load between the existing I/O groups and the new I/O groups, you can migrate your virtual disks (VDisks) to new I/O groups. Repeat this step for all VDisks that you want to assign to the new I/O group.

Adding a node to increase the size of a cluster

You can use the SAN Volume Controller Console to add a node to a cluster.

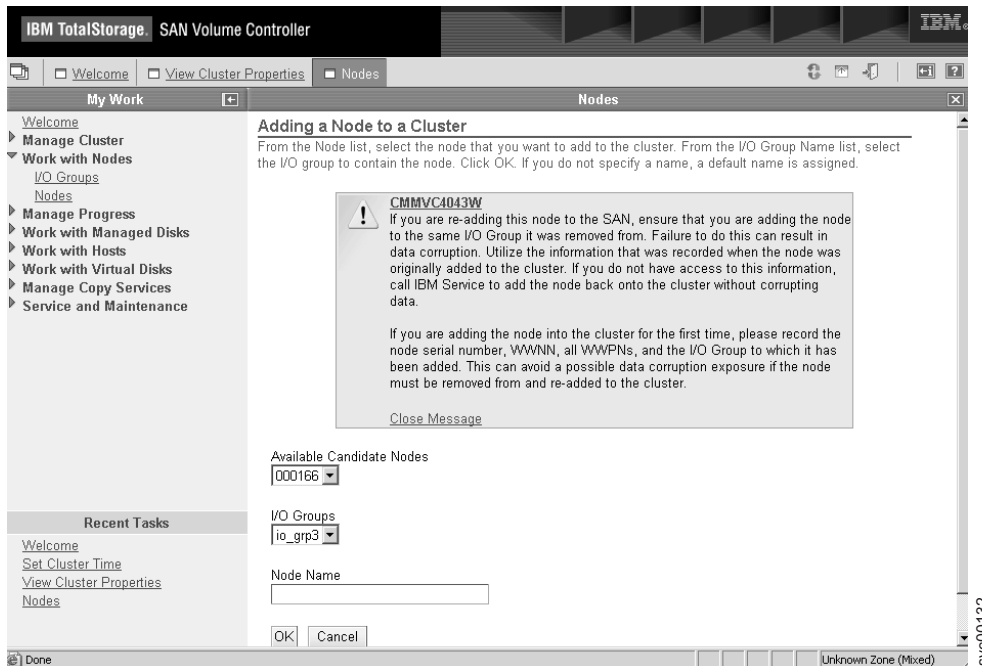
Attention: If you are adding a node that was previously removed from a cluster, ensure that you are adding the node to the same I/O group from which it was removed. Failure to do this can result in data corruption. If you do not know the I/O group name or ID that it was removed from, contact IBM Service to add the node to the cluster without corrupting data.

If you want to add a node that was previously removed from a cluster, you must have the following information about the node:

- Node serial number
- Worldwide node name (WWNN)
- All of the worldwide port names (WWPN)
- The name or ID of the I/O group from which the node was previously removed

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

1. Click **Work with Nodes** → **I/O groups** to determine the I/O group where the node will be added. The Viewing Input/Output Groups panel is displayed.
2. Record the name or ID of the first I/O group that has a node count of zero (0).
3. Click **Work with Nodes** → **Nodes**. The Viewing Nodes panel is displayed.
4. Select the node that you want to add from the list of available candidate nodes.
5. Select **Add a Node** from the task list and click **Go**. The Adding a Node to a Cluster panel is displayed.



6. Select the node that you want to add to the cluster from the **Available Candidate Nodes** list.
7. Select the I/O group from the **I/O Groups** list.

Important: If you are adding a node that was previously removed from the cluster, you must select the name of the I/O group from which the node was previously removed. If you are adding a node that has never been in a cluster, select the name of the I/O group that you recorded in step 2.

8. Click **OK**.
9. Verify that the node is online by refreshing the Viewing Nodes panel. You might have to close the panel and reopen it to refresh the panel.
10. Click the name of the node that you have added to the cluster. The Viewing General Details panel is displayed.
11. Click the **General**, **Ports** and **Vital Product Data** tabs and record the following information:
 - Node serial number
 - WWNN
 - WWPN
 - The name or ID of the I/O group that the node belongs to

12. Click **Close** to close the panel.

If the disk controller uses mapping to present RAID arrays or partitions to the cluster and the WWNNs or the WWPNNs have changed, you must modify the port groups that belong to the cluster.

Migrating a VDisk to a new I/O group

You can migrate a virtual disk (VDisk) to a new I/O group to manually balance the workload across the nodes in the cluster.

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

Perform the following steps to migrate a single VDisk:

1. Quiesce all I/O operations for the VDisk. You might have to determine the hosts that are using this VDisk.
2. Update the multipathing device driver configuration to remove all device identifiers that are presented by the VDisk you intend to move. If you are using the subsystem device driver (SDD), the device identifiers are referred to as virtual paths (vpaths).

Attention: Failure to perform this step can result in data corruption.
3. Stop and delete all FlashCopy mappings or Metro Mirror relationships that use this VDisk. To check if the VDisk is part of a mapping or relationship, perform the following steps:
 - a. Click **Work with Virtual Disks** → **Virtual Disks** in the portfolio. The Filtering Virtual Disks panel is displayed.
 - b. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Virtual Disks panel is displayed.
 - c. Click on the name of the VDisk that you want to migrate. The View VDisk Details panel is displayed.
 - d. Look for the **FlashCopy ID** and **Metro Mirror ID** fields. If these fields are not blank, the VDisk is part of a mapping or relationship.
 - e. Click **Close** to close the panel.
4. Migrate the VDisk by selecting the VDisk from the Viewing Virtual Disks panel and select **Modify a VDisk** from the task list and click **Go**. The Modifying Virtual Disk panel is displayed.
5. Select the new I/O group from the **I/O Group** list and click **OK**.
6. Follow your multipathing device drivers instructions for discovering new device identifiers. For example, if you are using SDD, see the *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide* and follow the instructions for discovering vpaths.

Replacing a faulty node with a spare node

You can use the SAN Volume Controller Console and the SAN Volume Controller front panel to replace a faulty node in a cluster.

Before you attempt to replace a faulty node with a spare node you must ensure that you meet the following requirements:

- SAN Volume Controller version 1.1.1 or later is installed on the cluster and on the spare node.
- You know the name of the cluster that contains the faulty node.
- A spare node is installed in the same rack as the cluster that contains the faulty node.
- You make a record of the last five characters of the original worldwide node name (WWNN) of the spare node. You will need this information, if and when, you want to stop using this node as a spare node.

If a node fails, the cluster continues to operate with degraded performance, until the faulty node is repaired. If the repair operation takes an unacceptable amount of time, it is useful to replace the faulty node with a spare node. However, the appropriate procedures must be followed and precautions must be taken so you do *not* interrupt I/O operations and compromise the integrity of your data.

The following table describes the changes that are made to your configuration when you replace a faulty node in the cluster:

Node attributes	Description
Front panel ID	This is the number that is printed on the front of the node and is used to select the node that is added to a cluster. This number can change.
Node ID	This is the ID that is assigned to the node. A new node ID is assigned each time a node is added to a cluster; the node name remains the same following service activity on the cluster. You can use the node ID or the node name to perform management tasks on the cluster. However, if you are using scripts to perform those tasks, use the node name rather than the node ID. This ID can change.
Node name	This is the name that is assigned to the node. If you do not specify a name, the SAN Volume Controller assigns a default name. The SAN Volume Controller creates a new default name each time a node is added to a cluster. If you choose to assign your own names, you must type the node name on the Adding a node to a cluster panel. If you are using scripts to perform management tasks on the cluster and those scripts use the node name, you can avoid the need to make changes to the scripts by assigning the original name of the node to a spare node. This name can change.
Worldwide node name	This is the WWNN that is assigned to the node. The WWNN is used to uniquely identify the node and the fibre-channel ports. The WWNN of the spare node changes to that of the faulty node. The node replacement procedures must be followed exactly to avoid any duplication of WWNNs. This name can change.

Node attributes	Description												
Worldwide port names	<p>These are the WWPNs that are assigned to the node. WWPNs are derived from the WWNN that is written to the spare node as part of this procedure. For example, if the WWNN for a node is 50050768010000F6, the four WWPNs for this node are derived as follows:</p> <table data-bbox="740 373 1458 535"> <tr> <td>WWNN</td> <td>50050768010000F6</td> </tr> <tr> <td>WWNN displayed on front panel</td> <td>000F6</td> </tr> <tr> <td>WWPN Port 1</td> <td>50050768014000F6</td> </tr> <tr> <td>WWPN Port 2</td> <td>50050768013000F6</td> </tr> <tr> <td>WWPN Port 3</td> <td>50050768011000F6</td> </tr> <tr> <td>WWPN Port 4</td> <td>50050768012000F6</td> </tr> </table> <p>These names do not change.</p>	WWNN	50050768010000F6	WWNN displayed on front panel	000F6	WWPN Port 1	50050768014000F6	WWPN Port 2	50050768013000F6	WWPN Port 3	50050768011000F6	WWPN Port 4	50050768012000F6
WWNN	50050768010000F6												
WWNN displayed on front panel	000F6												
WWPN Port 1	50050768014000F6												
WWPN Port 2	50050768013000F6												
WWPN Port 3	50050768011000F6												
WWPN Port 4	50050768012000F6												

This task assumes that you have already launched the SAN Volume Controller application.

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

1. Verify the name and ID of the node that you want to replace.

Perform the following steps to verify the name and ID:

 - a. Make sure that the SAN Volume Controller Console application is running on the cluster that contains the faulty node.
 - b. Click **Work with Nodes** → **Nodes** in the portfolio. The Viewing Nodes panel is displayed. If the node is faulty, it is shown as offline.
 - c. Ensure the partner node in the I/O group is online.
 - If the other node in the I/O group is offline, start the Directed Maintenance Procedures (DMPs) to determine the fault.
 - If you have been directed here by the DMPs, and subsequently the partner node in the I/O group has failed, recover the offline VDisks.
 - If you are replacing the node for other reasons, determine the node that you want to replace and ensure that the partner node in the I/O group is online.
 - If the partner node is offline, you will lose access to the VDisks that belong to this I/O group. Start the DMPs and fix the other node before proceeding to the next step.
2. Click the name of the faulty (offline) node. The Viewing General Details panel is displayed.
3. Click the **General**, **Ports** and **Vital Product Data** tabs and record the following information:
 - Node serial number
 - Worldwide node name
 - All of the worldwide port names
 - Name or ID of the I/O group that contains the node
 - Front panel ID
 - UPS serial number
4. Disconnect all four fibre-channel cables from the node.

Important: Do not plug the fibre-channel cables into the spare node until the spare node is configured with the WWNN of the faulty node.

5. Connect the power and signal cables from the spare node to the uninterruptible power supply (UPS) that has the serial number you recorded in step 3 on page 103.

Note: The signal cable can be plugged into any vacant position on the top row of serial connectors on the UPS. If no spare serial connectors are available on the UPS, disconnect the cables from the faulty node.

6. Power on the spare node.
7. Display the node status on the service panel. See the *IBM TotalStorage SAN Volume Controller: Service Guide* for more information.
8. Change the WWNN of the spare node to match the WWNN of the faulty node by performing the following steps:
 - a. With the node status displayed on the front panel; press and hold the down button; press and release the select button; release the down button. WWNN is displayed on line 1 of the display. Line 2 of the display contains the last five characters of the WWNN.
 - b. With the WWNN displayed on the service panel; press and hold the down button; press and release the select button; release the down button. The display switches into edit mode.
 - c. Change the WWNN that is displayed to match the last five digits of the WWNN you recorded in step 3 on page 103.

Note: To edit the displayed number, use the up and down buttons to increase or decrease the displayed numbers. Use the left and right buttons to move between fields.

- d. When the five characters match the number that you recorded in step 3 on page 103, press the select button twice to accept the number.
9. Connect the four fibre-channel cables that you disconnected from the faulty node and connect them to the spare node.
 - If an Ethernet cable has not been connected to the spare node, disconnect the Ethernet cable from the faulty node and connect it to the spare node.
10. Remove the faulty node from the cluster.

Remember: You must record the following information to avoid data corruption when this node is re-added to the cluster:

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

11. Add the spare node to the cluster.
12. Use the tools that are provided with your multipathing device driver on the host systems to verify that all paths are now online. See the documentation that is provided with your multipathing device driver for more information. For example, if you are using the subsystem device driver (SDD), see the *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide* for instructions on how to use the SDD management tool on host systems.
13. Repair the faulty node.

Attention: When the faulty node is repaired, do not connect the fibre-channel cables to it. Connecting the cables might cause data corruption.

14. If you want to use the repaired node as a spare node, perform the following steps:

- a. Display the node status on the front panel display of the node. See the *IBM TotalStorage SAN Volume Controller: Service Guide* for more information.
- b. With the node status displayed on the front panel, press and hold the down button; press and release the select button; release the down button. WWNN is displayed on line 1 of the display. Line 2 of the display contains the last five characters of the WWNN.
- c. With the WWNN displayed on the service panel, press and hold the down button; press and release the select button; release the down button. The display switches into edit mode.
- d. Change the displayed number to 00000.

Note: To edit the displayed number, use the up and down buttons to increase or decrease the displayed numbers. Use the left and right buttons to move between fields.

- e. Press the select button twice to accept the number.
This node can now be used as a spare node.

Attention: Never connect a node with a WWNN of 00000 to the cluster. If this node is no longer required as a spare and is to be used for normal attachment to a cluster, you must change the WWNN to the number you recorded when a spare was created. Using any other number might cause data corruption.

Renaming a node

You can rename a node from the Renaming Node panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to rename a node:

1. Click **Work with Nodes** → **Nodes** in the portfolio. The Viewing Nodes panel is displayed.
2. Select the node you want to rename and select **Rename a Node** from the list. Click **Go**. The Renaming Node panel is displayed.
3. Type the new name of the node and click **OK**.

Deleting a node from a cluster

You might have to delete a node from a cluster if the node has failed and is being replaced with a new node or if the repair that has been performed has caused that node to be unrecognizable by the cluster.

This task assumes that you have already launched the SAN Volume Controller application.

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 if the partner node fails.
- When you delete a node, you remove all redundancy from the I/O group. As a result, new or existing failures can cause I/O errors on the hosts. The following failures can occur:
 - Host configuration errors
 - Zoning errors
 - Multipathing software configuration errors
- If you are deleting the last node in an I/O group and there are virtual disks (VDisks) assigned to the I/O group, you cannot delete the node from the cluster if the node is online. If the node is offline, you can delete the node.
- If you are deleting the last node in an I/O group and there are no VDisks assigned to the I/O group, the cluster is destroyed. You must back up or migrate all data that you want to save before you delete the node.

Perform the following steps to delete a node from a cluster:

1. Determine the VDisks that are still assigned to this I/O group:
 - a. Request a filtered view of VDisks where the filter attribute is the name of the I/O group.
 - b. Determine which hosts the VDisk is mapped to.
 - If you do not want to maintain access to these VDisks proceed to step 2.
 - If you are deleting the last node in the I/O group and some or all of these VDisks contain data that you want to maintain access to, you must migrate the VDisk to a new I/O group.
2. Power off the node that you want to remove, unless this is the last node in the cluster. This ensures that the multipathing device driver does not rediscover the paths that are manually removed before you issue the delete node request.

Attention:

- Deleting or shutting down the configuration node might cause the Secure Shell (SSH) command to hang. If this occurs, wait for the SSH command to end or stop the command and issue the **ping** command with the cluster IP address. When the **ping** command times out, you can start to issue commands.
 - If you power on the node that has been removed and it is still connected to the same fabric or zone, it attempts to rejoin the cluster. At this point the cluster tells the node to remove itself from the cluster and the node becomes a candidate for addition to this cluster or another cluster.
 - If you are adding this node into the cluster, ensure that you add it to the same I/O group that it was previously a member of. Failure to do so can result in data corruption.
3. Update the multipathing device driver configuration to remove all device identifiers that are presented by the VDisk you intend to move. If you are using the subsystem device driver (SDD), the device identifiers are referred to as virtual paths (vpaths).

Attention: Failure to perform this step can result in data corruption.
 4. Click **Work with Nodes** → **Nodes** in the portfolio. The Viewing Nodes panel is displayed.

Viewing Nodes

Click on a node to view its details, or select a node and an action from the list and click Go. Add a node to the cluster by selecting that action from the list and clicking Go.

Refresh Last Refresh : Aug 19, 2005 1:10:10 PM

Select ^	Name ^	Status ^	World Wide Node Name (WWNN) ^	I/O Group Name ^	Config Node ^
<input type="radio"/>	Alina_node	Online	5005076801000174	io_grp0	No
<input type="radio"/>	node1	Online	5005076801000193	io_grp0	No
<input type="radio"/>	node18	Online	5005076801000183	io_grp2	No
<input type="radio"/>	node_7	Online	5005076801000123	io_grp3A	No
<input type="radio"/>	nodek3	Online	5005076801000166	io_grp1	No
<input type="radio"/>	nodek6	Online	5005076801000117	io_grp3A	No
<input type="radio"/>	test1	Online	50050768010000C0	io_grp1	Yes

Page 1 of 1 Total: 7 Filtered: 7 Displayed: 7 Selected: 0

svc00139

5. Select the node that you want to delete and select **Delete a Node** from the task list. Click **Go**. The Deleting Node from Cluster panel is displayed.
6. Click **Yes** to delete the node.

Renaming an I/O group

You can rename an I/O group from the Viewing I/O Groups panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to rename an I/O group:

1. Click **Work with Nodes** → **I/O Groups** in the portfolio. The Viewing Input/Output Groups panel is displayed.
2. Select the I/O group that you want to rename and select **Rename an I/O Group** from the list. Click **Go**. The Renaming I/O Group panel is displayed.
3. Type the new name of the I/O Group in the **New Name** field.
4. Click **OK**.

Renaming a cluster

You can rename a cluster from the Renaming Cluster panel.

This task assumes that you are at the Welcome panel for the SAN Volume Controller Console.

Perform the following steps to rename a cluster:

1. Click **Clusters** in the portfolio. The Viewing Clusters panel is displayed.
2. Select the cluster that you want to rename and select **Rename a Cluster** from the task list. Click **Go**. The Renaming Cluster panel is displayed.
3. Type a new name for the cluster.
4. Click **OK**.

The cluster is renamed to the name that you selected.

Shutting down a cluster

You can shut down a SAN Volume Controller cluster from the Shutting Down cluster panel.

If you want to remove all input power to a cluster (for example, the machine room power must 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 uninterruptible power supply (UPS), the SAN Volume Controller nodes detect the loss of power and continue to run on battery power until all data that is held in memory is saved to the internal disk drive. This increases the time that is required to make the cluster operational when input power is restored and severely increases the time that is required to recover from an unexpected loss of power that might occur before the UPS batteries have fully recharged.

When input power is restored to the UPSs, they start to recharge. However, the SAN Volume Controller nodes do not permit any I/O activity to be performed to the virtual disks (VDisks) until the UPS is charged enough to enable all the data on the SAN Volume Controller nodes to be saved in the event of an unexpected power loss. This might take as long as three hours. Shutting down the cluster prior to removing input power to the UPS units prevents the battery power from being drained and makes it possible for I/O activity to resume as soon as input power is restored.

Before shutting down a cluster, quiesce all I/O operations that are destined for this cluster. Failure to do so can result in failed I/O operations being reported to your host operating systems.

Attention:

- If you are shutting down the entire cluster, you lose access to all VDisks that are provided by this cluster. Shutting down the cluster also shuts down all SAN Volume Controller nodes. This shutdown causes the hardened data to be dumped to the internal hard drive.
- Ensure that you have stopped all FlashCopy, Metro Mirror and data migration operations before you attempt a cluster shutdown. Also ensure that all asynchronous deletion operations have completed prior to a shutdown operation.

Begin the following process of quiescing all I/O to the cluster by stopping the applications on your hosts that are using the VDisks that are provided by the cluster.

1. Determine which hosts are using the VDisks that are provided by the cluster.
2. Repeat the previous step for all VDisks.

You can shut down a cluster by stopping I/O activity and either pressing the power buttons on the front of each SAN Volume Controller node or by issuing a shutdown command to the cluster.

Attention: You must press and hold the power button for one second to shutdown the SAN Volume Controller node.

When input power is restored, you must press the power button on the UPS units before you press the power buttons on the SAN Volume Controller nodes.

Perform the following steps to shut down a cluster:

1. Click **Manage Clusters** → **Shut down Cluster** in the portfolio. The Shutting Down cluster panel is displayed.
2. Click **Yes**.

Shutting down a node

You can shut down a SAN Volume Controller node from the Shutting Down Node panel.

If you are shutting down the last SAN Volume Controller node in an I/O group, quiesce all I/O operations that are destined for this SAN Volume Controller node. Failure to do so can result in failed I/O operations being reported to your host operating systems.

Attention: Ensure that you have stopped all FlashCopy, Metro Mirror and data migration operations before you attempt a SAN Volume Controller node shutdown. Also ensure that all asynchronous deletion operations have completed prior to a shutdown operation.

You can shut down a SAN Volume Controller node by stopping I/O activity and either pressing the power buttons on the front of each SAN Volume Controller node or by issuing a shutdown command.

Attention: You must press and hold the power button for one second to shutdown the SAN Volume Controller node.

When input power is restored, you must press the power button on the uninterruptible power supply units before you press the power button on the SAN Volume Controller node.

Perform the following steps to use the shutdown command to shut down a SAN Volume Controller node:

1. Click **Work with Nodes** → **Nodes** in the portfolio. The Viewing Nodes panel is displayed.
2. Select the node that you want to shut down.
3. Select **Shut Down a Node** from the task list and click **Go**. The Shutting Down Node panel is displayed.
4. Click **Yes**.

Discovering MDisks

You can have the cluster rescan the fibre-channel network. The rescan discovers any new managed disks (MDisks) that might have been added to the cluster and rebalances MDisk access across the available controller device ports.

Perform the following steps to discover MDisks:

1. Click **Work with Managed Disks** → **Managed Disks** in the portfolio. The Filtering Managed Disks panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Managed Disks panel is displayed.
3. Select **Discover MDisks** from the list and click **Go**. The Discovering Managed Disks panel is displayed. The newly discovered MDisks are displayed in a table on the Discovering Managed Disks panel.

4. Click **Close** to return to the Viewing Managed Disks panel.

Renaming MDisks

You can rename a managed disk (MDisk) from the Renaming Managed Disk panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to rename an MDisk:

1. Click **Work with Managed Disks** → **Managed Disks** in the portfolio. The Filtering Managed Disks panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing MDisks panel is displayed.
3. Select the MDisk you want to rename and select **Rename an MDisk** from the list. Click **Go**. The Renaming Managed Disk panel is displayed.
4. Type a new name for the MDisk.
5. Click **OK**.

Adding excluded MDisks to a cluster

You can add managed disks (MDisks) that have been excluded from the cluster back into the cluster from the Including Managed Disk panel.

You must fix the fabric-related problem that caused the MDisk to become excluded from the cluster before you can add the MDisk to the cluster.

This task assumes that you have already launched the SAN Volume Controller application.

The MDisk might have been excluded from the cluster because of multiple I/O failures that are caused by noisy links.

Perform the following steps to add an excluded MDisk to a cluster:

1. Click **Work with Managed Disks** → **Managed Disks** in the portfolio. The Filtering Managed Disks panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing MDisks panel is displayed.
3. Select the excluded MDisk that you want to add to the cluster and select **Include an MDisk** from the list. Click **Go**. The Including Managed Disk panel is displayed.
4. Follow the instructions that are displayed on the Including Managed Disk panel.

Setting quorum disks

You can set a managed disk (MDisk) as a quorum disk from the Setting a Quorum Disk panel.

This task assumes that you have already launched the SAN Volume Controller application.

Attention: You must set quorum disks on multiple controllers to avoid the possibility of losing all of the quorum disks with a single failure.

Perform the following steps to set an MDisk as a quorum disk:

1. Click **Work with Managed Disks** → **Managed Disks** in the portfolio. The Filtering Managed Disks panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Managed Disks panel is displayed.
3. Select the MDisk that you want to set as a quorum disk and select **Set a Quorum Disk** from the list. Click **Go**. The Setting a Quorum Disk panel is displayed.
4. Select a quorum index number from the **Quorum Index** list and click **OK**.

Determining the relationship between MDisks and VDIsks

You can use the SAN Volume Controller Console to determine the relationship between managed disks (MDisks) and virtual disks (VDIsks).

This task assumes that you have already launched the SAN Volume Controller application.

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

1. Click **Work with Managed Disks** → **Managed Disks** in the portfolio. The Filtering Managed Disks panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Managed Disks panel is displayed.
3. Select the MDisk that you want to view.
4. Select **Show VDIsks** from the task list and click **Go**. The Viewing Virtual Disks panel is displayed. This panel lists the VDIsks that use this MDisk.

Determining the relationship between MDisks and RAID arrays or LUNs

Each managed disk (MDisk) corresponds with a single RAID array, or a single partition on a given RAID array. Each RAID controller defines a LUN number for this disk. The LUN number and controller name or ID are needed to determine the relationship between MDisks and RAID arrays or partitions.

This task assumes that you have already launched the SAN Volume Controller application.

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

1. Click **Work with Managed Disks** → **Managed Disks** in the portfolio. The Filtering Managed Disks panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Managed Disks panel is displayed.
3. Click the name of the MDisk that you want to view. The Viewing Managed Disk (MDisk) Details panel is displayed.
4. Record the controller name and controller LUN number.
5. Click **Work with Managed Disks** → **Disk Controller Systems** in the portfolio.
6. Click the name of the controller that you recorded in step 4 to show the detailed view of the controller. The Viewing General Details panel is displayed.

7. Record the vendor ID, the product ID and worldwide node name (WWNN).
8. Use the vendor ID, the product ID and WWNN to determine which controller presents this MDisk.
9. From the native user interface for the controller that presents this MDisk, list the LUNs that the controller presents and match the LUN number with that noted in step 3 on page 111. This is the exact RAID array and partition that corresponds with the MDisk.

Displaying MDisk groups

You can display the managed disk (MDisk) group that an MDisk is a part of from the Viewing Managed Disk Groups panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to display the MDisk group:

1. Click **Work with Managed Disks** → **Managed Disks** in the portfolio. The Filtering Managed Disks panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Managed Disks panel is displayed.
3. Select the MDisk that you want information about and select **Show MDisk Group** from the list. Click **Go**. The Viewing Managed Disk Groups panel is displayed. The MDisk group is displayed in a table on the Viewing Managed Disk Groups panel.

Creating MDisk groups

You can create a new managed disk (MDisk) group using the Create a Managed Disk Group wizard.

If you intend to keep the virtual disk (VDisk) allocation within one disk controller system, ensure that the MDisk group that corresponds with a single disk controller system is presented by that disk controller system. This also enables nondisruptive migration of data from one disk controller system to another disk controller system and simplifies the decommissioning process if you want to decommission a disk controller system at a later time.

Ensure all MDisks that are allocated to a single MDisk group are of the same RAID-type. Using the same RAID-type 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 have three RAID-5 arrays in one group and add a non-RAID disk to this group, you lose access to all the data that is striped across the group if the non-RAID disk fails. Similarly, for performance reasons, you should not mix RAID types.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to create a new MDisk group:

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

2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Managed Disk Groups panel is displayed.
3. Select **Create an MDisk Group** from the task list and click **Go**. The Create a Managed Disk Group wizard begins.
4. Complete the Create a Managed Disk Group wizard.

Adding MDisks to MDisk groups

You can add managed disks (MDisks) to an MDisk group from the Adding Managed Disks to Managed Disk Group panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to add MDisks to an MDisk group:

1. Click **Work with Managed Disks** → **Managed Disk Groups** in the portfolio. The Filtering Managed Disk Groups panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Managed Disk Groups panel is displayed.
3. Select the MDisk group that you want to add MDisks to and select **Add MDisks** from the list. Click **Go**. The Adding Managed Disks to Managed Disk Group panel is displayed.
4. Select the MDisks that you want to add and click **OK**.

Removing MDisks from an MDisk group

You can remove managed disks (MDisks) from an MDisk group.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to remove an MDisk from an MDisk group:

1. Click **Work with Managed Disks** → **Managed Disk Groups** in the portfolio. The Filtering Managed Disk Groups panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Managed Disk Groups panel is displayed.
3. Select the MDisk Group that you want to delete MDisks from and select **Remove MDisks** from the list. Click **Go**. The Deleting Managed Disks from Managed Disk Group panel is displayed.
4. Select the MDisk that you want to remove.
5. Click **OK**.

Viewing the progress of an MDisk removal

You can view the progress of an MDisk removal from the Viewing MDisk Removal Progress panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to view the progress of an MDisk removal:

1. Click **Manage Progress** → **View Progress**. The View Progress panel is displayed.
2. Click the **MDisk Removal** link. The Viewing MDisk Removal Progress panel is displayed.

Renaming MDisk groups

You can rename a managed disk (MDisk) group from the Renaming Managed Disk Group panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to rename an MDisk group:

1. Click **Work with Managed Disks** → **Managed Disk Groups** in the portfolio. The Filtering Managed Disk Groups panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Managed Disk Groups panel is displayed.
3. Select the MDisk Group that you want to rename and select **Rename an MDisk Group** from the list. Click **Go**. The Renaming Managed Disk Group panel is displayed.

Displaying VDIs

You can display the virtual disks (VDIs) that use a managed disk (MDisk) group from the Viewing Virtual Disks panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to display the VDIs that use an MDisk group:

1. Click **Work with Managed Disks** → **Managed Disk Groups** in the portfolio. The Filtering Managed Disk Groups panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Managed Disk Groups panel is displayed.
3. Select the MDisk group that you want to display VDIs for and select **Show VDIs Using This Group** from the list. Click **Go**. The Viewing Virtual Disks panel is displayed.

Deleting MDisk groups

You can delete a managed disk (MDisk) group using the Deleting a Managed Disk Group panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to delete an MDisk group:

1. Click **Work with Managed Disks** → **Managed Disk Groups** in the portfolio. The Filtering Managed Disk Groups panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Managed Disk Groups panel is displayed.

3. Select the MDisk group you want to delete and select **Delete an MDisk Group** from the list. Click **Go**. The Deleting a Managed Disk Group panel is displayed.

Creating VDisks

You can create virtual disks (VDisks) using the Create Virtual Disks wizard.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to create VDisks:

1. Click **Work with Virtual Disks** → **Virtual Disks** in the portfolio. The Filtering Virtual Disks panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Virtual Disks panel is displayed.
3. Select **Create VDisks** from the task list and click **Go**. The Create Virtual Disks wizard begins.
4. Complete the Create Virtual Disks wizard.

Viewing the progress of VDisk formatting

You can view the progress of virtual disk (VDisk) formatting from the Viewing VDisk Formatting Progress panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to view the progress of VDisk formatting:

1. Click **Manage Progress** → **View Progress**. The View Progress panel is displayed.
2. Click the **VDisk Formatting** link. The Viewing VDisk Formatting Progress panel is displayed.

Migrating VDisks

You can migrate a virtual disk (VDisk) from one managed disk (MDisk) group to another from the Migrating VDisks panel.

The SAN Volume Controller provides various data migration features. You can use these features to move the placement of data both within MDisk groups and between MDisk groups. These features can be used concurrently with I/O operations. There are two ways that you can migrate data:

1. Migrate data (extents) from one MDisk to another MDisk within the same MDisk group. This can be used to remove active or overutilized MDisks. This can only be performed using the command-line interface (CLI).
2. Migrate VDisks from one MDisk group to another. This can be used to remove active MDisk groups; for example, you can reduce the utilization of a group of MDisks.

You can determine the usage of MDisks by gathering I/O statistics about MDisks and VDisks. After you have gathered this data, you can analyze it to determine which VDisks or MDisks are active.

When a migrate command is issued, a check ensures that the destination of the migrate has enough free extents to satisfy the command. If there are enough free extents, the command proceeds.

Note: You cannot use the SAN Volume Controller data migration function to move a VDisk between MDisk groups that have different extent sizes.

While the migration proceeds, 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 situation, 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 (by marking the error as fixed).
2. Migrate one or more VDIsks that are already created from the MDisk group to another group. This frees up extents in the group and allows the original migrations to be restarted.

Perform the following steps to migrate VDIsks between MDisk groups:

1. Perform the following steps to determine if VDIsks are overused:
 - a. Click **Manage Cluster** → **Start statistics collection** in the portfolio. The Starting the Collection of Statistics panel is displayed.
 - b. Enter 15 minutes for the interval and click **OK**. This generates a new I/O statistics dump file approximately every 15 minutes.
 - c. Wait at least 15 minutes before you proceed to the next step.
2. View the I/O statistics log.
 - a. Click **Service and Maintenance** → **List dumps** in the portfolio. The List Dumps panel is displayed.
 - b. Click **I/O Statistics Logs**. This lists the I/O statistics files that have been generated. These are prefixed with m and Nm for MDisk statistics and v for VDisk statistics.
 - c. Click a filename to view the contents of the log.
 - d. Analyze the dumps to determine which VDIsks are active. It might be helpful to also determine which MDisks are heavily utilized so you can spread the data that they contain more evenly across all the MDisks in the group. Either create a new MDisk group or determine an existing group that is not yet over used. You can do this by checking the I/O statistics files that were previously generated and ensuring that the MDisks or VDIsks in the target MDisk group are less utilized than the source group.
3. Stop the statistics collection by clicking **Manage Cluster** → **Stop statistics collection** in the portfolio.
4. Migrate the VDisk.
 - a. Click **Work with Virtual Disks** → **Virtual Disks** in the portfolio. The Filtering Virtual Disks panel is displayed.
 - b. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Virtual Disks panel is displayed.
 - c. Select the VDisk that you want to migrate and select **Migrate a VDisk** from the task list. Click **Go**. The Migrating Virtual Disks panel is displayed.
 - d. Select a target MDisk group from the **Target MDisk Group** list.

- e. Click **OK**.

Viewing the progress of VDisk migration

You can view the progress of virtual disk (VDisk) migration from the Viewing VDisk Migration Progress panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to view the progress of VDisk migration:

1. Click **Manage Progress** → **View Progress**. The View Progress panel is displayed.
2. Click the **VDisk Migration** link. The Viewing VDisk Migration Progress panel is displayed.

Shrinking VDisks

You can shrink a virtual disk (VDisk) from the Shrinking VDisks panel.

VDisks can be reduced in size, if necessary. However, if the VDisk contains data that is being used, *do not attempt to shrink a VDisk without first backing up your data*. The SAN Volume Controller arbitrarily reduces the capacity of the VDisk by removing one or more extents from those that are allocated to the VDisk. You cannot control which extents are removed so you cannot guarantee that it is unused space that is removed.

Attention: Use this feature *only* to make a target or auxiliary VDisk the same size as the source or master VDisk when you create FlashCopy mappings or Metro Mirror relationships. Ensure that the target VDisk is not mapped to any hosts prior to performing this operation.

This task assumes that you have already launched the SAN Volume Controller application.

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. Issue the following CLI command to determine the exact capacity of the source or master VDisk:

```
svcinfolsvdisk -bytes vdiskname
```

Where *vdiskname* is the name of the VDisk for which you want to determine the capacity.

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

3. Click **Work with Virtual Disks** → **Virtual Disks** in the portfolio. The Filtering Virtual Disks panel is displayed.
4. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Virtual Disks panel is displayed.
5. Select the VDisk that you want to shrink and select **Shrink a VDisk** from the task list. Click **Go**. The Shrinking Virtual Disks panel is displayed.

Viewing virtual disk-to-host mappings

You can view the virtual disk-to-host mappings from the Virtual Disk-to-Host Mappings panel.

This task assumes that you have already launched the SAN Volume Controller application.

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

1. Click **Work with Virtual Disks** → **Virtual Disk-to-Host Mappings** in the portfolio. The Filtering Virtual Disk-to-Host Mappings panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Virtual Disk-to-Host Mappings panel is displayed.
3. Click **Close** to close the panel.

Determining the relationship between VDIs and MDIs

You can use the SAN Volume Controller Console to determine the relationship between virtual disks (VDIs) and managed disks (MDIs).

This task assumes that you have already launched the SAN Volume Controller application.

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

1. Click **Work with Virtual Disks** → **Virtual Disks** in the portfolio. The Filtering Virtual Disks panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Virtual Disks panel is displayed.
3. Select the VDI that you want to view.
4. Select **Show MDIs This VDI is Using** from the task list and click **Go**. The Viewing Managed Disks panel is displayed. This panel lists the MDIs that the selected VDI uses.

Recovering from offline VDIs

You can use the SAN Volume Controller Console to recover from an offline virtual disk (VDI) after a node or an I/O group has failed.

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

Data loss scenario 1

One node in an I/O group has failed and failover has started on the second node. During the failover process, the second node in the I/O group fails before the data in the write cache is written to hard disk. The first node is successfully repaired but its hardened data is not the most recent version that is committed to the data store; therefore, it cannot be used. The second node is repaired or replaced and has lost its hardened data, therefore, the node has no way of recognizing that it is part of the cluster.

Perform the following steps to recover from an offline VDisk when one node has down-level hardened data and the other node has lost hardened data:

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.

Data loss scenario 2

Both nodes in the I/O group have failed and have been repaired. The nodes have lost their hardened data, therefore, the nodes have no way of recognizing that they are part of the cluster.

Perform the following steps to recover from an offline VDisk when both nodes have lost their hardened data and cannot be recognized by 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.

Moving offline VDIs to the recovery I/O group

After a node or an I/O group fails, you can move the offline virtual disks (VDIs) to the recovery I/O group.

This task assumes that you have already launched the SAN Volume Controller application.

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

1. Click **Work with Virtual Disks** → **Virtual Disks** in the portfolio. The Filtering Virtual Disks panel is displayed.
2. Type the name of the I/O group in the **I/O Group** filter box and select Offline from the **Status** list. Click **OK**. The Viewing Virtual Disks panel is displayed.
3. For each VDisk that is returned, select the VDisk and select **Modify a VDisk** from the task list and click **Go**. The Modifying Virtual Disk panel is displayed.
4. From the **I/O Group** list, select the name of the recovery I/O group. You might be asked to confirm and force the move; select to force the move. Click **OK**. The Viewing Virtual Disks panel is displayed.
5. Verify that the VDIs are in the recovery I/O group.

Moving offline VDIs to their original I/O group

After a node or an I/O group fails, you can move offline virtual disks (VDIs) to their original I/O group.

This task assumes that you have already launched the SAN Volume Controller application.

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

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

1. Click **Work with Virtual Disks** → **Virtual Disks** in the portfolio. The Filtering Virtual Disks panel is displayed.
2. Select Offline from the **Status** list and click **OK**. The Viewing Virtual Disks panel is displayed.

3. For each VDisk that is returned, select the VDisk and select **Modify a VDisk** from the task list and click **Go**. The Modifying Virtual Disk panel is displayed.
4. From the **I/O Group** list, select the name of the VDisk's original I/O group. You might be asked to confirm and force the move; select to force the move. Click **OK**. The Viewing Virtual Disks panel is displayed.
5. Verify that the VDIs are online.

Deleting VDIs

You can delete a virtual disk (VDisk) from the Deleting Virtual Disk panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to delete a VDisk:

1. Click **Work with Virtual Disks** → **Virtual Disks** in the portfolio. The Filtering Virtual Disks panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Virtual Disks panel is displayed.
3. Select the VDisk you want to delete and select **Delete a VDisk** from the list. Click **Go**. The Deleting Virtual Disk panel is displayed.
4. Click **OK**.

Using image mode VDIs

Ensure that you are familiar with using image mode virtual disks (VDIs).

An image mode VDisk provides a direct block-for-block translation from the managed disk (MDisk) to the VDisk with no virtualization. This mode is intended to allow virtualization of MDIs that already contain data that was written directly, not through a SAN Volume Controller. Image mode VDIs have a minimum size of 1 block (512 bytes) and always occupy at least one extent.

Image mode MDIs are members of an MDisk group but, they do not contribute to free extents.

In order for a VDisk to be online, it is necessary for *all* MDIs in the MDisk group that are associated with the VDisk to be online. This applies to image mode VDIs as well as managed mode VDIs. A VDisk will be offline if any MDisk in the MDisk group is offline even if that MDisk does not contribute any extents to the VDisk or even if the MDisk has no allocated extents.

An image mode VDisk behaves just as a managed mode VDisk in terms of Metro Mirror and FlashCopy. Image mode VDIs are different from managed mode in two ways:

- Migration. An image mode disk can be migrated to another image mode disk. It becomes managed while the migration is ongoing, but returns to image mode when the migration is complete.
- Quorum disks. Image mode disks cannot be quorum disks. This means that a cluster with only image mode disks does not have a quorum disk.

Creating an image mode VDisk

You can import storage that contains existing data and continue to use this storage but make use of the cache and advanced functions, such as Copy Services and data migration. These disks are known as image mode virtual disks (VDisks).

Make sure that you are aware of the following before you create image mode VDisks:

- Unmanaged-mode managed disks (MDisks) that contain existing data cannot be differentiated from unmanaged-mode MDisks 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 logical unit from your RAID controller to the cluster and refresh the view of MDisks. The newly detected disk is displayed.
- Do *not* manually add an unmanaged-mode MDisk that contains existing data to an MDisk group. If you do, the data is lost. When you use the command to convert an image mode VDisk from an unmanaged-mode disk, select the MDisk group where you want to add the VDisk.

See the following Web site for more information:

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

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to create an image mode VDisk:

1. Stop all I/O operations from the hosts.
2. Unmap the logical disks that contain the data from the hosts.
3. Perform the following steps to create one or more MDisk groups:
 - a. Click **Work with Managed Disks** → **Managed Disk Groups** in the portfolio. The Filtering Managed Disk Groups panel is displayed.
 - b. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Managed Disk Groups panel is displayed.
 - c. Select **Create an MDisk Group** from the task list and click **Go**. The Create Managed Disk Group wizard begins.
 - d. Use the wizard to create the MDisk group.
4. Perform the following steps to refresh the list of MDisks from the SAN Volume Controller Console:
 - a. Click **Work with Managed Disks** → **Managed Disks** in the portfolio. The Filtering Managed Disks panel is displayed.
 - b. From the **Mode** list, select **Unmanaged** and click **OK**. The Viewing Managed Disks panel is displayed.
 - If the new unmanaged-mode MDisk is not listed, you can perform a fabric-level discovery. Select **Discover MDisks** from the task list and click **Go**. When this process is complete, refresh the list of MDisks, and the unmanaged-mode MDisk should appear in the list.
5. Perform the following steps to convert the unmanaged-mode MDisk to an image mode VDisk:
 - a. Click **Work with Managed Disks** → **Managed Disks** in the portfolio. The Filtering Managed Disks panel is displayed.

- b. From the **Mode** list, select **Unmanaged** and click **OK**. The Viewing Managed Disks panel is displayed.
 - c. Select the unmanaged-mode MDisk and select **Create VDisk in Image Mode** from the task list. Click **Go**. The Create Image mode Virtual Disk wizard begins.
 - d. Use the wizard to select the MDisk group where the image mode VDisk should be added and the I/O group that will provide the data path for the VDisk.
6. Perform the following steps to map the new VDisk to the hosts that were previously using the data that the MDisk now contains:
 - a. Click **Work with Virtual Disks** → **Virtual Disks** in the portfolio. The Filtering Virtual Disks panel is displayed.
 - b. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Virtual Disks panel is displayed.
 - c. Select the VDisks and select **Map VDisks to a host** from the task list. Click **Go**. The Creating Virtual Disk-to-Host Mappings panel is displayed.
 - d. Select the host that you want to map the VDisk to and click **OK**.

Once it is mapped to a host object, the image mode VDisk is detected as a disk drive with which the host can perform I/O operations.

If you want to virtualize the storage on an image mode VDisk, you can transform it into a striped VDisk. Migrate the data on the image mode VDisk to managed-mode disks in another MDisk group.

Migration methods

Several methods can be used to migrate image mode virtual disks (VDisks) into managed mode VDisks.

In order to perform any type of migration activity on an image mode VDisk, the image mode VDisk must first be converted into a managed mode disk. The VDisk is automatically converted into a managed mode disk whenever any kind of migration activity is attempted. After the image mode to managed mode migration operation has occurred, the VDisk becomes a managed mode VDisk and is treated the same way as any other managed mode VDisk.

If the image mode disk has a partial last extent, this last extent in the image mode VDisk must be the first to be migrated. This migration is processed as a special case. After this special migration operation has occurred, the VDisk becomes a managed mode VDisk and is treated in the same way as any other managed mode VDisk. If the image mode disk does not have a partial last extent, no special processing is performed. The image mode VDisk is changed into a managed mode VDisk and is treated the same way as any other managed mode VDisk.

An image mode disk can also be migrated to another image mode disk. The image mode disk becomes managed while the migration is ongoing, but returns to image mode when the migration is complete.

You can perform the following types of migrations:

- Migrate extents
- Migrate a VDisk
- Migrate to image mode

Perform the following steps to migrate VDisks:

1. Dedicate one MDisk group to image mode VDisks.
2. Dedicate one MDisk group to managed mode VDisks.
3. Use the migrate VDisk function to move the VDisks.

Viewing the progress of image mode migration

You can view the progress of image mode migration from the Viewing Image Mode Migration Progress panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to view the progress of image mode migration:

1. Click **Manage Progress** → **View Progress**. The View Progress panel is displayed.
2. Click the **Image Mode Migration** link. The Viewing Image Mode Migration Progress panel is displayed.

Viewing the progress of extent migration

You can view the progress of image mode migration from the Viewing Extent Migration Progress panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to view the progress of extent migration:

1. Click **Manage Progress** → **View Progress**. The View Progress panel is displayed.
2. Click the **Extent Migration** link. The Viewing Extent Migration Progress panel is displayed.

Creating hosts

You can create a new host object from the Creating Hosts panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to create a new host object:

1. Click **Work with Hosts** → **Hosts** in the portfolio. The Filtering Hosts panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Hosts panel is displayed.
3. Select **Create a Host** from the task list and click **Go**. The Creating Hosts panel is displayed.
4. Type the name that you want to call the host in the **Host Name** field. If you do not specify a name, a default name is assigned.
5. Select the type of host from the **Type** list.
6. Select the I/O groups to map to this host from the **I/O Groups** list.
7. Assign a worldwide port name (WWPN). A WWPN consists of 16 hexadecimal digits (for example, 210100e08b251dd4). You can select a WWPN from the list

of candidates, or you can enter a WWPN that is not in the list. You can assign one or more WWPNs to a single logical host object.

8. Click **OK**.
9. Repeat steps 3 on page 123 through 8 for each host object that you want to create.

Filtering hosts

You can filter hosts from the Filtering Hosts panel. The criteria that you specify controls which hosts are displayed on the panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to specify filter criteria:

1. Click **Work with Hosts** → **Hosts** in the portfolio. The Filtering Hosts panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Hosts panel is displayed.

Viewing general details

You can view general details about a host object from the Viewing General Details panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to view the general details for a host object:

1. Click **Work with Hosts** → **Hosts** in the portfolio. The Filtering Hosts panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Hosts panel is displayed.
3. Click the name of the host for which you want to view general details. The Viewing General Details panel is displayed.
4. Click **Close** to return to the Viewing Hosts panel.

Viewing port details

You can view the ports that are attached to a host object from the Viewing Port Details panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to view the ports for a host object:

1. Click **Work with Hosts** → **Hosts** in the portfolio. The Filtering Hosts panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Hosts panel is displayed.
3. Click the name of the host for which you want to view port details. The Viewing General Details panel is displayed.
4. Click **Ports** to view the ports that are attached to the host object. The Viewing Port Details panel is displayed.

5. Click **Close** to return to the Viewing Hosts panel.

Viewing mapped I/O groups

You can view the I/O groups that are mapped to a host object from the Viewing Mapped I/O Groups panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to view the I/O groups that are mapped to a host object:

1. Click **Work with Hosts** → **Hosts** in the portfolio. The Filtering Hosts panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Hosts panel is displayed.
3. Click the name of the host for which you want to view the mapped I/O groups. The Viewing General Details panel is displayed.
4. Click **Mapped I/O Groups** to view the I/O groups that are mapped to the host object. The Viewing Mapped I/O Groups panel is displayed.
5. Click **Close** to return to the Viewing Hosts panel.

Displaying VDIsks that are mapped to a host

You can display the virtual disks (VDIsks) that are mapped to a host by using the Viewing Virtual Disks panel.

If a number of new VDIsks are mapped to a host and a number of devices are already running I/O operations, a significant number of errors might be logged. When the new VDisk is mapped, multiple recoverable errors can be logged in the event log. The event log displays the errors that are caused by a check condition. The errors state that there has been a change to the device information since the last logical unit number (LUN) operation.

This task assumes that you have already launched the SAN Volume Controller application.

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

1. Click **Work with Hosts** → **Hosts** in the portfolio. The Filtering Hosts panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Hosts panel is displayed.
3. Select the host and select **Show the VDIsks Mapped to this Host** from the task list. Click **Go**.

Modifying a host

You can modify a host from the Modifying Host panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to modify a host:

1. Click **Work with Hosts** → **Hosts** in the portfolio. The Filtering Hosts panel is displayed.

2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Hosts panel is displayed.
3. Select the host that you want to modify and select **Modify a Host** from the task list. Click **Go**. The Modifying Host panel is displayed.
You can modify the following attributes for a host:
 - Name
 - Type
 - I/O group
4. Click **OK** after you have selected the new attributes. If you are modifying a host-to-I/O group mapping that results in the loss of a VDisk-to-host mapping, the Forcing the Deletion of a Host to I/O Group Mappings panel is displayed. Perform one of the following steps:
 - Click **Force Remove** to remove the host-to-I/O group mapping.
 - Click **Cancel** to preserve the host-to-I/O group mapping.

Adding ports to a host

You can add ports to a host from the Adding Ports panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to add ports to a host:

1. Click **Work with Hosts** → **Hosts** in the portfolio. The Filtering Hosts panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Hosts panel is displayed.
3. Select the host that you want to add ports to and select **Add Ports** from the task list. Click **Go**. The Adding Ports panel is displayed.
4. Perform one of the following steps to add the ports:
 - Select the ports that you want to add from the **Available Ports** list and click **Add**.
 - Type the worldwide port names (WWPNs) that you want to add in the **Additional Ports** field.
5. Click **OK**.

Deleting ports from a host

You can delete ports from the Deleting Ports panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to delete ports from a host:

1. Click **Work with Hosts** → **Hosts** in the portfolio. The Filtering Hosts panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Hosts panel is displayed.
3. Select the host that you want to delete ports from and select **Delete Ports** from the task list. Click **Go**. The Deleting Ports panel is displayed.
4. Select the ports that you want to delete from the **Available Ports** list and click **Add**.

5. Click **OK**.

Replacing an HBA in a host

It is sometimes necessary to replace the host bus adapter (HBA) that connects the host to the SAN. You must notify the SAN Volume Controller of the new worldwide port number (WWPNs) that this HBA contains.

Before you begin this task, you must ensure that the switch is zoned correctly.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to notify the SAN Volume Controller of a change to a defined host object:

1. Locate the host object that corresponds with the host in which you have replaced the HBA.
2. Click **Work with Hosts** → **Hosts** in the portfolio. The Filtering Hosts panel is displayed.
3. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Hosts panel is displayed.
4. Select the host object and then select **Add Ports** from the task list. Click **Go**. The Adding ports panel is displayed.
5. Select the candidate WWPNs from the **Available Ports** list and click **Add**. Click **OK**. The Viewing Hosts panel is displayed.
6. Select the host object and select **Delete Ports** from the task list. Click **Go**. The Deleting Ports panel is displayed.
7. Select the WWPNs that you want to remove (the ones that correspond with the old HBA that was replaced) and click **Add**. Click **OK**.

Any mappings that exist between the host object and VDisks are automatically applied to the new WWPNs. Therefore, the host sees the VDisks as the same SCSI LUNs as before. See the *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide* or your multipathing device driver user's guide for adding device identifiers (virtual paths if you are using SDD) to existing device identifiers.

Deleting hosts

You can delete a host object from the Deleting Hosts panel.

A deletion fails if there are any virtual disk (VDisk)-to-host mappings for the host. If you attempt to delete the host and it fails due to the existence of VDisk mappings, you are presented with the opportunity to perform a forced deletion, which deletes the VDisk mappings before the host is deleted.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to delete a host object:

1. Click **Work with Hosts** → **Hosts** in the portfolio. The Filtering Hosts panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Hosts panel is displayed.

3. Select the host that you want to delete and select **Delete a host** from the task list. Click **Go**. The Deleting Hosts panel is displayed.
4. Verify that you are deleting the correct host and click **OK**.

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

Creating FlashCopy mappings

You can create a FlashCopy mapping using the Create a FlashCopy Mapping wizard.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to create FlashCopy mappings:

1. Click **Manage Copy Services** → **FlashCopy mappings** in the portfolio. The Filtering FlashCopy Mappings panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing FlashCopy Mappings panel is displayed.
3. Select **Create a Mapping** from the task list and click **Go**. The Create a FlashCopy Mapping wizard begins.
4. Complete the Create a FlashCopy Mapping wizard.

Filtering FlashCopy mappings

You can specify filter criteria from the Filtering FlashCopy mappings panel. The criteria that you select determines which FlashCopy mappings are displayed on the panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to specify filter criteria:

1. Click **Manage Copy Services** → **FlashCopy Mappings** in the portfolio. The Filtering FlashCopy mappings panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing FlashCopy Mappings panel is displayed.

Starting FlashCopy mappings

You can start FlashCopy mappings from the FlashCopy Mappings panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to start a FlashCopy mapping:

1. Click **Manage Copy Services** → **FlashCopy Mappings** in the portfolio. The Filtering FlashCopy mappings panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing FlashCopy Mappings panel is displayed.

3. Select the appropriate mapping's row from the table.
4. Select **Start a Mapping** from the task list and click **Go**. The Starting FlashCopy mappings panel is displayed.

Viewing the progress of a FlashCopy

You can view the progress of a FlashCopy from the Viewing FlashCopy Progress panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to view the progress of a FlashCopy:

1. Click **Manage Progress** → **View Progress**. The View Progress panel is displayed.
2. Click the **FlashCopy** link. The Viewing FlashCopy Progress panel is displayed.

Stopping FlashCopy mappings

You can stop FlashCopy mappings from the FlashCopy Mappings panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to stop a FlashCopy mapping:

1. Click **Manage Copy Services** → **FlashCopy Mappings** in the portfolio. The Filtering FlashCopy Mappings panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Mappings panel is displayed.
3. Select the appropriate mapping's row from the table.
4. Select **Stop a mapping** from the task list and click **Go**. The Stopping FlashCopy mappings panel is displayed.

Modifying FlashCopy mappings

You can change the attributes for a FlashCopy mapping from the Modifying FlashCopy Mappings panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to change the attributes for a FlashCopy mapping:

1. Click **Manage Copy Services** → **FlashCopy Mappings** in the portfolio. The Filtering FlashCopy Mappings panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing FlashCopy Mappings panel is displayed.
3. Select **Modify a mapping** from the task list and click **Go**. The Modifying FlashCopy Mappings panel is displayed.

Deleting FlashCopy mappings

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

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to delete a FlashCopy mapping:

1. Click **Manage Copy Services** → **FlashCopy mappings** in the portfolio. The Filtering FlashCopy mappings panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing FlashCopy mappings panel is displayed.
3. Select the appropriate mapping's row from the table.
4. Select **Delete a mapping** from the task list and click **Go**. The Deleting FlashCopy mapping panel is displayed.

Note: If the FlashCopy mapping is in active state, the Forcing the Deletion of a FlashCopy Mapping panel is displayed. Follow the instructions that are displayed on the Forcing the Deletion of a Flashy Copy Mapping panel.

Creating FlashCopy consistency groups

You can create a FlashCopy consistency group from the Creating FlashCopy Consistency Groups panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to create a FlashCopy consistency group:

1. Click **Manage Copy Services** → **FlashCopy Consistency Groups** in the portfolio. The Filtering FlashCopy Consistency Groups panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing FlashCopy Consistency Groups panel is displayed.
3. Select **Create a Consistency Group** from the task list and click **Go**. The Creating FlashCopy Consistency Groups panel is displayed.
4. Type the name of the FlashCopy consistency group in the **FlashCopy Consistency Group Name** field. If you do not specify a name, a default name is assigned to the FlashCopy consistency group.
5. Select the mappings that you want in the consistency group from the **FlashCopy Mappings** list and click **OK**.

Note: You can create the FlashCopy consistency group before you create the mappings and then add the FlashCopy mappings to the consistency group. To add FlashCopy mappings this way, you must use the Modifying FlashCopy Mapping panel or the Creating FlashCopy Mappings panel.

Filtering FlashCopy consistency groups

You can specify filter criteria from the Filtering FlashCopy Consistency Groups panel. The filter criteria that you specify controls the FlashCopy consistency groups that are displayed on the panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to specify filter criteria:

1. Click **Manage Copy Services** → **FlashCopy consistency groups** in the portfolio. The Filtering FlashCopy Consistency Groups panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing FlashCopy Consistency Groups panel is displayed.

Starting FlashCopy consistency groups

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

This task assumes that you have already launched the SAN Volume Controller application.

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

1. Click **Manage Copy Services** → **FlashCopy Consistency Groups** in the portfolio. The Filtering FlashCopy Consistency Groups panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The FlashCopy Consistency Groups panel is displayed.
3. Select the appropriate mapping's row from the table.
4. Select **Start a Consistency Group** from the task list and click **Go**. The Starting FlashCopy Consistency Groups panel is displayed.

Stopping FlashCopy consistency groups

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

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to stop a FlashCopy consistency group:

1. Click **Manage Copy Services** → **FlashCopy Consistency Groups** in the portfolio. The Filtering FlashCopy Consistency Groups panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The FlashCopy Consistency Groups panel is displayed.
3. Select the appropriate mapping's row from the table.
4. Select **Stop a Consistency Group** from the task list and click **Go**. The Stopping Consistency Groups panel is displayed.

Renaming FlashCopy consistency groups

You can rename a FlashCopy consistency group from the Renaming FlashCopy Consistency Group panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to rename a consistency group:

1. Click **Manage Copy Services** → **FlashCopy Consistency Groups** in the portfolio. The Filtering FlashCopy Consistency Groups panel is displayed.

2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing FlashCopy Consistency Groups panel is displayed.
3. Select **Rename a Consistency Group** from the task list and click **Go**. The Renaming FlashCopy Consistency Group panel is displayed.

Deleting FlashCopy consistency groups

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

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to delete a FlashCopy consistency groups:

1. Click **Manage Copy Services** → **FlashCopy Consistency Groups** in the portfolio. The Filtering FlashCopy consistency groups panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The FlashCopy consistency groups panel is displayed.
3. Select the appropriate mapping's row from the table.
4. Select **Delete a Consistency Group** from the task list and click **Go**. The Delete Consistency Groups panel is displayed.

Creating Metro Mirror relationships

You can create a Metro Mirror relationship using the Creating Metro Mirror Relationships wizard.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to create a Metro Mirror relationship:

1. Click **Manage Copy Services** → **Metro Mirror Relationships** in the portfolio. The Filtering Metro Mirror Relationships panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Metro Mirror Relationships panel is displayed.
3. Select **Create a Relationship** from the list and click **Go**. The Create a Metro Mirror Relationship wizard begins.
4. Complete the Create a Metro Mirror Relationship wizard.

Filtering Metro Mirror relationships

You can use the Filtering Metro Mirror Relationships panel to display a subset of the relationships that are configured on your cluster.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to specify some filter criteria:

1. Click **Manage Copy Services** → **Metro Mirror Relationships** in the portfolio. The Filtering Metro Mirror Relationships panel is displayed.

2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Metro Mirror Relationships panel is displayed.

Starting a Metro Mirror copy

You can start a Metro Mirror copy from the Starting Copy Process panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to start a Metro Mirror copy:

1. Click **Manage Copy Services** → **Metro Mirror Relationships** in the portfolio. The Filtering Metro Mirror Relationships panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Metro Mirror Relationships panel is displayed.
3. Select the relationship for which you want to start the copy process.
4. Select **Start Copy Process** and click **Go**. The Starting Copy Process panel is displayed.

Viewing the progress of Metro Mirror copy processes

You can view the progress of a Metro Mirror copy process from the Viewing Metro Mirror Progress panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to view the progress of a Metro Mirror copy process:

1. Click **Manage Progress** → **View Progress**. The View Progress panel is displayed.
2. Click the **Metro Mirror** link. The Viewing Metro Mirror Progress panel is displayed.

Stopping a Metro Mirror copy

You can stop a Metro Mirror copy from the Stopping Copy Process panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to stop a Metro Mirror copy:

1. Click **Manage Copy Services** → **Metro Mirror Relationships** in the portfolio. The Filtering Metro Mirror Relationships panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Metro Mirror Relationships panel is displayed.
3. Select the relationship for which you want to stop the copy process.
4. Select **Stop Copy Process** and click **Go**. The Stopping Copy Process panel is displayed.

Modifying Metro Mirror relationships

You can modify the attributes for a Metro Mirror Relationship from the Modifying Metro Mirror Relationships panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to modify the attributes for a Metro Mirror relationship:

1. Click **Manage Copy Services** → **Metro Mirror Relationships** in the portfolio. The Filtering Metro Mirror Relationships panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Metro Mirror Relationships panel is displayed.
3. Select the relationship that you want to modify.
4. Select **Modify a Relationship** from the task list and click **Go**. The Modifying Metro Mirror Relationship panel is displayed.

Switching the copy direction of a Metro Mirror relationship

You can reverse the roles of the primary and secondary virtual disks (VDisks) in a Metro Mirror relationship from the Switching the Direction of a Metro Mirror Relationship panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to reverse the roles of the primary and secondary VDisks:

1. Click **Manage Copy Services** → **Metro Mirror Relationships** in the portfolio. The Filtering Metro Mirror Relationships panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Metro Mirror Relationships panel is displayed.
3. Select **Switch Copy Direction** from the task list and click **Go**. The Switching the Direction of Metro Mirror Relationship panel is displayed.

Deleting Metro Mirror relationships

You can delete a Metro Mirror relationship from the Deleting Metro Mirror relationships panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to delete a Metro Mirror relationship:

1. Click **Manage Copy Services** → **Metro Mirror Relationships** in the portfolio. The Filtering Metro Mirror Relationship panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Metro Mirror Relationships panel is displayed.
3. Select the relationship that you want to delete by clicking on the appropriate line in the **Select** column.
4. Select **Delete a Relationship** from the task list and click **Go**. The Deleting Metro Mirror Relationship panel is displayed.

Creating Metro Mirror consistency groups

You can create a Metro Mirror consistency group using the wizard.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to create a Metro Mirror consistency group:

1. Click **Manage Copy Services** → **Metro Mirror Consistency Groups** in the portfolio. The Filtering Metro Mirror Consistency Groups panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type.
3. Select **Create a Consistency Group** from the task list and click **Go**. The wizard begins.
4. Complete the wizard.

Filtering Metro Mirror consistency groups

You can use the Filtering Metro Mirror Consistency Groups panel to display a subset of the consistency groups that are configured on your cluster.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to specify filter criteria:

1. Click **Manage Copy Services** → **Metro Mirror Consistency Groups** in the portfolio. The Filtering Metro Mirror Consistency Groups panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type.

Renaming a Metro Mirror consistency group

You can rename a Metro Mirror consistency group from the Renaming Metro Mirror Consistency Group panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to rename a Metro Mirror consistency group:

1. Click **Manage Copy Services** → **Metro Mirror Consistency Groups** in the portfolio. The Filtering Metro Mirror Consistency Groups panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type.
3. Select the group that you want to change.
4. Select **Rename a Consistency Group** from the task list and click **Go**. The Renaming Metro Mirror Consistency Group panel is displayed.

Starting a Metro Mirror consistency group copy

You can start a Metro Mirror copy from the Starting Copy Process panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to start a Metro Mirror copy:

1. Click **Manage Copy Services** → **Metro Mirror Consistency Groups** in the portfolio. The Filtering Metro Mirror Consistency Groups panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Metro Mirror Consistency Groups panel is displayed.
3. Select the relationship for which you want to start the copy process.
4. Select **Start Copy Process** and click **Go**. The Starting Copy Process panel is displayed.

Stopping a Metro Mirror consistency group copy

You can stop a Metro Mirror copy from the Stopping Copy Process panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to stop a Metro Mirror copy:

1. Click **Manage Copy Services** → **Metro Mirror Consistency Groups** in the portfolio. The Filtering Metro Mirror Consistency Groups panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Metro Mirror Consistency Groups panel is displayed.
3. Select the group for which you want to stop the copy process.
4. Select **Stop Copy Process** and click **Go**. The Stopping Copy Process panel is displayed.

Deleting Metro Mirror consistency groups

You can delete a Metro Mirror consistency group from the Deleting Metro Mirror Consistency Groups panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to delete a Metro Mirror consistency group:

1. Click **Manage Copy Services** → **Metro Mirror Consistency Groups** in the portfolio. The Filtering Metro Mirror Consistency Groups panel is displayed.
2. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type.
3. Select the group that you want to delete.
4. Select **Delete a Consistency Group** from the task list and click **Go**. The Deleting Metro Mirror Consistency Group panel is displayed.

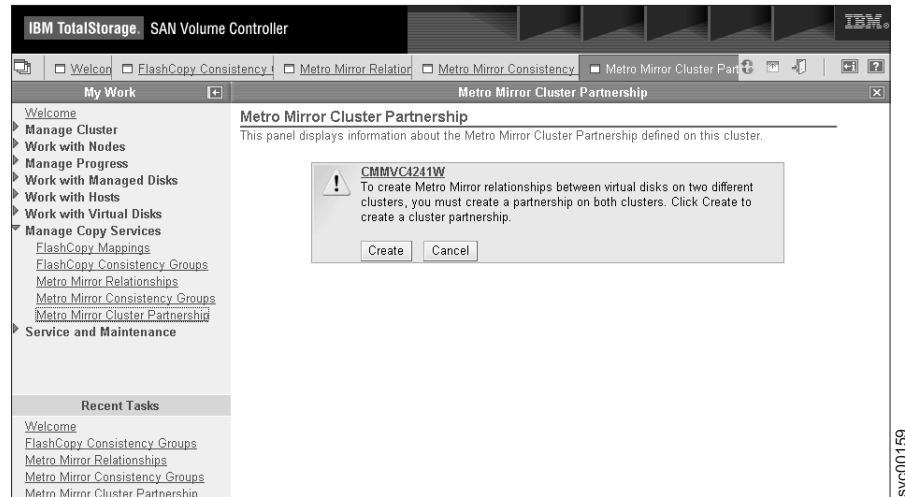
Creating Metro Mirror partnerships

You can create a Metro Mirror partnership from the Creating Metro Mirror Partnership panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to create a Metro Mirror partnership:

1. Click **Manage Copy Services** → **Metro Mirror Cluster Partnership** in the portfolio. The Metro Mirror Cluster Partnership panel is displayed.



2. Click **Create**. The Create Cluster Partnerships panel is displayed.

Modifying Metro Mirror partnerships

You can change the bandwidth available for background copy in a Metro Mirror partnership.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to change a Metro Mirror partnership:

1. Click **Manage Copy Services** → **Metro Mirror Cluster Partnerships** in the portfolio. The Metro Mirror Cluster Partnership panel is displayed.
2. Click **Modify**. The Modify Cluster Partnership panel is displayed.

Note: You can set the bandwidth attribute for the path from cluster A to cluster B to a different setting from the setting used for the path from cluster B to cluster A.

Deleting Metro Mirror partnerships

You can delete a Metro Mirror partnership on the local cluster from the Delete Cluster Partnership panel.

This task assumes that you have already launched the SAN Volume Controller application.

The Metro Mirror partnership must be deleted on both the local and remote cluster for the partnership to be completely removed.

Perform the following steps to delete a Metro Mirror partnership on the local cluster:

1. Click **Manage Copy Services** → **Metro Mirror Cluster Partnerships** in the portfolio. The Metro Mirror Cluster Partnership panel is displayed.
2. Click **Delete**. The Delete Cluster Partnership panel is displayed.

3. Click **Delete** to delete the Partnership on the local cluster or click **Cancel** to return to the Metro Mirror Cluster Partnership panel.

Viewing the feature log

You can view the feature log for the cluster from the Feature Log panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following step to view the feature log for the cluster:

Click **Service and Maintenance** → **View Feature Log** in the portfolio. The Feature Log panel is displayed.

Viewing and updating feature settings

You can view and update the feature settings in the Featurization Settings panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to view and update the feature settings:

1. Click **Service and Maintenance** → **Set Features** in the portfolio. The Featurization Settings panel is displayed.
2. Disable or Enable current features.
3. Click **Update Feature Settings**.

Running the cluster maintenance procedure

You can use the SAN Volume Controller Console to run the cluster maintenance procedure.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to run the cluster maintenance procedure:

1. Click **Service and Maintenance** → **Run Maintenance Procedures** in the portfolio. The Maintenance Procedures panel is displayed.
2. Click **Start Analysis** to analyze the cluster error log. The Maintenance panel is displayed.

If you click the error code of a particular error log entry, you are guided through a series of actions that help you estimate the state of the cluster and determine if the error was an isolated event or a component failure. If a component has failed, it might be necessary to exchange that component. Where necessary, images of the failing component are displayed. If a repair is performed successfully, the state of an error record in the error log changes from an unfixed error to a fixed error.

Configuring error notification settings

You can configure the error notification settings for the cluster from the Modify SNMP Error Notification Settings panel.

This task assumes that you have already launched the SAN Volume Controller application.

The error notification settings apply to the entire cluster. You can specify the types of errors that cause the cluster to send a notification. The cluster sends a Simple Network Management Protocol (SNMP) notification. The SNMP setting represents the kind of error.

Perform the following steps to configure the error notification settings:

1. Click **Service and Maintenance** → **Set Error Notification** in the portfolio. The Modify SNMP Error Notification Settings panel is displayed.

The following table describes the three types of notification:

Notification type	Description
All	Report all errors at or above the threshold limit, including information events.
Hardware only	Report all errors at or above the threshold limit, excluding information events.
None	Do not report any errors or information events. This option disables error notification.

If you specify *All* or *Hardware Only*, errors are reported to the SNMP destinations of your choice. To specify an SNMP destination, you *must* provide a valid IP address and SNMP community string.

Note: A valid community string can contain up to 60 letters or digits, without any spaces. A maximum of six SNMP destinations can be specified. When you create the cluster or enable error notification for the first time, you are asked to specify only one SNMP destination. You can add five additional destinations by using the Error Notification options.

The SAN Volume Controller uses the error notifications settings to call Home if errors occur. You must specify *All* or *Hardware Only* and send the trap to the master console if you want the SAN Volume Controller to call Home when errors occur.

2. Click **Modify Settings** to update the settings.

Displaying and saving log and dump files

You can save the log and dump files for nodes.

You can save dump data for any node in the cluster. When you use this procedure to display dump data only, the dump files for the configuration node are displayed. An option on the dumps menu allows you to display data from other nodes. If you choose to display or save data from another node, that data is first copied to the configuration node.

The software dump files contain dumps of the SAN Volume Controller 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.

The **List dumps** option supports the following file types:

- Error logs
- Configuration logs

- I/O statistic logs
- I/O trace logs
- Feature logs
- Software dumps

Perform the following steps to display log and dump files:

1. Click **Service and Maintenance** → **List Dumps** in the portfolio. The List Dumps panel is displayed.

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, 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 are displayed. Dumps and logs on all nodes in the cluster can be deleted on or copied to the configuration 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.

2. Copy the files to your local workstation by right-clicking on the filename and using the **Save Link As...** (Netscape) or **Save Target As...** (Internet Explorer) option from the Web browser.

Analyzing the error log

You can analyze the error log from the Analyze Error Log panel.

This task assumes that you have already launched the SAN Volume Controller application.

Note: Log files that are copied to the configuration node are *not* automatically deleted by the SAN Volume Controller. Therefore, clear any files that are copied to the configuration node to ensure that the configuration node does not run out of space.

Perform the following steps to analyze the error log:

1. Click **Service and Maintenance** → **Analyze Error Log** in the portfolio. The Error log analysis panel is displayed.

The Error log analysis panel lets you 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 that the table is sorted by either error priority or time. For error priority, the most serious errors are the lowest-numbered errors. Therefore, they are displayed first in the table.

Either the oldest or the latest entry can be displayed first in the table. You can also select how many error log entries are 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 99.

2. 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 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 as unfixed.

3. Click **Clear log** to erase the entire cluster error log.

Note: Clicking **Clear log** does *not* fix the existing errors.

Recovering a node and returning it to the cluster

After a node or an I/O group fails, you can use the SAN Volume Controller to recover a node and return it to the cluster.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to recover a node and return it to the cluster:

1. Click **Work with Nodes** → **Nodes** in the portfolio. The Viewing Nodes panel is displayed.
2. Verify that the node is offline.
3. Select the offline node.
4. Select **Delete a Node** from the task list and click **Go**. The Deleting Node from Cluster panel is displayed.
5. Click **Yes**.
6. Verify that the node can be seen on the fabric.
7. If the nodes are repaired by replacing the front panel module or a node is repaired by replacing it with another node, the worldwide node name (WWNN) for the node changes. In this case, you must follow these additional steps:
 - a. At the end of the recovery process it is necessary to follow the subsystem device driver (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 Multipath Subsystem Device Driver: User's Guide*.
 - b. You might also have 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 must modify the port groups that belong to the cluster because the WWNN or worldwide port names (WWPNs) 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 from which it was removed. 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. If you are adding the node into the cluster for the first time, you must record the following information:

- Node serial number
 - WWNN
 - All WWPNs
 - I/O group that the node belongs to
8. Add the node back into the cluster.

- a. From the Viewing Nodes panel, select **Add a Node** from the task list and click **Go**. The Adding a Node to a Cluster panel is displayed.
 - b. 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.
 - c. Click **OK**.
9. Verify that the node is online by refreshing the Viewing Nodes panel.

Note: If the panel does not refresh, close the panel and reopen it.

Managing SSH keys

You can manage SSH keys from the SAN Volume Controller Console.

The communication between the SAN Volume Controller Console software and the SAN Volume Controller cluster is through the Secure Shell (SSH) protocol. In this protocol, the SAN Volume Controller Console software acts as the SSH client and the SAN Volume Controller cluster acts as the SSH host server.

As an SSH client, the SAN Volume Controller Console must use an SSH2 RSA key pair composed of a public key and a private key which are coordinated when the keys are generated. The SSH client public key is stored on each SAN Volume Controller cluster with which the SAN Volume Controller Console communicates. The SSH client private key is known to the SAN Volume Controller Console software by being stored in a specific directory with a specific name. If the SSH protocol detects the key pair is mismatched, the SSH communication fails.

The SAN Volume Controller Console externalizes the status of a mismatched or invalid SAN Volume Controller Console client key pair in the **Availability Status** column of the Cluster panel.

You can use the SAN Volume Controller Console to perform the following SSH key management tasks:

- Add SSH keys to other hosts
- Add additional keys to the SAN Volume Controller Console
- Replace the client SSH key private key
- Replace the SSH key pair
- Reset the SSH fingerprint
- Reset a refused SSH key

Adding SSH keys for hosts other than the master console

You can add Secure Shell (SSH) keys on other hosts.

Perform the following steps to add SSH keys on hosts other than the master console:

1. Generate the public private key pair on each host that you want to use the SAN Volume Controller command-line interface. See the information that came with your SSH client for specific details about using the key generation program that comes with your SSH client.
2. Copy the public keys from each of these hosts to the master console.
3. Securely copy these public keys from the master console to the cluster.
4. Repeat for each public key copied onto the master console in step 2.

Adding subsequent SSH public keys to the SAN Volume Controller

You can add subsequent Secure Shell (SSH) public keys to the SAN Volume Controller from the SSH Public Key Maintenance panel.

This task assumes that you are at the Welcome panel for the SAN Volume Controller Console.

The SSH key allows the master console (where the SAN Volume Controller Console is running) to access the cluster.

During the cluster creation wizard, you added a SSH key to the cluster. You can add additional SSH keys to grant SSH access to other servers.

Perform the following steps to add additional SSH keys:

1. Click **Clusters** in the portfolio.
2. Click the cluster whose SSH keys you want to maintain.
3. Select **Maintain SSH Keys** from the task list and click **Go**. The SSH Public Key Maintenance panel is displayed.
4. Follow the instructions that are on the SSH Public Key Maintenance panel.
5. Click Add Key when you have completed the SSH Public Key Maintenance panel.

After the initial configuration of the cluster has been performed using the SAN Volume Controller Console and at least one SSH client key has been added, the remainder of the configuration can either be performed using the SAN Volume Controller Console or the command-line interface.

Replacing the client SSH private key known to the SAN Volume Controller software

You can replace the client SSH private key that is known to the SAN Volume Controller software.

Attention: If you have successfully contacted other SAN Volume Controller clusters, you will break that connectivity if you replace the client SSH private key that is known to the SAN Volume Controller software.

Perform the following steps to replace the client SSH private key:

1. Sign off the SAN Volume Controller Console.
2. Using the Windows Services facility, perform the following steps to stop the IBM CIM Object Manager:
 - a. Click **Start** → **Settings** → **Control Panel**.
 - b. Double-click **Administrative Tools**.
 - c. Double-click **Services**.
 - d. Select **IBM CIM Object Manager** in the list of services, right click, and select **Stop**.
 - e. Leave the Services panel open.
3. Perform the following steps to copy the client SSH private key into the appropriate SAN Volume Controller Console directory:
 - a. Open a command prompt window.
 - b. Issue the following command:

```
copy filename C:\Program Files\IBM\svconsole\cimom\icat.ppk
```

- Where *filename* is the path and file name of the client SSH private key.
4. Select **IBM CIM Object Manager** in the list of services, right click and select **Start**.
 5. Log on to the SAN Volume Controller Console.
 6. Click **Clusters** in the portfolio.
 7. Check the status of the cluster.

Replacing the SSH key pair

You can use the SAN Volume Controller Console to replace the Secure Shell (SSH) key pair.

Hypothetical scenarios where you must replace the SSH key pair

The following scenarios require you to replace the SSH key pair:

- If you change the SSH keys that are used by the master console to communicate with the SAN Volume Controller Console, you must store the client SSH private key in the SAN Volume Controller Console software and then store the client SSH public key on the SAN Volume Controller cluster.
- If you change the IP address of your SAN Volume Controller cluster after you have added the cluster to SAN Volume Controller Console, the SAN Volume Controller Console is not aware of the existence of the cluster.

Using the SAN Volume Controller Console to replace the SSH key pair

Perform the following steps to remove the cluster and replace the SSH key pair:

1. Click **Clusters** in the portfolio.
2. Select the cluster for which you want to replace the key.
3. Click **Maintain SSH Keys** from the task list and click **Go**. The SSH Public Key Maintenance panel is displayed.
4. Type your user name and password.
5. Click the **Maintain SSH Keys** option. The window opens to enable you to enter the client SSH public key information that is to be stored on the cluster.
 - If you are adding the SSH client key for the master console, click **Browse** and locate the public key you generated earlier.
 - If you are adding an SSH client key for another system, either click **Browse** and locate the public key or cut and paste the public key into the direct input field.
6. Click **Administrator**.
7. Type a name of your choice in the **ID** field that uniquely identifies the key to the cluster.
8. Click **Add Key**.
9. Click **Maintain SSH Keys**.
10. Click **Show IDs** to see all key IDs that are loaded on the SAN Volume Controller.

Resetting a refused SSH key

You can reset a refused SSH key relationship between the SAN Volume Controller Console and the SAN Volume Controller cluster.

Because the client SSH key pair must be coordinated across two systems, you might have to take one or more actions to reset the pair of keys.

Perform one or more of the following actions to reset the refused client SSH key pair:

- Replace the client SSH public key on the SAN Volume Controller cluster.
- Replace the client SSH private key known to the SAN Volume Controller software.

Resetting the SSH fingerprint

You can reset the Secure Shell (SSH) fingerprint for a cluster that is managed by the SAN Volume Controller Console for your configuration by using the Resetting the SSH Fingerprint panel.

You must have superuser administrator authority to reset the SSH fingerprint.

If you have changed the name of the master console, you must also change the master console host name in the IBM WebSphere® Application Server files.

The SAN Volume Controller Console and the cluster communicate through the SSH protocol. In this protocol, the SAN Volume Controller Console 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 that result in the SSH client flagging the SSH host fingerprint as changed and, therefore, no longer valid.

The SAN Volume Controller Console displays the status of the cluster SSH server key in the **Availability Status** column of the Viewing Clusters panel.

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.
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 you are prompted with the message CMMVC3201W.

Availability status is changed to OK.

Chapter 4. Using the CLI

The SAN Volume Controller cluster command-line interface (CLI) is a collection of commands that you can use to manage the SAN Volume Controller.

Overview

The CLI commands use the secure shell (SSH) connection between the SSH client software on the host system and the SSH server on the SAN Volume Controller cluster.

Before you can use the CLI, you must have already created a cluster.

You must perform the following actions to use the CLI from a client system:

- Install and set up SSH client software on each system that you plan to use to access the CLI.
- Generate an SSH key pair on each SSH client.
- Store the SSH public key for each SSH client on the SAN Volume Controller.

Note: After the first SSH public key is stored, you can add additional SSH public keys using either the SAN Volume Controller Console or the CLI.

You can use the CLI to perform the following functions:

- Set up of the cluster, its nodes, and the I/O groups
- Analyze error logs
- Set up and maintenance of managed disks (MDisk) and MDisk groups
- Set up and maintenance of client public SSH keys on the cluster
- Set up and maintenance of virtual disks (VDisks)
- Set up of logical host objects
- Map VDisks to hosts
- Navigate from managed hosts to VDisks and to MDisks, and the reverse direction up the chain
- Set up and start Copy Services:
 - FlashCopy and FlashCopy consistency groups
 - Synchronous Metro Mirror and Metro Mirror consistency groups

Preparing the SSH client system

Before you can issue command-line interface (CLI) commands from the host to the cluster, you must prepare the Secure Shell (SSH) client system.

Windows operating systems

If you have purchased the master console hardware and software from IBM, PuTTY for Windows operating systems has been previously installed for you.

If you are installing the master console on your own hardware with a Windows operating system, you can download PuTTY from the following Web site:

<http://www.chiark.greenend.org.uk/~sgtatham/putty/>

The following Web site offers SSH client alternatives for Windows:

<http://www.openssh.com/windows.html>

Cygwin software has an option to install an OpenSSH client. You can download OpenSSH from the following Web site:

<http://www.cygwin.com/>

AIX operating systems

For AIX 5L Power 5.1 and 5.2, you can get OpenSSH from the Bonus Packs but you also need its prerequisite, OpenSSL, from the AIX toolbox for Linux applications for Power Systems. For AIX 4.3.3, you can get the software from the AIX toolbox for Linux applications.

You can also get the AIX installation images from IBM DeveloperWorks at the following Web site:

<http://oss.software.ibm.com/developerworks/projects/openssh>

Linux operating systems

OpenSSH is installed by default on most Linux distributions. If it is not installed on your system, consult your installation media or visit the following Web site:

<http://www.openssh.org/portable.html>

OpenSSH is able to run on a wide variety of additional operating systems. For more information visit the following Web site:

<http://www.openssh.org/portable.html>

Preparing the SSH client system to issue CLI commands

In order to issue command-line interface (CLI) commands to the cluster from a host, you must prepare the Secure Shell (SSH) client on the host so that the host is accepted by the SSH server on the cluster.

If you want to use a host that requires a different type of SSH client (for example, OpenSSH), follow the instructions for that software.

Perform the following steps to enable your host to issue CLI commands:

1. For the master console and Windows hosts:
 - a. Generate an SSH key pair using the PuTTY key generator.
 - b. Store the SSH clients public key on the cluster (using a browser pointing to the SAN Volume Controller Console).
 - c. Configure the PuTTY session for the CLI.
2. For other types of hosts:
 - a. Follow the instructions specific to the SSH client to generate an SSH key pair.
 - b. Store the SSH clients public key on the cluster (using a browser pointing to the SAN Volume Controller Console or the CLI from an already established host).

- c. Follow the instructions that are specific to the SSH client to establish an SSH connection to the SAN Volume Controller cluster.

Preparing the SSH client on an AIX host

When you use AIX hosts, Secure Shell (SSH) logins are authenticated on the SAN Volume Controller cluster using the RSA-based authentication that is supported in the OpenSSH client available for AIX.

RSA-based authentication uses public-key cryptography to allow the encryption and decryption to use separate keys. Therefore, it is not possible to derive the decryption key from the encryption key. Initially, the user creates a public/private key pair for authentication purposes. The server (the SAN Volume Controller cluster in this case) knows the public key, and only the user (the AIX host) knows the private key. Because physical possession of the public key allows access to the cluster, the public key must be kept in a protected place. You can store the public key in the `/.ssh` directory on the AIX host with restricted access permissions.

When you use the AIX host to log into the SAN Volume Controller cluster, the SSH program on the SAN Volume Controller cluster sends the AIX host the key pair that it wants to use for authentication. The AIX server checks if this key is permitted, and if so, sends the SSH program that is running on behalf of the user a challenge. The challenge is a random number that is encrypted by the user's public key. The challenge can only be decrypted using the correct private key. The user's client (the AIX host) uses the private key to decrypt the challenge and prove that the user has the private key. The private key is not shown to the server (the SAN Volume Controller cluster) or to anyone who might be intercepting the transmissions between the AIX host and the SAN Volume Controller cluster.

Perform the following steps to set up an RSA key pair on the AIX host and the SAN Volume Controller cluster:

1. Create an RSA key pair by issuing a command on the AIX host that is similar to the following command:

```
ssh-keygen -t rsa1
```

Tip: Issue the command from the `$HOME/.ssh` directory.

This process generates two user named files. If you select the name *key*, the files are named *key* and *key.pub*. Where *key* is the name of the private key and *key.pub* is the name of the public key.

2. Store the private key from this key pair on the AIX host, in the `$HOME/.ssh` directory, in the `$HOME.ssh/identity` file. If you are using multiple keys, all of the keys must appear in the identity file.
3. Store the public key on the master console of the SAN Volume Controller cluster. Typically this can be done with ftp; however, the master console might have ftp disabled for security reasons, in which case an alternative method, such as secure copy is required. You can then use the SAN Volume Controller Console, to transfer the public key to the cluster. Select an access level of either administrator or service.

You can now access the cluster from the AIX host using an SSH command similar to the following:

```
ssh admin@my_cluster
```

Where *admin* means that you associated the key with an administrative ID and *my_cluster* is the name of the cluster IP.

Refer to your client's documentation for SSH on your host system for more host specific details regarding this task.

Issuing CLI commands from a PuTTY SSH client system

You can issue command-line interface (CLI) commands from a PuTTY SSH client system.

Perform the following steps to issue CLI commands:

1. Open a command prompt.
2. Issue the following command to set the path environment variable to include the PuTTY directory:

```
set path=C:\Program Files\putty;%path%
```

Where *Program Files* is the directory where PuTTY is installed.

3. Use the PuTTY plink utility to connect to the SSH server on the cluster.

Running the PuTTY and plink utilities

Ensure that you are familiar with how to run the PuTTY and plink utilities.

The Secure Shell (SSH) protocol specifies that the first access to a new host server sends 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 if you accept the responsibility of connecting with this host. If you type *y*, the host fingerprint and IP address are saved by the SSH client.

When you use PuTTY, you must also type *y* to accept this host fingerprint. However, the host fingerprint and IP address are stored in the registry for the user name that is logged onto Windows.

The SSH protocol also specifies that once the SSH server public key is accepted, another challenge is presented if the fingerprint of an SSH server changes from the one previously accepted. In this case, you must decide if you want to accept this changed host fingerprint.

Note: The SSH server keys on the SAN Volume Controller are regenerated when a microcode load is performed on the cluster. As a result, a challenge is sent because the fingerprint of the SSH server has changed.

All command-line interface (CLI) commands are run in an SSH session. You can run the commands in one of the following modes:

- An interactive prompt mode
- A single line command mode, which is entered one time to include all parameters

Interactive mode

For interactive mode, you can use the PuTTY executable to open the SSH restricted shell.

The following is an example of the command that you can issue to start interactive mode:

```
C:\support utils\putty admin@svconsoleip
```

Where *support utils\putty* is the location of your *putty.exe* file and *svcconsoleip* is the IP address of your SAN Volume Controller Console.

If you were to issue the **svcinfo lsshkeys** command, which lists the SSH client public keys that are stored on the SAN Volume Controller cluster, the following output is displayed:

```
IBM_2145:your_cluster_name:admin>svcinfo lsshkeys -user all -delim :
id:userid:key identifier
1:admin:smith
2:admin:jones
```

You can type **exit** and press **Enter** to escape the interactive mode command.

The following is an example of the host fingerprint challenge when using **plink** 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_2145:your_cluster_name:admin>
```

Single line command

For single line command mode, you can type the following all on one command line:

```
C:\Program Files\IBM\svcconsole\cimom>
plink admin@9.43.225.208 svcinfo lsshkeys
-user all -delim :
Authenticating with public key "imported-openssh-key"
id:userid:key identifier
1:admin:smith
2:admin:jones
```

Note: If you are submitting a CLI command with all parameters in single line command mode, you are challenged upon first appearance of the SSH server host fingerprint. Ensure that the SSH server host fingerprint is accepted before you submit a batch script file.

The following is an example of the host fingerprint challenge when using **plink** 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

```

Setting the cluster time using the CLI

You can use the command-line interface (CLI) to set the cluster time.

Perform the following steps to set the cluster time:

1. Issue the **svcinfo showtimezone** CLI command to display the current time-zone settings for the cluster. The cluster ID and the associated time zone are displayed.
2. Issue the **svcinfo lstimezones** CLI command to list the time zones that are available on the cluster. A list of valid time zone settings are displayed. The specific cluster ID and the assigned time zone are indicated in the list.
3. Issue the following CLI command to set the time zone for the cluster.

```
svctask settimezone -timezone time_zone_setting
```

Where *time_zone_setting* is the new time zone that have you chosen from the list of time zones that are available on the cluster.

4. Issue the following CLI command to set the time for the cluster:

```
svctask setclustertime -time 031809142005
```

Where *031809142005* is the new time that you want to set for the cluster. You must use the *MMDDHHmmYYYY* format to set the time for the cluster.

Reviewing and setting the cluster features using the CLI

You can use the command-line interface (CLI) to set up the cluster features.

Perform the following steps to set up the cluster features:

1. Issue the **svcinfo lslicense** CLI command to return the current license (featurization) settings for the cluster. The feature settings are displayed in a list that indicates if the feature is enabled or disabled.
2. Issue the **svctask chlicense** CLI command to change the licensed settings of the cluster. Because the feature settings are entered when the cluster is first created, update the settings only if you have changed your license. You can change the following values:
 - FlashCopy: disabled or enabled
 - Metro Mirror: disabled or enabled

- Virtualization limit: number, in gigabytes (1073741824 bytes)

Displaying cluster properties using the CLI

You can use the command-line interface (CLI) to display the properties for a cluster.

Perform the following steps to display cluster properties:

Issue the **svcinfolcluster** command to display the properties for a cluster.

For example, issue the following command:

```
svcinfolcluster -delim : 10030a007e5
```

Where *10030a007e5* is the name of the cluster.

The output from this command includes the following information for each cluster on the fabric:

- Cluster ID
- Cluster name
- Cluster IP address
- Cluster service mode IP address

Maintaining passwords for the front panel using the CLI

You can use the command-line interface (CLI) to view and change the status of the password reset feature for the SAN Volume Controller front panel.

The menu on the SAN Volume Controller front panel provides an option to reset the administrator password. This option resets the administrator password to a random string and displays the new administrator password on the SAN Volume Controller front panel. You can use this new administrator password to access the system. For password protection, change the administrator password at the next login.

Perform the following steps to view and change the status of the password reset feature:

1. Issue the **svctask setpwdreset** CLI command to view and change the status of the password reset feature for the SAN Volume Controller front panel. Passwords can consist of A - Z, a - z, 0 - 9, and underscore.
2. Record the administrator password because you cannot access the cluster without it.

Adding nodes to a cluster using the CLI

You can use the command-line interface (CLI) to add nodes to a cluster.

Before you add a node to a cluster, you must 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. If you are replacing a node and the switch is zoned by worldwide port name (WWPN) rather than by switch port, make sure that the switch is configured such that the node being added is in the same VSAN/zone.

Special procedures when adding a node to a cluster

Applications on the host systems direct I/O operations to file systems or logical volumes that are mapped by the operating system to virtual paths (vpaths), which are pseudo disk objects supported by the Subsystem Device Driver (SDD). SDD maintains an association between a VPath and a SAN Volume Controller virtual disk (VDisk). This association uses an identifier (UID) which is unique to the VDisk and is never reused. The UID allows SDD to directly associate vpaths with VDIs.

SDD operates within a protocol stack that contains disk and fibre channel device drivers that allow it to communicate with the SAN Volume Controller 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 worldwide name (WWN) for the fibre channel node and ports.

If an error occurs, the 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 that were previously used.

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

Before you add a node to the cluster, you must 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 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 back-end storage.

If any of the previous conditions are true, the following special procedures apply:

- The node must be added to the same I/O group that it was previously in. You can use the command-line interface (CLI) command **svcinfolnode** or the SAN Volume Controller Console to determine the WWN of the cluster nodes.
- Before you add the node back into the cluster, you must shut down all of the hosts using the cluster. The node must then be added before the hosts are rebooted. If the I/O group information is unavailable or it is inconvenient to shut down and reboot all of the hosts using the cluster, then do the following:
 - On all of the hosts connected to the cluster, unconfigure the fibre-channel adapter device driver, the disk device driver and multipathing driver before you add the node to the cluster.
 - Add the node to the cluster and then reconfigure the fibre-channel adapter device driver, the disk device driver, and multipathing driver.

Scenarios where the special procedures can apply

The following two scenarios describe situations where the special procedures can apply:

- Four nodes of an eight-node cluster have been lost because of the failure of a pair of 2145 uninterruptible power supply (2145 UPS) or four 2145 uninterruptible power supply-1U (2145 UPS-1U). In this case, the four nodes must be added back into the cluster using the CLI command **svctask addnode** or the SAN Volume Controller Console.

- A user decides to delete four nodes from the cluster and add them back into the cluster using the CLI command **svctask addnode** or the SAN Volume Controller Console.

Perform the following steps to add nodes to a cluster:

1. Issue the **svcinfolnode** CLI command to list the nodes that are currently part of the cluster and determine the I/O group for which to add the node.

The following is an example of the output that is displayed:

```
svcinfolnode -delim :
id:name:UPS_serial_number:WWNN:status:IO_group_id:
IO_group_name:config_node:UPS_unique_id
1:node1:10L3ASH:500507680100002C:online:0:io_grp0:yes:202378101C0D18D8
....
```

2. Issue the **svcinfolnodecandidate** CLI command to list nodes that are not assigned to a cluster and to verify that when a second node is added to an I/O group, it is attached to a different UPS.

The following is an example of the output that is displayed:

```
svcinfolnodecandidate -delim :
id:panel_name:UPS_serial_number:UPS_unique_id
5005076801000001:000341:10L3ASH:202378101C0D18D8
5005076801000009:000237:10L3ANF:202378101C0D1796
50050768010000F4:001245:10L3ANF:202378101C0D1796
....
```

3. Issue the **svctask addnode** CLI command to add a node to the cluster.

Attention: If you are re-adding a node to the cluster, ensure that you are adding the node to the same I/O group from which it was removed. 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.

Important: Each node in an I/O group must be attached to a different UPS.

The following is an example of the CLI command you can issue to add a node to the cluster using the panel name parameter:

```
svctask addnode -panelname 000237
-iogrp io_grp0 -name group1node2
```

Where *000237* is the panel name of the node, *io_grp0* is the name of the I/O group that you are adding the node to and *group1node2* is the name that you want to give to the node.

The following is an example of the CLI command you can issue to add a node to the cluster using the WWNN parameter:

```
svctask addnode -wwnodename 5005076801000001
-iogrp io_grp1 -name group2node2
```

Where *5005076801000001* is the WWNN of the node, *io_grp1* is the name of the I/O group that you are adding the node to and *group2node2* is the name that you want to give to the node.

You can specify a name for the node or use the default name.

If you do not specify the name for the node, the node can later be identified by using the front panel name, which is printed on a label on the front of the SAN Volume Controller, or by using the WWNN of that node .

Record the following information for the new node:

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

4. Issue the **svctask chnode** CLI command to change the default name of a node to a name that can make it easy to identify in the cluster. The following is an example of the CLI command that you can issue:

```
svctask chnode -name group1node1 node1
```

Where *group1node1* is the new name for the node and *node1* is the default name that was assigned to the node.

5. Issue the **svcinfolnode** CLI command to verify the final configuration.

The following is an example of the output that is displayed:

```
svcinfolnode -delim :  
  
id:name:UPS_serial_number:WWNN:status:I0_group_id:  
I0_group_name:config_node:UPS_unique_id  
1:group1node1:10L3ASH:500507680100002C:online:0:io_grp0:yes:202378101C0D18D8  
2:group1node2:10L3ANF:5005076801000009:online:0:io_grp0:no:202378101C0D1796  
3:group2node1:10L3ASH:5005076801000001:online:1:io_grp1:no:202378101C0D18D8  
4:group2node2:10L3ANF:50050768010000F4:online:1:io_grp1:no:202378101C0D1796  
....
```

Note: If this command is issued quickly after you have added nodes to the cluster, the status of the nodes might be adding. The status is shown as adding if the process of adding the nodes to the cluster is still in progress. You do not have to wait for the status of all the nodes to be online before you continue with the configuration process.

Remember: Record the following information:

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

The nodes have been added to the cluster.

Related tasks

“Registering the SAN Volume Controller ports with the EMC CLARiiON” on page 240

You must register the SAN Volume Controller ports with an EMC CLARiiON controller if Access Logix is installed.

Displaying node properties using the CLI

You can use the command-line interface (CLI) to display node properties.

Perform the following steps to display the node properties:

1. Issue the **svcinfolnode** CLI command to display a concise list of nodes in the cluster.

The following is an example of the CLI command you can issue to list the nodes in the cluster:

```
svcinfolnode -delim :
```

The following is an example of the output that is displayed:

```
id:name:UPS_serial_number:WWNN:status:IO_group_id:
IO_group_name:config_node:UPS_unique_id
1:group1node1:10L3ASH:500507680100002C:online:0:io_grp0:yes:202378101C0D18D8
2:group1node2:10L3ANF:5005076801000009:online:0:io_grp0:no:202378101C0D1796
3:group2node1:10L3ASH:5005076801000001:online:1:io_grp1:no:202378101C0D18D8
4:group2node2:10L3ANF:50050768010000F4:online:1:io_grp1:no:202378101C0D1796
```

2. Issue the **svcinfo lsnode** CLI command and specify the node ID or name of the node that you want to receive detailed output.

The following is an example of the CLI command you can issue to list detailed output for a node in the cluster:

```
svcinfo lsnode -delim : group1_node1
```

Where *group1_node1* is the name of the node for which you want to view detailed output.

Where *1* is the name of the node for which you want to view detailed output.

The following is an example of the output that is displayed:

```
id:1
name:group1node1
UPS_serial_number:10L3ASH
WWNN:500507680100002C
status:online
IO_group_id:0
IO_group_name:io_grp0
partner_node_id:2
partner_node_name:group1node2
config_node:yes
UPS_unique_id:202378101C0D18D8
port_id:500507680110002C
port_status:active
port_id:500507680120002C
port_status:active
port_id:500507680130002C
port_status:active
port_id:500507680140003C
port_status:active
```

The output includes the following information:

- Node ID
- Node name
- Worldwide node name (WWNN)
- Details about the uninterruptible power supply (UPS) to which the node is attached
- Details about the input/output (I/O) group of which the node is a member
- Detailed fibre-channel port status information.

Discovering MDisks using the CLI

You can use the command-line interface (CLI) to discover managed disks (MDisks).

When back-end controllers are added to the fibre-channel SAN and are included in the same switch zone as a SAN Volume Controller cluster, the cluster automatically discovers the back-end controller and integrates the controller to determine the storage that is presented to the SAN Volume Controller nodes. The SCSI logical units (LUs) that are presented by the back-end controller are displayed as unmanaged MDisks. However, if the configuration of the back-end controller is modified after this has occurred, the SAN Volume Controller cluster might be

unaware of these configuration changes. You can request that the SAN Volume Controller cluster rescans the fibre-channel SAN to update the list of unmanaged MDisks.

Note: The automatic discovery that is performed by SAN Volume Controller cluster does not write anything to an unmanaged MDisk. You must instruct the SAN Volume Controller cluster to add an MDisk to an MDisk group or use an MDisk to create an image mode virtual disk (VDisk).

Perform the following steps to discover MDisks:

1. Issue the **svctask detectmdisk** CLI command to manually scan the fibre-channel network. The scan discovers any new MDisks that might have been added to the cluster and rebalances MDisk access across the available controller device ports.
2. Issue the **svcinfolismdiskcandidate** CLI command to show the unmanaged MDisks. These MDisks have not been assigned to an MDisk group.
 - If you want to view all of the MDisks, issue the **svcinfolismdisk** CLI command.

You have now seen that the back-end controllers and switches have been set up correctly and that the SAN Volume Controller cluster recognizes the storage that is presented by the back-end controller.

The following example describes a scenario where a single back-end controller is presenting eight SCSI LUs to the SAN Volume Controller cluster:

1. Issue `svctask detectmdisk`.
2. Issue `svcinfolismdiskcandidate`.

The following output is displayed:

```
id
0
1
2
3
4
5
6
7
```

3. Issue `svcinfolismdisk -delim : -filtervalue mode=unmanaged`

The following output is displayed:

```
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
```

Creating MDisk groups using the CLI

You can use the command-line interface (CLI) to create a managed disk (MDisk) group.

Attention: If you add an MDisk to an MDisk group as an MDisk, any data on the MDisk is lost. If you want to keep the data on an MDisk (for example because you want to import storage that was previously not managed by a SAN Volume Controller), you must create image mode virtual disks (VDisks) instead.

Assume that the cluster has been set up and that a back-end controller has been configured to present new storage to the SAN Volume Controller.

Before you create MDisk groups, consider how you plan to use your storage. The SAN Volume Controller allows you to create up to 128 MDisk groups and to add up to 128 MDisks to an MDisk group. Consider the following factors as you decide how many MDisk groups to create:

- A VDisk can only be created using the storage from one MDisk group. Therefore, if you create small MDisk groups, you might lose the benefits that are provided by virtualization, namely more efficient management of free space and a more evenly distributed workload for better performance.
- If any MDisk in an MDisk group goes offline, all the VDIs in the MDisk group go offline. Therefore you might want to consider using different MDisk groups for different back-end controllers or for different applications.
- If you anticipate regularly adding and removing back-end controllers or storage, this task is made simpler by grouping all the MDisks that are presented by a back-end controller into one MDisk group.
- All the MDisks in an MDisk group should have similar levels of performance or reliability, or both. If an MDisk group contains MDisks with different levels of performance, the performance of the VDIs in this group is limited by the performance of the slowest MDisk. If an MDisk group contains MDisks with different levels of reliability, the reliability of the VDIs in this group is that of the least reliable MDisk in the group.

Even with the best planning, circumstances can change and you must reconfigure your MDisk groups after they have been created. The data migration facilities that are provided by the SAN Volume Controller 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 can have different extent sizes; however, this places restrictions on the use of data migration. The choice of extent size affects the total amount of storage that a SAN Volume Controller cluster can manage. Table 12 shows the maximum amount of storage that can be managed by a cluster for each extent size. Because the SAN Volume Controller allocates a whole number of extents to each VDisk that is created, using a larger extent size might increase the amount of storage that is wasted at the end of each VDisk. Larger extent sizes also reduces the ability of the SAN Volume Controller to distribute sequential I/O workloads across many MDisks and hence can reduce the performance benefits of virtualization.

Table 12. Extent size

Extent Size	Maximum storage capacity of cluster
16 MB	64 TB
32 MB	128 TB
64 MB	256 TB

Table 12. Extent size (continued)

128 MB	512 TB
256 MB	1 PB
512 MB	2 PB

Important: You can specify different extent sizes for different MDisk groups; however, you cannot migrate VDisks between MDisk groups with different extent sizes. If possible, create all your MDisk groups with the same extent size.

Perform the following steps to create an MDisk group:

Issue the **svctask mkmdiskgrp** CLI command to create an MDisk group.

The following is an example of the CLI command you can issue to create an MDisk group:

```
svctask mkmdiskgrp -name maindiskgroup -ext 32  
-mdisk mdsk0:mdsk1:mdsk2:mdsk3
```

Where *maindiskgroup* is the name of the MDisk group that you want to create, 32 MB is the size of the extent you want to use, and *mdsk0*, *mdsk1*, *mdsk2*, *mdsk3* are the names of the four MDisks that you want to add to the group.

You created and added MDisks to an MDisk group.

The following example provides a scenario where you want to create an MDisk group, but you do not have any MDisks available to add to the group. You plan to add the MDisks at a later time.

1. Issue `svctask mkmdiskgrp -name bkpmdiskgroup -ext 32`.
Where *bkpmdiskgroup* is the name of the MDisk group that you want to create and 32 MB is the size of the extent you want to use.
2. You find four MDisks that you want to add to the MDisk group.
3. Issue `svctask addmdisk -mdisk mdsk4:mdsk5:mdsk6:mdsk7 bkpmdiskgroup`.
Where *mdsk4*, *mdsk5*, *mdsk6*, *mdsk7* are the names of the MDisks that you want to add to the MDisk group and *bkpmdiskgroup* is the name of the MDisk group for which you want to add MDisks.

You used the **svctask mkmdiskgrp** CLI command to create the MDisk group *bkpmdiskgroup* and later used the **svctask addmdisk** CLI command to add *mdsk4*, *mdsk5*, *mdsk6*, *mdsk7* to the MDisk group.

Adding MDisks to MDisk groups using the CLI

You can use the command-line interface (CLI) to add managed disks (MDisks) to MDisk groups.

The MDisks must be in unmanaged mode. Disks that already belong to an MDisk group cannot be added to another MDisk group until they have been deleted from their current MDisk group. You can delete an MDisk from an MDisk group under the following circumstances:

- If the MDisk does not contain any extents in use by a virtual disk (VDisk)
- If you can first migrate the extents in use onto other free extents within the group

Important: Do not add the MDisk using this procedure if you want to make an image mode VDisk with it.

Perform the following steps to add MDisks to MDisk groups:

1. Issue the **svcinfolsmdiskgrp** CLI command to list the existing MDisk groups. The following is an example of the CLI command you can issue to list the existing MDisk groups:

```
svcinfolsmdiskgrp -delim :
```

The following is an example of the output that is displayed:

```
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
```

2. Issue the **svctask addmdisk** CLI command to add MDisks to the MDisk group. The following is an example of the CLI command you can issue to add MDisks to an MDisk group:

```
svctask addmdisk -mdisk mdisk4:mdisk5:mdisk6:mdisk7 bkpmdiskgroup
```

Where *mdisk4:mdisk5:mdisk6:mdisk7* are the names of the MDisks that you want to add to the MDisk group and *bkpmdiskgroup* is the name of the MDisk group for which you want to add the MDisks.

Creating VDisks

You can use the command-line interface (CLI) to create a virtual disk (VDisk).

This task assumes that the cluster has been setup and that you have created managed disk (MDisk) groups. You must establish an empty MDisk group to hold the MDisks used for image mode VDisks.

Note: If you want to keep the data on an MDisk, you should instead create image mode VDisks. This task describes how to create a VDisk with striped virtualization.

Perform the following steps to create VDisks:

1. Issue the **svcinfolsmdiskgrp** CLI command to list the available MDisk groups and the amount of free storage in each group.

The following is an example of the CLI command you can issue to list MDisk groups:

```
svcinfolsmdiskgrp -delim :
```

The following is an example of the output that is displayed:

```
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
```

2. Decide which MDisk group you want to provide the storage for the VDisk.
3. Issue the **svcinfolsiogrp** CLI command to show the I/O groups and the number of VDisks 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. You can use FlashCopy to make copies of VDisks regardless of whether the source and

destination VDisk are in the same I/O group. If you plan to use intracluster Metro Mirror, make sure that both the master and auxiliary VDisk are in the same I/O group.

The following is an example of the CLI command you can issue to list I/O groups:

```
svcinfolsiogrp -delim :
```

The following is an example of the output that is displayed:

```
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. Decide which I/O group you want to assign the VDisk to. This determines which SAN Volume Controller nodes in the cluster process the I/O requests from the host systems. If you have more than one I/O group, make sure you distribute the VDIs between the I/O groups so that the I/O workload is shared evenly between all SAN Volume Controller nodes.

5. Issue the **svctask mkvdisk** CLI command to create a VDisk.

The following is an example of the CLI command you can issue to create a VDisk using the I/O group ID and MDisk group ID:

```
svctask mkvdisk -name mainvdisk1 -iogrp 0
  -mdiskgrp 0 -vtype striped -size 256 -unit gb
```

Where *mainvdisk1* is the name that you want to call the VDisk, *0* is the ID of the I/O group that want the VDisk to use, *0* is the ID of the MDisk group that you want the VDisk to use, and *256* is the capacity of the VDisk.

The following is an example of the CLI command you can issue to create a VDisk using the I/O group and MDisk group name:

```
svctask mkvdisk -name bkpvdisk1 -iogrp io_grp1
  -mdiskgrp bkpmdiskgroup -vtype striped -size 256 -unit gb
```

Where *bkpvdisk1* is the name that you want to call the VDisk, *io_grp1* is the name of the I/O group that want the VDisk to use, *bkpmdiskgroup* is the name of the MDisk group that you want the VDisk to use, and *256* is the capacity of the VDisk.

6. Issue the **svcinfolsvdisk** CLI command to list all the VDIs that have been created.

The following is an example of the CLI command you can issue to list VDIs:

```
svcinfolsvdisk -delim :
```

The following is an example of the output that is displayed:

```
id:name:IO_group_id:IO_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:bkpvdisk1:1:io_grp1:online:1:bkpmdiskgroup:
512.0GB:striped:::
```

Creating host objects using the CLI

You can use command-line interface (CLI) to create host objects.

Perform the following steps to create host objects:

1. Issue the **svctask mkhost** CLI command to create a logical host object. Assign your worldwide port name (WWPN) for the host bus adapters (HBAs) in the hosts.

The following is an example of the CLI command you can issue to create a host:

```
svctask mkhost -name demohost1 -hbawwpn 210100e08b251dd4
```

Where *demohost1* is the name of the host and *210100e08b251dd4* is the WWPN of the HBA.

2. Issue the **svctask addhostport** CLI command to add ports to the host.

The following is an example of the CLI command you can issue to add a port to the host:

```
svctask mkhost -name demohost1 -hbawwpn 210100e08b251dd5
```

This command adds another HBA WWPN called *210100e08b251dd5* to the host that was created in step 1.

The following is an example of the CLI command you can issue to create a second host:

```
svctask mkhost -hbawwpn 210100e08b251dd6:210100e08b251dd7 -name demohost2
```

Where *demohost2* is the name of the second host and *210100e08b251dd6*, *210100e08b251dd7* are the HBA's WWPNs.

Note: If you were to add a host with a faulty WWPN, or the WWPN had been assigned to the wrong host, you must issue the **svctask addhostport** CLI command to add that same host with the correct WWPN, then issue the **svctask rmhostport** CLI 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 must issue a command that is similar to the following:

```
svctask addhostport -hbawwpn 210100e08b251dd4 demohost1
```

This adds the host called *demohost1* with the WWPN, *210100e08b251dd4*. You must issue the **svctask rmhostport** CLI command to delete the host with the WWPN that had stopped working.

The following is an example of the CLI command you must issue to delete the host:

```
svctask rmhostport -hbawwpn 210100e08b251dd5 demohost1
```

The **svctask addhostport** and **svctask rmhostport** CLI commands have allowed you to delete the host using the WWPN *210100e08b251dd5*, and then add the deleted host using the WWPN *210100e08b251dd4*.

Creating VDisk-to-host mappings using the CLI

You can use the command-line interface (CLI) to create virtual disk (VDisk)-to-host mappings.

Perform the following steps to create VDisk-to-host mappings:

Issue the **svctask mkvdiskhostmap** CLI command to create VDisk-to-host mappings.

The following is an example of the CLI command you can issue to create VDisk-to-host mappings:

```
svctask mkvdiskhostmap -host demohost1 mainvdisk1
```

Where *demohost1* is the name of the host and *mainvdisk1* is the name of the VDisk.

Creating FlashCopy mappings using the CLI

You can use the command-line interface (CLI) to create FlashCopy mappings.

The FlashCopy mapping specifies the source and destination virtual disk (VDisk). The destination must be identical in size to the source, or the mapping fails. The source and destination cannot be in an existing mapping. That is, a VDisk 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.

Perform the following steps to create FlashCopy mappings:

1. Issue the **svcinfolsvdisk -bytes** CLI command to find the exact size of the source VDisk that you want to create a target disk of the same size.
2. Issue the **svctask mkfcmap** CLI command to create a FlashCopy mapping.

The following is an example of the CLI command you can issue to create FlashCopy mappings with the copy rate parameter:

```
svctask mkfcmap -source mainvdisk1 -target bkpvdisk1  
-name main1copy -copyrate 75
```

Where *mainvdisk1* is the name of the source VDisk, *bkpvdisk1* is the name of the target VDisk, *main1copy* is the name that you want to call the FlashCopy and *75* is the priority that you want to give the background copy rate.

The following is an example of the CLI command you can issue to create FlashCopy mappings without the copy rate parameter:

```
svctask mkfcmap -source mainvdisk2 -target bkpvdisk2  
-name main2copy
```

Where *mainvdisk2* is the name of the source VDisk, *bkpvdisk2* is the name of the target VDisk, *main2copy* is the name that you want to call the FlashCopy.

Note: The default copy rate of 50 is used if you do not specify a copy rate.

3. Issue the **svcinfolsfcmmap** CLI command to check the attributes of the FlashCopy mappings that have been created:

The following is an example of the CLI command you can issue to view the attributes of the FlashCopy mappings:

```
svcinfolsfcmmap -delim :
```

The following is an example of the output that is displayed:

```

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

```

Creating a FlashCopy consistency group and adding mappings using the CLI

You can use the command-line interface (CLI) to create and add mappings to a FlashCopy consistency group.

If you have created several FlashCopy mappings for a group of virtual disks (VDisks) that contain elements of data for the same application, it can be convenient to assign these mappings to a single FlashCopy consistency group. You can then issue a single prepare or trigger command for the whole group. For example, you can copy all of the files for a database at the same time.

Perform the following steps to create a FlashCopy mappings:

1. Issue the **svctask mkfcconsistgrp** CLI command to create a FlashCopy consistency group.

The following is an example of the CLI command you can issue to create a FlashCopy consistency group:

```
svctask mkfcconsistgrp -name maintobkpfcopy
```

Where *maintobkpfcopy* is the name that you want to call the FlashCopy consistency group.

2. Issue the **svcinfolsfconsistgrp** CLI command to display the attributes of the group that you have created.

The following is an example of the CLI command you can issue to display the attributes of a FlashCopy consistency group:

```
svcinfolsfconsistgrp -delim :
```

The following is an example of the output that is displayed:

```

id:name:status
1:maintobkpfcopy:idle_copied

```

3. Issue the **svctask chfcmap** CLI command to add FlashCopy mappings to the FlashCopy consistency group:

The following are examples of the CLI commands you can issue to add FlashCopy mappings to the FlashCopy consistency group:

```
svctask chfcmap -consistgrp maintobkpfcopy main1copy
svctask chfcmap -consistgrp maintobkpfcopy main2copy
```

Where *maintobkpfcopy* is the name of the FlashCopy consistence group and *main1copy*, *main2copy* are the names of the FlashCopy mappings.

4. Issue the **svcinfolsfcmmap** CLI command to display the new attributes of the FlashCopy mappings.

The following is an example of the output that is displayed:

```

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

```

5. Issue the **svcinfn lsfcconsistgrp** CLI command to display the detailed attributes of the group.

The following is an example of the CLI command you can issue to display detailed attributes:

```
svcinfn lsfcconsistgrp -delim : maintobkpfcopy
```

Where *maintobkpfcopy* is the name of the FlashCopy consistency group.

The following is an example of the output that is displayed:

```
id:1
name:maintobkpfcopy
status:idle_copied
FC_mapping_id:0
FC_mapping_name:main1copy
FC_mapping_id:1
FC_mapping_name:main2copy
```

Preparing and triggering a FlashCopy mapping using the CLI

Before you start the FlashCopy process using the command-line interface (CLI), you must prepare and trigger a FlashCopy mapping.

Triggering a FlashCopy mapping creates a point-in-time copy of the data on the source virtual disk (Vdisk) and writes it to the target Vdisk for the mapping.

Perform the following steps to prepare and trigger a FlashCopy mapping:

1. Issue the **svctask prestartfcmap** CLI command to prepare the FlashCopy mapping.

The following is an example of the CLI command you can issue to prepare a FlashCopy mapping:

```
svctask prestartfcmap main1copy
```

Where *main1copy* is the name of the FlashCopy mapping.

The mapping enters the preparing state and moves to the prepared state when it is ready.

2. Issue the **svcinfn lsfcmmap** CLI command to check the state of the mapping.

The following is an example of the output that is displayed:

```
svcinfn lsfcmmap -delim :
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:::prepared:0:50
```

3. Issue the **svctask startfcmap** CLI command to start (trigger) the FlashCopy mapping.

The following is an example of the CLI command you can issue to start the FlashCopy mapping:

```
svctask startfcmap main1copy
```

Where *main1copy* is the name of the FlashCopy mapping.

4. Issue the **svcinfn lsfcmmapprogress** CLI command to check the progress of the FlashCopy mapping.

The following is an example of the output that is displayed:

```
svcinfn lsfcmmapprogress -delim :
id:progress
0:47
```

You have created a point-in-time copy of the data on a source VDisk and written that data to a target VDisk. The data on the target VDisk is only recognized by the hosts that are mapped to it.

Preparing and triggering a FlashCopy consistency group using the CLI

You can use the command-line interface (CLI) to prepare and trigger a FlashCopy consistency group to start the FlashCopy process.

Starting the FlashCopy process creates a point-in-time copy of the data on the source virtual disk (VDisk) and writes it to the target VDisk for each mapping in the group. When you have assigned several mappings to a FlashCopy consistency group, you only have to issue a single prepare or trigger command for the whole group to prepare or trigger all the mappings at once.

Perform the following steps to prepare and trigger a FlashCopy consistency group:

1. Issue the **svctask prestartfcconsistgrp** CLI command to prepare the FlashCopy consistency group before the copy process can be started (triggered).

Remember: You only have to issue a single prepare command for the whole group in order to prepare all the mappings at once.

The following is an example of the CLI command you can issue to prepare the FlashCopy consistency group:

```
svctask prestartfcconsistgrp maintobkpfcopy
```

Where *maintobkpfcopy* is the name of the FlashCopy consistency group. The group enters the preparing state, and then moves to the prepared state when it is ready.

2. Issue the **svcinfolsfconsistgrp** command to check the status of the FlashCopy consistency group.

The following is an example of the CLI command you can issue to check the status of the FlashCopy consistency group:

```
svcinfolsfconsistgrp -delim :
```

The following is an example of the output that is displayed:

```
id:name:status
1:maintobkpfcopy:prepared
```

3. Issue the **svctask startfcconsistgrp** CLI command to start (trigger) the FlashCopy consistency group to make the copy.

Remember: You only have to issue a single start command for the whole group in order to trigger all the mappings at once.

The following is an example of the CLI command you can issue to start the FlashCopy consistency group mappings:

```
svctask startfcconsistgrp maintobkpfcopy
```

Where *maintobkpfcopy* is the name of the FlashCopy consistency group. The FlashCopy consistency group enters the copying state, and then returns to the `idle_copied` state when complete.

4. Issue the **svcinfolsfconsistgrp** command to check the status of the FlashCopy consistency group.

The following is an example of the CLI command you can issue to check the status of the FlashCopy consistency group:

```
svcinfolsfcconsistgrp -delim : maintobkpfcopy
```

Where *maintobkpfcopy* is the name of the FlashCopy consistency group.

The following is an example of the output that is displayed when the process is still copying:

```
id:name:status
1:maintobkpfcopy:copying
```

The following is an example of the output that is displayed when the process has finished copying:

```
id:1
name:maintobkpfcopy
state:idle_copied
FC_mapping_id:0
FC_mapping_name:main1copy
FC_mapping_id:1
FC_mapping_name:main2copy
```

Determining the WWPNs of a node using the CLI

You can determine the worldwide port names (WWPNs) of a node using the command-line interface (CLI).

Perform the following steps to determine the WWPNs of a node:

1. Issue the **svcinfolnode** CLI command to list the nodes in the cluster.
2. Record the name or ID of the node for which you want to determine the WWPNs.
3. Issue the **svcinfolnode** CLI command and specify the node name or ID that was recorded in step 2.

The following is an example of the CLI command you can issue:

```
svcinfolnode node1
```

Where *node1* is the name of the node for which you want to determine the WWPNs.

4. Record the four port IDs (WWPNs).

Determining the VDisk name from the device identifier on the host

You can use the command-line interface (CLI) to determine the virtual disk (VDisk) name from the device identifier on the host.

Each VDisk that is exported by the SAN Volume Controller is assigned a unique device identifier. The device identifier uniquely identifies the VDisk and can be used to determine which VDisk corresponds to the volume that the host sees.

Perform the following steps to determine the VDisk name from the device identifier:

1. Find the device identifier. For example, if you are using the subsystem device driver (SDD), the disk identifier is referred to as the virtual path (vpath) number. You can issue the following command to find the vpath serial number:

```
datapath query device
```

For other multipathing drivers, refer to the documentation that is provided with your multipathing driver to determine the device identifier.

2. Find the host object that is defined to the SAN Volume Controller and corresponds with the host that you are working with.

- a. Find the worldwide port numbers (WWPNs) by looking at the device definitions that are stored by your operating system. For example, on AIX the WWPNs are in the ODM and if you use Windows you have to go into the HBA Bios.
 - b. Verify which host object is defined to the SAN Volume Controller for which these ports belong. The ports are stored as part of the detailed view, so you must list each host by issuing the following CLI command:

```
svcinfo lshost name/id
```

Where *name/id* is the name or ID of the host.
 - c. Check for matching WWPNs.
3. Issue the following command to list the VDisk-to-host mappings:

```
svcinfo lshostvdiskmap hostname
```

Where *hostname* is the name of the host.
 4. Find the VDisk UID that matches the device identifier and record the VDisk name or ID.

Determining the host that a VDisk is mapped to

You can determine the host that a virtual disk (VDisk) is mapped to using the command-line interface (CLI).

Perform the following steps to determine the host that the VDisk is mapped to:

1. Find the VDisk name or ID that you want to check.
2. Issue the following CLI command to list the hosts that this VDisk is mapped to:

```
svcinfo lsvdiskhostmap vdiskname/id
```

Where *vdiskname/id* is the name or ID of the VDisk.
3. Find 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

You can determine the relationship between virtual disks (VDIs) and managed disks (MDIs) using the command-line interface (CLI).

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

Choose one of the following options to determine the relationship between VDIs and MDIs:

- If you want to determine the relationship between VDIs and MDIs, issue the following CLI command:

```
svcinfo lsvdiskmember vdiskname/id
```

Where *vdiskname/id* is the name or ID of the VDisk.

This CLI command returns a list of the IDs that correspond to the MDIs that comprise the VDisk.

- If you want to determine the relationship between MDIs and VDIs, issue the following CLI command:

```
svcinfo lsmdiskmember mdiskname/id
```

Where *mdiskname/id* is the name or ID of the MDI.

This CLI command returns a list of IDs that correspond to the VDisks that are using this MDisk.

- If you want to determine the relationship between VDisks and MDisks and the number of extents that are provided by each MDisk, issue the following CLI command:

```
svcinfolsvdiskextent vdiskname/id
```

Where *vdiskname/id* is the name or ID of the VDisk.

This CLI command returns a table of MDisk IDs and the corresponding number of extents that each MDisk provides as storage for the given VDisk.

- If you want to determine the relationship between MDisks and VDisks and the number of extents that are used by each VDisk, issue the following CLI command:

```
svcinfolsmdiskextent mdiskname/id
```

Where *mdiskname/id* is the name or ID of the MDisk.

This CLI command 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

You can determine the relationship between managed disks (MDisks) and RAID arrays or LUNs using the command-line interface (CLI).

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

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

1. Issue the following command to display a detailed view of the MDisk:

```
svcinfolsmdisk mdiskname
```

Where *mdiskname* is the name of the MDisk for which you want to display a detailed view.

2. Record the controller name or controller ID and the controller LUN number.
3. Issue the following command to display a detailed view of the controller:

```
svcinfolcontroller controllername
```

Where *controllername* is the name of the controller that you recorded in step 2.

4. Record the vendor ID, product ID, and WWNN. You can use this information to determine what is being presented to the MDisk.
5. From the native user interface for the given controller, list the LUNs it is presenting and match the LUN number with that noted in step 1. This tells you the exact RAID array or partition that corresponds with the MDisk.

Increasing the size of your cluster using the CLI

You can increase the size of a cluster using the command-line interface (CLI).

You can increase throughput by adding more nodes to the cluster. The nodes must be added in pairs and assigned to a new I/O group.

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

1. Add a node to your cluster and repeat this step for the second node.
2. If you want to balance the load between the existing I/O groups and the new I/O groups, you can migrate your virtual disks (VDisks) to new I/O groups. Repeat this step for all VDisks that you want to assign to the new I/O group.

Adding a node to increase the size of a cluster using the CLI

You can add a node to increase the size of a cluster using the command-line interface (CLI).

Attention: If you are adding a node that was previously removed from a cluster, ensure that you are adding the node to the same I/O group from which it was removed. Failure to do this can result in data corruption. If you do not know the I/O group name or ID that it was removed from, contact IBM Service to add the node to the cluster without corrupting data.

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

1. Issue the following command to verify that the node is detected on the fabric and to obtain the worldwide node names (WWNNs) of the nodes on the cluster:

```
svcinfolnodecandidate
```

2. Record the WWNN.

3. Issue the following command to determine the I/O group where the node should be added:

```
svcinfolsiogrp
```

4. Record the name or ID of the first I/O group that has a node count of zero (0). You will need the ID for the next step.

5. Record the following information for future reference:

- Node serial number.
- Worldwide node name.
- All of the worldwide port names.
- The name or ID of the I/O group that contains the node.

6. Issue the following command to add the node to the cluster:

```
svctask addnode -wwnodename WWNN -iogrp newiogrpname/id [-name newnodename]
```

Where *WWNN* is the WWNN of the node, *newiogrpname/id* is the name or ID of the I/O group that you want to add the node to and *newnodename* is the name that you want to assign to the node.

7. Issue the following command to verify that the node is online:

```
svcinfolnode
```

If the disk controller uses mapping to present RAID arrays or partitions to the cluster and the WWNNs or the worldwide port names have changed, you must modify the port groups that belong to the cluster.

Migrating a VDisk to a new I/O group using the CLI

You can use the command-line interface (CLI) to migrate a virtual disk (VDisk) to a new I/O group to increase the size of your cluster.

You can migrate a VDisk to a new I/O group to manually balance the workload across the nodes in the cluster. However, you might end up with a pair of nodes that are overworked and another pair that are underworked. Follow this procedure to migrate a single VDisk to a new I/O group. Repeat for other VDisks as required.

Attention: This is a disruptive procedure. Access to the VDisk is lost while you follow this procedure. Under no circumstances should VDisks be moved to an offline I/O group. To avoid data loss, you must ensure that the I/O group is online before you move the VDisks.

Perform the following steps to migrate a single VDisk:

1. Quiesce all I/O operations for the VDisk. You might have to determine the hosts that are using this VDisk in advance.
2. Before migrating the VDisk, it is essential that for each device identifier that is presented by the VDisk you intend to move, the subsystem device driver (SDD) or other multipathing driver configuration is updated to remove the device identifiers. Failure to do this can result in data corruption. See the *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide* or the documentation that is provided with your multipathing driver for details about how to dynamically reconfigure device identifiers for the given host operating system.
3. 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.

- a. Find the **FC_id** and **RC_id** fields. If these are not blank, the VDisk is part of a mapping or relationship.
 - b. Stop or delete any FlashCopy mappings or Metro Mirror relationships that use this VDisk.
4. Issue the following command to migrate the VDisk:

```
svctask chvdisk -iogrp newiogrpname/id vdiskname/id
```

Where *newiogrpname/id* is the name or ID of the I/O group where you want to migrate the VDisk and *vdiskname/id* is the name or ID of the VDisk that you want to migrate.
 5. Discover the new device identifiers and check that each device identifier presents the correct number of paths. See the *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide* or the documentation that is provided with your multipathing driver for details about how to discover device identifiers for the given host operating system.

Replacing a faulty node in the cluster using the CLI

You can use the command-line interface (CLI) to replace a faulty node in the cluster.

Before you attempt to replace a faulty node with a spare node you must ensure that you meet the following requirements:

- SAN Volume Controller version 1.1.1 or later is installed on the cluster and on the spare node.
- You know the name of the cluster that contains the faulty node.
- A spare node is installed in the same rack as the cluster that contains the faulty node.
- You must make a record of the last five characters of the original worldwide node name (WWNN) of the spare node. You will need this information, if and when, you want to stop using this node as a spare node.

If a node fails, the cluster continues to operate with degraded performance until the faulty node is repaired. If the repair operation takes an unacceptable amount of time, it is useful to replace the faulty node with a spare node. However, the appropriate procedures must be followed and precautions must be taken so you do *not* interrupt I/O operations and compromise the integrity of your data.

The following table describes the changes that are made to your configuration when you replace a faulty node in the cluster:

Node attributes	Description												
Front panel ID	This is the number that is printed on the front of the node and is used to select the node that is added to a cluster. This number changes.												
Node ID	This is the ID that is assigned to the node. A new node ID is assigned each time a node is added to a cluster; the node name remains the same following service activity on the cluster. You can use the node ID or the node name to perform management tasks on the cluster. However, if you are using scripts to perform those tasks, use the node name rather than the node ID. This ID changes.												
Node name	This is the name that is assigned to the node. If you do not specify a name, the SAN Volume Controller assigns a default name. The SAN Volume Controller creates a new default name each time a node is added to a cluster. If you choose to assign your own names, you must type the node name on the Adding a node to a cluster panel. If you are using scripts to perform management tasks on the cluster and those scripts use the node name, you can avoid the need to make changes to the scripts by assigning the original name of the node to a spare node. This number might change.												
Worldwide node name	This is the WWNN that is assigned to the node. The WWNN is used to uniquely identify the node and the fibre-channel ports. The WWNN of the spare node changes to that of the faulty node. The node replacement procedures must be followed exactly to avoid any duplication of WWNNs. This name changes.												
Worldwide port names	<p>These are the WWPNs that are assigned to the node. WWPNs are derived from the WWNN that is written to the spare node as part of this procedure. For example, if the WWNN for a node is 50050768010000F6, the four WWPNs for this node are derived as follows:</p> <table border="0"> <tr> <td>WWNN</td> <td>50050768010000F6</td> </tr> <tr> <td>WWNN displayed on front panel</td> <td>000F6</td> </tr> <tr> <td>WWPN Port 1</td> <td>50050768014000F6</td> </tr> <tr> <td>WWPN Port 2</td> <td>50050768013000F6</td> </tr> <tr> <td>WWPN Port 3</td> <td>50050768011000F6</td> </tr> <tr> <td>WWPN Port 4</td> <td>50050768012000F6</td> </tr> </table> <p>These names do not change.</p>	WWNN	50050768010000F6	WWNN displayed on front panel	000F6	WWPN Port 1	50050768014000F6	WWPN Port 2	50050768013000F6	WWPN Port 3	50050768011000F6	WWPN Port 4	50050768012000F6
WWNN	50050768010000F6												
WWNN displayed on front panel	000F6												
WWPN Port 1	50050768014000F6												
WWPN Port 2	50050768013000F6												
WWPN Port 3	50050768011000F6												
WWPN Port 4	50050768012000F6												

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

1. Verify the name and ID of the node that you want to replace.

Perform the following step to verify the name and ID:

- a. Issue the **svcinfo lsnode** CLI command to ensure that the partner node in the I/O group is online.

- If the other node in the I/O group is offline, start Directed Maintenance Procedures (DMPs) to determine the fault.
 - If you have been directed here by the DMPs, and subsequently the partner node in the I/O group has failed, see the procedure for recovering from offline VDisks after a node or an I/O group failed.
 - If you are replacing the node for other reasons, determine the node you want to replace and ensure that the partner node in the I/O group is online.
 - If the partner node is offline, you will lose access to the VDisks that belong to this I/O group. Start the DMPs and fix the other node before proceeding to the next step.
2. Find and record the following information about the faulty node:
 - Node serial number
 - Worldwide node name
 - All of the worldwide port names
 - Name or ID of the I/O group that contains the node
 - Front panel ID
 - Uninterruptible power supply serial number
 - a. Issue the **svcinfo lsnode** CLI command to find and record the node name and I/O group name. The faulty node will be offline.
 - b. Record the following information about the faulty node:
 - Node name
 - I/O group name
 - c. Issue the following CLI command:


```
svcinfo lsnodevpd nodename
```

 Where *nodename* is the name that you recorded in step 1 on page 173.
 - d. Find the WWNN field in the output.
 - e. Record the last five characters of the WWNN.
 - f. Find the `front_panel_id` field in the output.
 - g. Record the front panel ID.
 - h. Find the `UPS_serial_number` field in the output.
 - i. Record the UPS serial number.
3. Disconnect all four fibre-channel cables from the node.

Important: Do not plug the fibre-channel cables into the spare node until the spare node is configured with the WWNN of the faulty node.
4. Connect the power and signal cables from the spare node to the uninterruptible power supply that has the serial number you recorded in step 2i.

Note: The signal cable can be plugged into any vacant position on the top row of serial connectors on the uninterruptible power supply. If no spare serial connectors are available on the uninterruptible power supply, disconnect the cables from the faulty node.
5. Power on the spare node.
6. Display the node status on the service panel. See the *IBM TotalStorage SAN Volume Controller: Service Guide* for more information.
7. Change the WWNN of the spare node to match the WWNN of the faulty node by performing the following steps:

- a. With the node status displayed on the front panel; press and hold the down button; press and release the select button; release the down button. WWNN is displayed on line 1 of the display. Line 2 of the display contains the last five characters of the WWNN.
- b. With the WWNN displayed on the service panel; press and hold the down button; press and release the select button; release the down button. The display switches into edit mode.
- c. Change the WWNN that is displayed to match the WWNN recorded in step 2e on page 174.

Note: To edit the displayed number, use the up and down buttons to increase or decrease the displayed numbers. Use the left and right buttons to move between fields.

- d. When the five characters match the number that you recorded in step 2e on page 174, press the select button twice to accept the number.
8. Connect the four fibre-channel cables that were disconnected from the faulty node to the spare node.
9. Issue the following CLI command to remove the faulty node from the cluster:
`svctask rmnode nodename/id`

Where *nodename/id* is the name or ID of the faulty node.

Remember to record the following information to avoid data corruption when this node is re-added to the cluster:

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

10. Issue the following command to add the spare node to the cluster:

```
svctask addnode -wwnodename WWNN -iogrp iogroupname/id -name nodename
```

Where *WWNN* is the WWNN of the node, *iogroupname/id* is the name or ID of the I/O group and *nodename* is the name of the node.

11. Use the subsystem device driver (SDD) or your multipathing driver management tool on the host systems to verify that all paths are now online. See the *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide* or the documentation that is provided with your multipathing driver for more information.

Attention: When the faulty node is repaired do not connect the fibre-channel cables to it. Connecting the cables can cause data corruption.

12. If you want to use the repaired node as a spare node, perform the following steps:
 - a. Display the node status on the front panel display of the node. See the *IBM TotalStorage SAN Volume Controller: Service Guide* for more information.
 - b. With the node status displayed on the front panel, press and hold the down button; press and release the select button; release the down button. WWNN is displayed on line 1 of the display. Line 2 of the display contains the last five characters of the WWNN.
 - c. With the WWNN displayed on the service panel, press and hold the down button; press and release the select button; release the down button. The display switches into edit mode.
 - d. Change the displayed number to 00000.

Note: To edit the displayed number, use the up and down buttons to increase or decrease the displayed numbers. Use the left and right buttons to move between fields.

- e. Press the select button twice to accept the number.

This node can now be used as a spare node.

Attention: Never connect a node with a WWNN of 00000 to the cluster. If this node is no longer required as a spare and is to be used for normal attachment to a cluster, you must change the WWNN to the number you recorded when a spare was created. Using any other number might cause data corruption.

Recovering from offline VDIsks using the CLI

You can recover from an offline virtual disk (VDisk) after a node or an I/O group has failed using the command-line interface (CLI).

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

Data loss scenario 1

One node in an I/O group has failed and failover has started on the second node. During the failover process, the second node in the I/O group fails before the data in the write cache is written to hard disk. The first node is successfully repaired but its hardened data is not the most recent version committed to the data store; therefore, it cannot be used. The second node is repaired or replaced and has lost its hardened data, therefore, the node has no way of recognizing that it is part of the cluster.

Perform the following steps to recover from an offline VDisk when one node has down-level hardened data and the other node has lost hardened data:

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

Data loss scenario 2

Both nodes in the I/O group have failed and have been repaired. The nodes have lost their hardened data, therefore, the nodes have no way of recognizing that they are part of the cluster.

Perform the following steps to recover from an offline VDisk when both nodes have lost their hardened data and cannot be recognized by the cluster:

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

Recovering a node and returning it to the cluster using the CLI

After a node or an I/O group fails, you can use the command-line interface (CLI) to recover a node and return it to the cluster.

Perform the following steps to recover a node and return it to the cluster:

1. Issue the following command to verify that the node is offline:

```
svcinfolsnnode
```

2. Issue the following command to remove the old instance of the offline node from the cluster:

```
svctask rmnode nodename/id
```

Where *nodename/id* is the name or ID of the node.

3. Issue the following command to verify that the node can be seen on the fabric:

```
svcinfolsnodecandidate
```

Note: Remember the worldwide node names (WWNNs) for each node because you will need them in the following step.

4. If the nodes are repaired by replacing the front panel module or a 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, you must discover the new paths and check that each device identifier presents the correct number of paths. For example, if you are using the subsystem device driver (SDD), the device identifier is referred to as the virtual path (vpath) number. See the *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide* or the documentation that is provided with your multipathing driver for details about how to dynamically reconfigure and add device identifiers for the given host operating system.
 - b. You might 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 must modify the port groups that belong to the cluster because the WWNN or WWPNs 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 from which it was removed. 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 Support 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, you must record the following information:

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

5. Issue the following command to add the node back into the cluster:

```
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 and *NODENAME* is the name of the node.

6. Issue the following command to verify that the node is online:

```
svcinfo lsnode
```

Moving offline VDIs to the recovery I/O group using the CLI

You can move offline virtual disks (VDIs) to the recovery I/O group using the command-line interface (CLI).

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

1. Issue the following CLI command to list all VDIs that are offline and belong to the I/O group:

```
svcinfo lsvdisk -filtervalue IO_group_name=  
IOGRPNAME/ID:status=offline
```

Where *IOGRPNAME/ID* is the name of the I/O group that failed.

2. Issue the following CLI command to move the VDI to the recovery I/O group:

```
svctask chvdisk -iogrp recovery_io_grp -force  
vdiskname/ID
```

Where *vdiskname/ID* is the name of one of the VDIs that are offline.

3. Repeat step 2 for all VDIs that are offline.

Moving offline VDIs to their original I/O group using the CLI

You can move offline virtual disks (VDIs) to their original I/O group using the command-line interface (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: Do not move VDIs to an offline I/O group. Ensure that the I/O group is online before you move the VDIs back to avoid any further data loss.

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

1. Issue the following command to move the VDI back into the original I/O group:

```
svctask chvdisk -iogrp IOGRPNAME/ID -force  
vdiskname/ID
```

Where *IOGRPNAME/ID* is the name or ID of the original I/O group and *vdiskname/ID* is the name or ID of the offline VDI.

2. Issue the following command to verify that the VDIs are now online:

```
svcinfo lsvdisk -filtervalue IO_group_name=  
IOGRPNAME/ID
```

Where *IOGRPNAME/ID* is the name or ID of the original I/O group.

Informing the SAN Volume Controller of changes to host HBAs using the CLI

You can use the command-line interface (CLI) to inform the SAN Volume Controller of a change to a defined host object.

Because it is sometimes necessary to replace the HBA that connects the host to the SAN, you must inform the SAN Volume Controller of the new worldwide port names (WWPNs) that this HBA contains.

Ensure that your switch is zoned correctly.

Perform the following steps to inform the SAN Volume Controller of a change to a defined host object:

1. Issue the following CLI command to list the candidate HBA ports:

```
svcinfolshbaportcandidate
```

You should see a list of the HBA ports that are available for addition to host objects. One or more of these HBA ports should correspond with the one or more WWPNs that belong to the new HBA port.

2. Locate the host object that corresponds with the host in which you have replaced the HBA. The following CLI command lists all the defined host objects:

```
svcinfolshost
```

3. Issue the following CLI command to list the WWPNs that are currently assigned to the host object:

```
svcinfolshost hostobjectname
```

Where *hostobjectname* is the name of the host object.

4. Issue the following CLI command to add the new ports to the existing host object:

```
svctask addhostport -hbawwpn one or more existing WWPNs separated by : hostobjectname/ID
```

Where *one or more existing WWPNs separated by :* is the WWPNs that are currently assigned to the host object and *hostobjectname/ID* is the name or ID of the host object.

5. Issue the following CLI command to remove the old ports from the host object:

```
svctask rmhostport -hbawwpn one or more existing WWPNs separated by : hostobjectname/ID
```

Where *one or more existing WWPNs separated by :* is the WWPNs that are currently assigned to the host object and *hostobjectname/ID* is the name or ID of the host object.

Any mappings that exist between the host object and the virtual disks (VDisks) are automatically applied to the new WWPNs. Therefore, the host sees the VDisks as the same SCSI LUNs as before.

See the *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide* or the documentation that is provided with your multipathing driver for additional information about dynamic reconfiguration.

Expanding VDisks

You can use command-line interface (CLI) or the SAN Volume Controller Console to expand a virtual disk (VDisk).

A VDisk that is not mapped to any hosts and does not contain any data can be expanded at any time. If the VDisk contains data that is in use, you can expand the VDisks if your host has an AIX, Windows 2000 or Windows 2003 operating system.

The following table provides the supported operating systems and requirements for expanding VDisks that contain data:

Operating system	Supported	Requirement
AIX	Yes	AIX version 5.2 or later
HP-UX	No	-

Operating system	Supported	Requirement
Linux	No	-
SUN Solaris	No	-
Windows NT®	No	-
Windows 2000, 2003	Yes	-

Expanding a VDisk that is mapped to an AIX host

The SAN Volume Controller supports the ability to dynamically expand the size of a virtual disk (VDisk) if the AIX host is using AIX version 5.2 or later.

The **chvg** command options provide the ability to expand the size of a physical volume that the Logical Volume Manager (LVM) uses, without interruptions to the use or availability of the system. Refer to the *AIX System Management Guide: Operating System and Devices* for more information.

Expanding a VDisk that is mapped to a Windows 2000 host using the CLI

You can use the command-line interface (CLI) to expand a virtual disk (VDisk) that is mapped to a Windows 2000 host.

VDisks that are mapped for FlashCopy or that are in Metro Mirror relationships cannot be expanded.

Ensure that you have run Windows Update and have applied all recommended updates to your system before you attempt 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:

```
svcinfo lsvdisk -bytes vdiskname
```

Where *vdiskname* is the name of the VDisk for which you want to determine the exact size.

VDisks can be expanded under Windows 2000 concurrently with I/O operations.

You can expand VDisks for the following reasons:

- To increase the available capacity 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 so that it can be used in a FlashCopy mapping or Metro Mirror relationship.

Perform the following steps to expand a VDisk that is mapped to a Windows 2000 host:

1. Issue the following CLI command to expand the VDisk:

```
svctask expandvdisksize -size disk_size -unit
b | kb | mb | gb | tb | pb vdisk_name/vdisk_id
```

Where *disk_size* is the capacity by which you want to expand the VDisk, *b / kb / mb / gb / tb / pb* is the data unit that you want to use in conjunction with the capacity and *vdisk_name/vdisk_id* is the name of the VDisk or the ID of the VDisk that you want to expand.

2. On the Windows host, start the Computer Management application and open the Disk Management window under the Storage branch.

You will see the VDisk that you expanded now has some unallocated space at the end of the disk.

You can expand dynamic disks without stopping I/O operations in most cases. However, in some applications the operating system might report I/O errors. When this problem occurs, either of the following entries might 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 in 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.

If the Computer Management application was open before you expanded the VDisk, use the Computer Management application to issue a rescan command.

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.

Shrinking a virtual disk using the CLI

You can shrink a virtual disk (VDisk) using the command-line interface (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 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 Metro Mirror relationships. You should also ensure that the target VDisk is not mapped to any hosts prior to performing this operation.

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>
```

Migrating extents using the CLI

To improve performance, you can migrate extents using the command-line interface (CLI).

The SAN Volume Controller 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.
3. The source MDisk must not currently be the source MDisk for any other migrate extents operation.
4. The destination MDisk must not be the destination MDisk for any other migrate extents operation.

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 VDIs 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
```

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:

```
svcinfolsmdiskextent <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:

```
svcinfolsmdisk <mdiskname/ID>
```

Look for the `mdisk_grp_name` attribute.

- b. List the MDisks in the group by issuing the following command:

```
svcinfolsmdisk -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:

```
svcinfolsfreeextents <mdiskname>
```

You can issue the **`svcinfolsmdiskextent <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 step 1), 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. `<newmdiskname/ID>` is the name or ID of the MDisk to which you want to migrate this set of extents.

5. Repeat steps 2 to 4 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 VDisks between MDisk groups using the CLI

You can migrate virtual disks (VDisks) between managed disk (MDisk) groups using the command-line interface (CLI).

You can determine the usage of particular MDisks by gathering input/output (I/O) statistics about MDisks and VDisks. Once you have gathered this data, you can analyze it to determine which VDisks or MDisks are hot. This procedure then takes you through migrating VDisks 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.

Note: You cannot use the SAN Volume Controller data migration function to move a VDisk between MDisk groups that have different extent sizes. See “Managed disk groups” for more information on extents.

While the migration proceeds, 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 VDisks 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).

Perform the following steps to migrate VDisks between MDisk groups:

1. Isolate any VDisks 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.
- Analyze the dumps to determine which VDIsks 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.
 - Stop the statistics collection again. Issue the following command:

```
svctask stopstats
```

Once you have analyzed the I/O statistics data, you can determine which VDIsks 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 VDIsks in the target MDisk group are less utilized than the source group.

- 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  
<newmdiskgrname/ID> -threads 4
```

- You can check the progress of the migration by issuing the following command:

```
svcinfo lsmigrate
```

Migrating a VDisk between I/O groups using the CLI

Ensure that you are familiar with migrating a virtual disk (VDisk) between I/O groups.

Attention: These migration tasks are disruptive. The cached data that is held within the cluster must first be written to disk before the allocation of the VDisk can be changed.

Modifying the I/O group that services the VDisk cannot be done concurrently with I/O operations. It also requires a rescan at the host level to ensure that the multipathing driver is notified that the allocation of the preferred node has changed and the ports by which the VDisk is accessed has changed. This should only be done in the situation where one pair of nodes has become over utilized.

Perform the following steps to migrate a VDisk between I/O groups:

- Synchronize all file systems that are mounted on the given VDisk.
- Stop all I/O operations to the VDisk.
- Issue the following CLI command to migrate the VDisk into a new I/O group:

```
svctask chvdisk -iogrp new_io_grp_name_or_id  
vdisk
```

Where *new_io_grp_name_or_id* is the name or ID of the I/O group that you want to migrate the VDisk to and *vdisk* is the name of the VDisk that you want to migrate.

- Resynchronize the VDisk to host mapping. See the *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide* or the documentation that is provided with your multipathing driver for more information.
- Restart the I/O operations to the VDisk.

Creating an image mode VDisk using the CLI

You can use the command-line interface (CLI) to import storage that contains existing data and continue to use this storage. You can also make use of the advanced functions, such as Copy Services, data migration, and the cache. These disks are known as image mode virtual disks (VDisks).

Make sure you are aware of the following before you create image mode VDisks:

1. Unmanaged-mode managed disks (MDisks) that contain existing data cannot be differentiated from unmanaged-mode MDisks that are blank. Therefore, it is vital that you control the introduction of these MDisks to the cluster by adding these disks one at a time. For example, map a single LUN from your RAID controller to the cluster and refresh the view of MDisks. The newly detected MDisk is displayed.
2. *Do not* manually add an unmanaged-mode MDisk that contains existing data to an MDisk group. If you do, the data is lost. When you use the command to convert an image mode VDisk from an unmanaged-mode disk, you will select the MDisk group where it should be added.

See the following Web site for more information:

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

For complete instructions on the CLI commands, refer to the IBM TotalStorage SAN Volume Controller: Command-Line Interface User's Guide

Perform the following steps to create an image mode VDisk:

1. Stop all I/O operations from the hosts. Unmap the logical disks that contain the data from the hosts.
2. Create one or more MDisk groups.
3. Ensure that the MDisk groups have enough free capacity to contain all of the migrating data.
4. Map a single RAID array or logical unit from your RAID controller to the cluster. You can do this through a switch zoning or a RAID controller based on your host mappings. The array or logical unit appears as an unmanaged-mode MDisk to the SAN Volume Controller.
5. Issue the **svcinfo lsmdisk** command to list the unmanaged-mode MDisks. If the new unmanaged-mode MDisk is not listed, you can perform a fabric-level discovery. Issue the **svctask detectmdisk** command to scan the fibre-channel network for the unmanaged-mode MDisks.

Note: The **svctask detectmdisk** command also rebalances MDisk access across the available controller device ports.

6. Convert the unmanaged-mode MDisk to an image mode virtual disk. Issue the **svctask mkvdisk** command to create an image mode virtual disk object.
7. Map the new VDisk to the hosts that were previously using the data that the MDisk now contains. You can use the **svctask mkvdiskhostmap** command to create a new mapping between a VDisk and a host. This makes the image mode VDisk accessible for I/O operations to the host.

After the VDisk is mapped to a host object, the VDisk is detected as a disk drive with which the host can perform I/O operations.

If you want to virtualize the storage on an image mode VDisk, you can transform it into a striped VDisk. Migrate the data on the image mode VDisk to managed-mode disks in another MDisk group. Issue the **svctask migratevdisk** command to migrate an entire image mode VDisk from one MDisk group to another MDisk group.

Migrating to an image mode virtual disk using the CLI

You can use the command-line interface (CLI) to migrate data to an image mode virtual disk (VDisk).

The **svctask migratetoimage** CLI command allows you to migrate the data from an existing VDisk onto a different managed disk (MDisk).

When the **svctask migratetoimage** CLI command is issued, it migrates the data of the user specified source VDisk onto the specified target MDisk. When the command completes, the VDisk is classified as an image mode VDisk.

The MDisk specified as the target must be in an unmanaged state at the time the command is run. Issuing this command results in the inclusion of the MDisk into the user specified MDisk group.

Issue the following CLI command to migrate data to an image mode VDisk:

```
svctask migratetoimage -vdisk vdiskname/ID
  -mdisk newmdiskname/ID -mdiskgrp newmdiskgrpname/ID
  -threads 4
```

Where *vdiskname/ID* is the name or ID of the VDisk, *newmdiskname/ID* is the name or ID of the new MDisk, and *newmdiskgrpname/ID* is the name or ID of the new MDisk group.

Deleting a node from a cluster using the CLI

You can use the command-line interface (CLI) to delete a node from a cluster.

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 goes into write-through mode and that you are exposed to a single point of failure if the partner node fails.
- When you delete a node, you remove all redundancy from the I/O group. As a result, new or existing failures can cause I/O errors on the hosts. The following failures can occur:
 - Host configuration errors
 - Zoning errors
 - Multipathing software configuration errors
- If you are deleting the last node in an I/O group and there are virtual disks (VDisks) assigned to the I/O group, you cannot delete the node from the cluster if the node is online. If the node is offline, you can delete the node.
- If you are deleting the last node in an I/O group and there are no VDisks assigned to the I/O group, the cluster is destroyed. You must back up or migrate all data that you want to save before you delete the node.

Perform the following steps to delete a node:

1. Perform the following steps to determine the VDIsks that are still assigned to this I/O group:
 - a. Issue the following CLI command to request a filtered view of the VDIsks:


```
svcinfolsvdisk -filtervalue IO_group_name=name
```

Where *name* is the name of the I/O group for which you want to view the VDIsks.
 - b. Issue the following CLI command to list the hosts that this VDisk is mapped to:


```
svcinfolsvdiskhostmap vdiskname/id
```

Where *vdiskname/id* is the name or ID of the VDisk.

 - If there are VDIsks assigned to this I/O group that contain data that you want to continue to access, migrate the VDIsks to a new I/O group.
2. Power off the node that you want to remove, unless this is the last node in the cluster. This ensures that the subsystem device driver (SDD) does not rediscover the paths that are manually removed before you issue the delete node request.

Attention:

- Deleting or shutting down the configuration node might cause the Secure Shell (SSH) command to hang. If this occurs, wait for the SSH command to end or stop the command and issue the ping command for the cluster IP address. When the failover command completes, you can start to issue commands.
 - If you power on the node that has been removed and it is still connected to the same fabric or zone, it attempts to rejoin the cluster. At this point the cluster tells the node to remove itself from the cluster and the node becomes a candidate for addition to this cluster or another cluster.
 - If you are adding this node into the cluster, ensure that you add it to the same I/O group that it was previously a member of. Failure to do so can result in data corruption.
3. Before deleting the node, it is essential that for each vpath that is presented by the VDIsks you intend to remove, the SDD configuration is updated to remove these vpaths. Failure to do this can result in data corruption. See the *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide* for details about how to dynamically reconfigure SDD for the given host operating system.
 4. Issue the following CLI command to delete a node from the cluster:


```
svctask rmnode node_name_or_id
```

Where *node_name_or_id* is the name or ID of the node.

Performing the cluster maintenance procedure using the CLI

You can use the command-line interface (CLI) to perform the cluster maintenance procedure.

Perform the following steps for cluster maintenance:

1. 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.
2. Issue the `svctask dumperrlog` command to dump the contents of the error log to a text file.
3. Locate and fix the error.

4. Issue the `svctask clearerrlog` command to clear all entries from the error log, including status events and any unfixed errors. Only issue 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.

Note: Clearing the error log does not fix the errors.

5. Issue the `svctask cherrstate` command to toggle the state of an error between unfixed and fixed.

Changing the cluster IP address using the CLI

You can use the command-line interface (CLI) to change the IP address for a cluster.

Attention: When you specify a new IP address for a cluster, the existing communication with the cluster is broken. You must reconnect to the cluster with the new IP address.

Perform the following steps to change the cluster IP address:

1. Issue the **svcinfo lscluster** command to list the current IP address of the cluster.
2. Record the current IP address for future reference.
3. Issue the following command to change the cluster IP address:

```
svctask chcluster -clusterip cluster_ip_address
```

Where *cluster_ip_address* is the new IP address for the cluster.

Changing the cluster gateway address using the CLI

You can use the command-line interface (CLI) to change the gateway address for a cluster.

Perform the following steps to change the cluster gateway address:

1. Issue the **svcinfo lscluster** command to list the current gateway address of the cluster.
2. Record the current gateway address for future reference.
3. Issue the following command to change the cluster gateway address:

```
svctask chcluster -gw cluster_gateway_address
```

Where *cluster_gateway_address* is the new gateway address for the cluster.

Changing the cluster subnet mask using the CLI

You can use the command-line interface (CLI) to change the subnet mask for a cluster.

Perform the following steps to change the cluster subnet mask:

1. Issue the **svcinfo lscluster** command to list the current subnet mask of the cluster.
2. Record the current subnet mask for future reference.
3. Issue the following command to change the cluster subnet mask:

```
svctask chcluster -mask cluster_subnet_mask
```

Where *cluster_subnet_mask* is the new subnet mask for the cluster.

Maintaining SSH keys using the CLI

You can use the command-line interface (CLI) to maintain SSH keys.

Attention: After you add a cluster, close the Maintaining SSH Keys panel.

Perform the following steps to maintain SSH keys:

1. Issue the **svcinfo lsshkeys** CLI command to list the SSH keys that are available on the cluster.
2. Issue the **svctask addsshkey** CLI command to install a new SSH key on the cluster. The key file must first be copied onto the cluster. Each key is associated with an ID string that you define that can consist of up to 30 characters. Up to 100 keys can be stored on a cluster. You can add keys to provide either administrator access or service access. For example, issue the following:

```
svctask addsshkey -user service -file /tmp/id_rsa.pub  
-label testkey
```

Where */tmp/id_rsa.pub* is the name of the file that the SSH key will be saved in and *testkey* is the label to associate with this key.

3. You can issue the **svctask rmsshkey** CLI command to remove an SSH key from the cluster.
4. You can issue the **svctask rmallsshkeys** CLI command to remove all of the SSH keys from the cluster.

Setting up error notifications using the CLI

You can set up error notifications using the command-line interface (CLI).

The error notification settings apply to the entire cluster. You can specify the types of errors that cause the cluster to send a notification. The cluster sends a Simple Network Management Protocol (SNMP) notification. The SNMP setting represents the kind of error.

The following table describes the three types of SNMP notification:

Notification type	Description
All	Report all errors at or above the threshold limit, including information events.
Hardware only	Report all errors at or above the threshold limit, excluding information events.
None	Do not report any errors or information events. This option disables error notification.

If you specify *All* or *Hardware Only*, errors are reported to the SNMP destinations of your choice. To specify an SNMP destination, you *must* provide a valid IP address and SNMP community string.

Note: A valid community string can contain up to 60 letters or digits, without any spaces. A maximum of six SNMP destinations can be specified. When you create the cluster or enable error notification for the first time, you are asked to specify only one SNMP destination. You can add five additional destinations by using the Error Notification options.

The SAN Volume Controller uses the error notifications settings to call Home if errors occur. You must specify *All* or *Hardware Only* and send the trap to the master console if you want the SAN Volume Controller to call Home when errors occur.

Perform the following step to configure the error notification settings:

Issue the **svctask setevent** CLI command to specify the action that you want to take when an error or event is logged to the error log. You can select if 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. For example, you can issue the one of the following CLI commands to set up error notification:

```
svctask setevent -snmptrap all or hardware_only  
-snmpip 9.11.255.634,9.11.265.635 -community mysancommunity,myothersancommunity
```

Where *all or hardware_only* is the type of SNMP notification that you want to set, *9.11.255.634,9.11.265.635* are the IP addresses of the host systems that are running the SNMP manager software, and *mysancommunity,myothersancommunity* are the SNMP community strings that you want to use.

```
svctask setevent -snmptrap none
```

Where *none* indicates that you do not want to report any errors or information events.

Changing cluster passwords using the CLI

You can use the command-line interface (CLI) to change the administrator and service passwords.

Passwords only affect the SAN Volume Controller Console that accesses the cluster. To restrict access to the CLI, you must control the list of SSH client keys that are installed on the cluster.

Perform the following steps to change the administrator and service passwords:

1. Issue the following command to change the administrator users password:

```
svtask chcluster -admpwd admin_password
```

Where *admin_password* is the new administrator password that you want to use.

2. Issue the following command to change the service users password:

```
svtask chcluster -servicepwd service_password
```

Where *service_password* is the new service password that you want to use.

Changing the language setting using the CLI

You can use the command-line interface (CLI) to change the language settings.

Perform the following steps to change the language settings:

Issue the **svcservicetask setlocale** CLI command to change the locale setting for the cluster. This CLI command changes all interfaces output to the chosen language.

You can choose one of the following language settings:

- **0** US English (default)

- 1 Chinese (simplified)
- 2 Chinese (traditional)
- 3 Japanese
- 4 Korean
- 5 French
- 6 German
- 7 Italian
- 8 Spanish
- 9 Portuguese (Brazilian)

Note: This command does not change the front panel display settings.

The following is an example of the CLI command you can issue to change the default language of English to Japanese:

```
svcservicetask setlocale -locale 3
```

Where 3 is the language setting for Japanese.

Viewing the feature log using the CLI

You can use the command-line interface (CLI) to view the feature log.

Perform the following steps to view the feature log:

1. Issue the **svcinfolfeaturedumps** 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.
2. Issue the **svcservicemodeinfo lfeaturedumps** 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

You can use the command-line interface (CLI) to analyze the error log.

Perform the following steps to analyze the error log:

Issue any of the following CLI commands to list error log files:

- **svcinfolerrlogbydisk**
- **svcinfolerrlogbydiskgroup**
- **svcinfolerrlogbyvdisk**
- **svcinfolerrlogbyhost**
- **svcinfolerrlogbynode**
- **svcinfolerrlogbyiogrp**
- **svcinfolerrlogbyfcconsistgrp**
- **svcinfolerrlogbyfcmap**
- **svcinfolerrlogbyrcconsistgrp**
- **svcinfolerrlogbyrcrelationship**

These CLI commands list the error log by type and return a list of dumps in the appropriate directory. For example, the **svcinfolerrlogbymdisk** CLI command displays the error log by managed disks (MDisks).

You can display the whole log or filter the log so that only errors, events, or unfixed errors are displayed. You can also request that the output is sorted either by error priority or by time. For error priority, the most serious errors are the lowest-numbered errors. Therefore, the most serious errors are 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 using the CLI

You can use the command-line interface (CLI) to shut down a cluster.

If you want to remove all input power to a cluster (for example, the machine room power must 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 uninterruptible power supply (UPS), the SAN Volume Controller nodes detect the loss of power and continue to run on battery power until all data that is held in memory is saved to the internal disk drive. This increases the time that is required to make the cluster operational when input power is restored and severely increases the time that is required to recover from an unexpected loss of power that might occur before the UPS batteries have fully recharged.

When input power is restored to the UPSs, they start to recharge. However, the SAN Volume Controller nodes do not permit any I/O activity to be performed to the virtual disks (VDisks) until the UPS is charged enough to enable all the data on the SAN Volume Controller nodes to be saved in the event of an unexpected power loss. This might take as long as three hours. Shutting down the cluster prior to removing input power to the UPS units prevents the battery power from being drained and makes it possible for I/O activity to resume as soon as input power is restored.

Before shutting down a cluster, quiesce all I/O operations that are destined for this cluster. Failure to do so can result in failed I/O operations being reported to your host operating systems.

Attention:

- If you are shutting down the entire cluster, you lose access to all VDisks that are provided by this cluster. Shutting down the cluster also shuts down all SAN Volume Controller nodes. This shutdown causes the hardened data to be dumped to the internal hard drive.
- Ensure that you have stopped all FlashCopy, Metro Mirror and data migration operations before you attempt a cluster shutdown. Also ensure that all asynchronous deletion operations have completed prior to a shutdown operation.

Begin the following process of quiescing all I/O to the cluster by stopping the applications on your hosts that are using the VDisks that are provided by the cluster.

1. Determine which hosts are using the VDisks that are provided by the cluster.
2. Repeat the previous step for all VDisks.

You can shut down a cluster by stopping I/O activity and either pressing the power buttons on the front of each SAN Volume Controller node or by issuing a shutdown command to the cluster.

Attention: You must press and hold the power button for one second to shutdown the SAN Volume Controller node.

When input power is restored, you must press the power button on the UPS units before you press the power buttons on the SAN Volume Controller nodes.

Perform the following steps to shut down a cluster:

1. Issue the following command to shut down a cluster:

```
svctask stopcluster
```

The following output is displayed:

```
Are you sure that you want to continue with the shut down?
```

2. Type `y` to shut down the entire cluster.

Chapter 5. Backing up and restoring the cluster configuration

You can back up and restore the cluster configuration data after preliminary tasks are completed.

Cluster configuration data provides information about your cluster and the objects that are defined in it. The backup and restore functions of the **svconfig** command can only back up and restore your cluster configuration data. You must regularly back up your application data using the appropriate backup methods.

Information about the following objects is included in the cluster configuration data:

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

You can maintain your cluster configuration data by performing the following tasks:

- Backing up the cluster configuration data
- Restoring the cluster configuration data
- Deleting unwanted backup configuration data files

Backing up the cluster configuration

You can backup your cluster configuration data from the Backing up a Cluster Configuration panel.

Before you backup your cluster configuration data, the following prerequisites must be met:

- No independent operations that change the cluster configuration can be running while the backup command is running.
- No object name can begin with an underscore.
- All objects should have non-default names, that is, names that are not assigned by the SAN Volume Controller.

Note:

- The default object names for controllers, I/O groups and managed disks (MDisks) do not restore correctly if the ID of the object is different than what is recorded in the current cluster configuration data file.

- All other objects with default names are renamed during the restore process. The new names appear in the format *name_r*. Where *name* is the name of the object in your cluster.

This task assumes that you have already launched the SAN Volume Controller application.

The backup function is designed to back up information about your cluster configuration, such as virtual disks (VDisks), local Metro Mirror information, managed disk (MDisk) groups, and nodes. All other data that you have written to the VDisks is *not* backed up. Any application that uses the VDisks on the cluster as storage, must back up its application data using appropriate backup methods.

You must regularly back up your cluster configuration data and your application data to avoid data loss. If a cluster is lost after a severe failure occurs, both cluster configuration and application data is lost. You must reinstate the cluster to the exact state it was in prior to the failure and then recover the application data.

The backup function creates three files that provide information about the backup process and cluster configuration. When you use the SAN Volume Controller Console to perform the backup, these files are created in the `\console\backup\cluster` directory of the master console. Where *console* is the directory where the SAN Volume Controller Console is installed and *cluster* is the name of the cluster for which you want to back up the cluster configuration data.

The following table describes the three files that are created by the backup process:

File name	Description
svc.config.backup.xml	This file contains your cluster configuration data.
svc.config.backup.sh	This file contains the names of the commands that were issued to create the backup of the cluster.
svc.config.backup.log	This file contains details about the backup, including any error information that might have been reported.

If the `svc.config.backup.xml` file already exists in the directory, it is renamed to `svc.config.backup.bak`. After the file is renamed the new `svc.config.backup.xml` is written.

Perform the following steps to backup your cluster configuration data:

1. Click **Service and Maintenance** → **Backup Configuration** in the portfolio. The Backing up a Cluster Configuration panel is displayed.
2. Click **Backup**.

Backing up the cluster configuration using the CLI

You can backup your cluster configuration data using the command-line interface (CLI).

Before you backup your cluster configuration data, the following prerequisites must be met:

- No independent operations that change the cluster configuration can be running while the backup command is running.
- No object name can begin with an underscore.
- All objects should have non-default names, that is, names that are not assigned by the SAN Volume Controller.

Note:

- The default object names for controllers, I/O groups and managed disks (MDisks) do not restore correctly if the ID of the object is different than what is recorded in the current cluster configuration data file.
- All other objects with default names are renamed during the restore process. The new names appear in the format *name_r*. Where *name* is the name of the object in your cluster.

The backup feature of the **svcconfig** CLI command is designed to back up information about your cluster configuration, such as virtual disks (VDisks), local Metro Mirror information, managed disk (MDisk) groups, and nodes. All other data that you have written to the VDIs is *not* backed up. Any application that uses the VDIs on the cluster as storage, must back up its application data using the appropriate backup methods.

You must regularly back up your cluster configuration data and your application data to avoid data loss. If a cluster is lost after a severe failure occurs, both cluster configuration and application data is lost. You must reinstate the cluster to the exact state it was in prior to the failure and then recover the application data.

Perform the following steps to backup your cluster configuration data:

1. Back up all of the application data that you have stored on your VDIs using your preferred backup method.
2. Open a command prompt.
3. Issue the following command to log onto the cluster:

```
ssh -l admin your_cluster_name -p 22
```

Where *your_cluster_name* is the name of the cluster for which you want to backup cluster configuration data.

4. Issue the following CLI command to remove all of the existing cluster configuration backup and restore files that are on your cluster:

```
svcconfig clear -all
```

5. Issue the following CLI command to backup your cluster configuration:

```
svcconfig backup
```

The following output is an example of the messages that are displayed during the backup process:

```
CMMVC6112W io_grp io_grp1 has a default name
CMMVC6112W io_grp io_grp2 has a default name
CMMVC6112W mdisk mdisk14 ...
CMMVC6112W node node1 ...
CMMVC6112W node node2 ...
.....
CMMVC6136W No SSH key file svc.config.renee.admin.key
CMMVC6136W No SSH key file svc.config.service.service.key
```

The **svcconfig backup** CLI command creates three files that provide information about the backup process and cluster configuration. These files are created in the /tmp directory of the configuration node.

The following table describes the three files that are created by the backup process:

File name	Description
svc.config.backup.xml	This file contains your cluster configuration data.
svc.config.backup.sh	This file contains the names of the commands that were issued to create the backup of the cluster.
svc.config.backup.log	This file contains details about the backup, including any error information that might have been reported.

6. Issue the following command to exit the cluster:

```
exit
```

7. Issue the following command to copy the backup files to a location that is not in your cluster:

```
scp -P 22 admin@your_cluster:/tmp/svc.config.backup.*  
/offclusterstorage/
```

Where *your_cluster* is the name of your cluster and *offclusterstorage* is the location where you want to store the backup files.

You must copy these files to a location outside of your cluster because the /tmp directory on this node becomes inaccessible if the configuration node changes. The configuration node might change in response to an error recovery action or to a user maintenance activity.

Tip: To maintain controlled access to your cluster configuration data, copy the backup files to a location that is password protected.

8. Ensure that the copies of the backup files are stored in the location that you specified in step 7.

You can rename the backup files to include the configuration node name either at the start or end of the file names so you can easily identify these files when you are ready to restore your configuration.

Issue the following command to rename the backup files that are stored on a Linux or AIX host:

```
mv /offclusterstorage/svc.config.backup.xml  
/offclusterstorage/svc.config.backup.xml_myconfignode
```

Where *offclusterstorage* is the name of the directory where the backup files are stored and *myconfignode* is the name of your configuration node.

To rename the backup files that are stored on a Windows host, right-click on the name of the file and select **Rename**.

Downloading backup configuration data files

You can use the SAN Volume Controller Console to download backup configuration data files to your master console.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to download the backup configuration data files to your master console:

1. Click **Service and Maintenance** → **List Dumps** in the portfolio. The List Dumps panel is displayed.
2. Click **Software Dumps**. The Software Dumps panel is displayed.
3. Find the name of your backup configuration data file.
4. Right-click on your backup configuration data file and click **Save Target As**.
5. Select the location where you want to save the file and click **Save**.

Restoring the cluster configuration using the CLI

You can restore your cluster configuration data using the command-line interface (CLI).

Before you restore your cluster configuration data, the following prerequisites must be met:

- You have superuser administrator authority.
- You have a copy of your backup cluster configuration files on a server that is accessible to the cluster.
- You have a backup copy of your application data.
- You know the current feature settings for your cluster.
- You have not removed any hardware since the last backup of your cluster configuration. If you had to replace a faulty node, the new node must use the same worldwide node name (WWNN) as the faulty node that it replaced.

Note: You can add new hardware, but you must not remove any hardware because the removal can cause the restore process to fail.

- No changes have been made to the fabric of the cluster since the last backup of your cluster configuration. If changes are made, you must back up your cluster configuration again.

The restore must be performed to a single node cluster. You can restore the configuration using any node as the configuration node. However, if you do not use the node that was the configuration node when the cluster was first created, the SCSI inquiry identifiers of the I/O groups can change. This can affect IBM TotalStorage Productivity Center for Fabric (TPC for Fabric), VERITAS Volume Manager and any other programs that record this information.

The SAN Volume Controller analyzes the backup configuration data file and the cluster to verify that the required disk controller system nodes are available.

Before you begin, hardware recovery must be complete. The following hardware must be operational: hosts, SAN Volume Controller, disk controller systems, disks, and the SAN fabric.

Important: There are two phases during the restore process: prepare and execute. You must not make any changes to the fabric or cluster between these two phases.

Perform the following steps to restore your cluster configuration data:

1. Select delete cluster from the front panel on each node in the cluster that does *not* display `Cluster :` on the front panel. If the front panel of the node displays `Cluster :`, the node is already a candidate node.

2. Create a new cluster from the front panel of any node in the cluster. If possible, use the node that was originally the configuration node for the cluster.
3. Generate a Secure Shell (SSH) key pair for the SAN Volume Controller Console.
4. Generate an SSH key pair for all of the hosts that you want to use to access the CLI.
5. Log on to the SAN Volume Controller Console.
6. Finish creating the cluster by using the SAN Volume Controller Console.
After the cluster is created and configured, you should be able to connect to the cluster using the master console or the CLI.
7. Issue the following command to log onto the cluster:

```
ssh -l admin your_cluster_name -p 22
```


Where *your_cluster_name* is the name of the cluster for which you want to restore the cluster configuration.
8. Issue the following CLI command to ensure that only the configuration node is online.

```
svcinfolnode
```


The following is an example of the output that is displayed:

```
id name status IO_group_id IO_group_name config_node
1 node1 online 0 io_grp0 yes
```
9. Issue the following CLI command to remove all of the existing backup and restore cluster configuration files that are on your cluster:

```
svcconfig clear -all
```
10. Issue the following command to exit the cluster:

```
exit
```
11. Copy the `svc.config.backup.xml` file from the master console to the `/tmp` directory of the cluster using the PuTTY `pscp` program. Perform the following steps to use the PuTTY `pscp` program to copy the file:
 - a. Open a command prompt from the master console.
 - b. Set the path in the command line to use `pscp` with the following format:

```
set PATH=C:\path\to\putty\directory;%PATH%
```
 - c. Issue the following command to specify the location of your private SSH key for authentication:

```
pscp <private key location> source [source...] [user@]host:target
```
12. Issuing the following CLI command to compare the current cluster configuration with the backup configuration data file:

```
svcconfig restore -prepare
```


This CLI command creates a log file in the `/tmp` directory of the configuration node. The name of the log file is `svc.config.restore.prepare.log`.
13. Issue the following command to copy the log file to another server that is accessible to the cluster:

```
pscp -i <private key location> [user@]host:source target
```
14. Open the log file from the server where the copy is now stored.
15. Check the log file for errors.
 - If there are errors, correct the condition which caused the errors and reissue the command. You must correct all errors before you can proceed to step 16 on page 201.
 - If you need assistance, contact IBM Customer Support.

16. Issue the following CLI command to restore the cluster configuration:

```
svconfig restore -execute
```

Note: Issuing this CLI command on a single node cluster adds the other nodes and hosts to the cluster.

This CLI command creates a log file in the /tmp directory of the configuration node. The name of the log file is `svc.config.restore.execute.log`.

17. Issue the following command to copy the log file to another server that is accessible to the cluster:

```
scp -P 22 admin@your_cluster:/tmp/svc.config.restore.execute.log  
/offclusterstorage/
```

Where *your_cluster* is the name of your cluster and *offclusterstorage* is the location where you want to store the log file.

18. Open the log file from the server where the copy is now stored.
19. Check the log file to ensure that no errors or warnings have occurred.

Note: You might receive a warning that states a featurization is not enabled. This means that after the recovery process, the current feature settings do not match the previous feature settings. The recovery process continues normally and you can enter the correct feature settings in the SAN Volume Controller Console at a later time.

The following output is displayed after a successful cluster configuration restore:

```
IBM_2145:your_cluster_name:admin>
```

You can remove any unwanted configuration backup and restore files from the cluster by issuing the `svconfig clear -all` CLI command.

Related tasks

“Backing up the cluster configuration using the CLI” on page 196

You can backup your cluster configuration data using the command-line interface (CLI).

“Creating a cluster from the front panel” on page 79

After you have created a pair of nodes, you can use the front panel of the SAN Volume Controller to create a cluster.

“Creating a cluster using the SAN Volume Controller Console” on page 86

After you have created a pair of nodes, you can create and configure a cluster.

Related information

Chapter 2, “Creating a SAN Volume Controller cluster,” on page 79

You must complete the two phases that are required to create a cluster before you can configure the SAN Volume Controller.

Deleting backup configuration files

You can delete a backup cluster configuration from the Deleting a Cluster Configuration panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to delete backup configuration files:

1. Click **Service and Maintenance** → **Delete Backup** in the portfolio. The Deleting a Cluster Configuration panel is displayed.
2. Click **OK**.

Deleting backup configuration files using the CLI

You can use the command-line interface (CLI) to delete backup configuration files.

Perform the following steps to delete backup configuration files:

1. Issue the following command to log onto the cluster:

```
ssh -l admin your_cluster_name -p 22
```

Where *your_cluster_name* is the name of your cluster.

2. Issue the following CLI command to erase all of the files that are stored in the `/tmp` directory:

```
svconfig clear -all
```

Chapter 6. Upgrading the SAN Volume Controller software

The SAN Volume Controller software can be upgraded while you run day-to-day operations.

However, performance is degraded during the software upgrade process.

Note: Applying a software upgrade takes approximately one hour because there is a 30 minute delay to allow the multipathing software to recover.

Software and microcode for the SAN Volume Controller and its attached adapters is tested and released as a single package. The package number increases each time a new release is made. The package includes Linux, Apache and the SAN Volume Controller software.

If you upgrade to more than one level above your current level, you might be required to install the intermediate level. For example, if you are upgrading from level 1 to level 3, you might have to install level 2 before you can install level 3. Details for any prerequisite levels are provided with the source files.

Attention:

- If you apply the software upgrade while a node is in service mode, the node is deleted from the cluster. All status information that is stored on the node is deleted and data loss can occur if the cluster depends solely on this node.
- Ensure that you have no unfixed errors in the log and that the cluster date and time are correctly set. Start the Directed Maintenance Procedures (DMPs) and ensure that you fix any outstanding errors before you attempt to concurrently upgrade the software.

Metro mirror

When you upgrade software where the cluster participates in one or more intercluster relationships, update the clusters one at a time. Do not upgrade the clusters concurrently because you can lose synchronization and availability.

You can create new metro mirror partnerships between two clusters with different major software version numbers.

Related tasks

“Copying the SAN Volume Controller software upgrade files using PuTTY scp” on page 205

PuTTY scp (pscp) provides a file transfer application for secure shell (SSH) to copy files either between two directories on the configuration node or between the configuration node and another host.

“Upgrading the SAN Volume Controller software using the SAN Volume Controller Console” on page 206

You can upgrade the cluster software using the SAN Volume Controller Console.

“Upgrading the SAN Volume Controller software using the CLI” on page 209

You can use the command-line interface (CLI) to install software upgrades.

“Performing a disruptive software upgrade using the CLI” on page 211

You can use the command-line interface (CLI) to perform a disruptive software upgrade.

“Performing the node rescue” on page 212

If it is necessary to replace the hard disk drive or if the software on the hard

disk drive is corrupted, you can reinstall the software on the SAN Volume Controller by using the node rescue procedure.

“Recovering from software upgrade problems manually” on page 213
When a new software level is committed, you might not be able to return to a previous software level because some data structures might have been changed such that they cannot be used with the previous software level. Therefore, if you have any problems, you must install the newest level of the software.

Related reference

“Installing or upgrading the SAN Volume Controller software”

The SAN Volume Controller software can be installed or upgraded after you have downloaded the software package from the SAN Volume Controller Web site.

“Upgrading the SAN Volume Controller software automatically” on page 206
When new nodes are added to the cluster, the software upgrade file is automatically downloaded to the new nodes from the SAN Volume Controller cluster.

“Recovering from software upgrade problems automatically” on page 213
The cluster automatically stops the software upgrade process if any of the nodes fail to upgrade to the new software level.

Installing or upgrading the SAN Volume Controller software

The SAN Volume Controller software can be installed or upgraded after you have downloaded the software package from the SAN Volume Controller Web site.

Software package

The software installation or upgrade procedure copies the new software level to the cluster and starts an automatic installation process. During the installation process, each node is restarted. While each node restarts, there might be some degradation in the maximum I/O rate that can be sustained by the cluster. The amount of time that is needed to install or upgrade the software is dependent on the size of the cluster and the size of the software update package. The size of the software update package is determined by the number of components that are being replaced. After all the nodes in the cluster are successfully restarted with the new software level, the new software level is automatically committed.

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, these restrictions are documented on the SAN Volume Controller Web site that you use to download the software packages. During the software upgrade procedure, only the following SAN Volume Controller commands are operational from the time the install process starts to the time that the new software level is committed, or until the process has been backed-out. All other commands fail with a message that indicates a software upgrade is in progress. The object type in the following commands is xxxx:

- `svcinfo lsxxxx`
- `svcinfo lsxxxxcandidate`
- `svcinfo lsxxxxprogress`
- `svcinfo lsxxxxmember`
- `svcinfo lsxxxxextent`

- `svcinfolsxrxdumps`
- `svcinfolcaterrlog`
- `svcinfolerrlogbyxxxx`
- `svcinfolcaterrlogbyseqnum`
- `svctaskrmnode`
- `svcservicetaskrmnode`

Because of the operational limitations that occur during the software upgrade process, the software installation is a user task.

Related reference

Chapter 6, “Upgrading the SAN Volume Controller software,” on page 203
The SAN Volume Controller software can be upgraded while you run day-to-day operations.

Copying the SAN Volume Controller software upgrade files using PuTTY scp

PuTTY scp (pscp) provides a file transfer application for secure shell (SSH) to copy files either between two directories on the configuration node or between the configuration node and another host.

To use pscp, you must have the appropriate permissions on the source and destination directories on your respective hosts.

The pscp application is available when you install an SSH client on your host system. You can access the pscp application through a command prompt.

Perform the following steps to use pscp:

1. Start a PuTTY session.
2. Configure your PuTTY session to access your SAN Volume Controller Console cluster.
3. Save your PuTTY configuration session. For example, you can name your saved session SVCPUTTY.
4. Open a command prompt.
5. Issue the following command to set the path environment variable to include the PuTTY directory:

```
set path=C:\Program Files\putty;%path%
```

Where *Program Files* is the directory where PuTTY is installed.

6. Issue the following command to copy the package onto the node where the CLI runs:

```
directory_software_upgrade_files pscp -load saved_putty_configuration
software_upgrade_file_name admin@cluster_ip_address:home/admin/upgrade
```

Where *directory_software_upgrade_files* is the directory that contains the software upgrade files, *saved_putty_configuration* is the name of the PuTTY configuration session, *software_upgrade_file_name* is the name of the software upgrade file and *cluster_ip_address* is the IP address of your cluster.

If there is insufficient space to store the software upgrade file on the cluster, the copy process fails. Perform one of the following steps:

- Issue the **svctask cleardumps** CLI command to free space on the cluster and repeat step 6.

- Issue the following command from the cluster to transfer the error logs to the master console:

```
pscp -unsafe -load saved_putty_configuration  
admin@cluster_ip_address:/dump/elog/* your_desired_directory
```

Where *saved_putty_configuration* is the name of the PuTTY configuration session, *cluster_ip_address* is the IP address of your cluster and *your_desired_directory* is the directory where you want to transfer to the error logs.

After you have transferred the error logs to the master console, repeat step 6 on page 205.

Related reference

Chapter 6, “Upgrading the SAN Volume Controller software,” on page 203
The SAN Volume Controller software can be upgraded while you run day-to-day operations.

Upgrading the SAN Volume Controller software automatically

When new nodes are added to the cluster, the software upgrade file is automatically downloaded to the new nodes from the SAN Volume Controller cluster.

If you add a new node that has or requires a software level that is higher than the software level available on the cluster, the new node is *not* configured into the cluster. The new node must be downgraded to the software level of the cluster before it can join the cluster. If a node is added to the cluster that does not have software installed or has an old software level that cannot be recognized by the cluster, a node rescue must be performed to force a reinstallation of the software.

Error counts

During the software upgrade, increased I/O error counts are displayed by the **datapath query device** or **datapath query adapter** commands if active I/O operations exist between the hosts and the SANs. See the *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide* for more information about the **datapath query** commands.

During the software upgrade, each node of a working pair is upgraded sequentially. The node that is being upgraded is temporarily unavailable and all I/O operations to that node fail. As a result, the I/O error counts increase and the failed I/O operations are directed to the partner node of the working pair. Applications should not see any I/O failures.

Related reference

Chapter 6, “Upgrading the SAN Volume Controller software,” on page 203
The SAN Volume Controller software can be upgraded while you run day-to-day operations.

Upgrading the SAN Volume Controller software using the SAN Volume Controller Console

You can upgrade the cluster software using the SAN Volume Controller Console.

Perform the following steps if you are using Internet Explorer:

1. Click **Tools** in the menu bar.
2. Select the **Internet Options** → **Connections** tab.

3. Click on **LAN Settings...** and ensure that the box marked **Use a proxy server for your LAN (These settings will not apply to dial-up or VPN connections)** is unchecked.
4. Click **OK** twice to accept the settings.

Perform the following steps if you are using Netscape:

1. Click **Edit** in the menu bar.
2. Click on **Preferences....** Expand the Advanced section and select **Proxies.**
3. Select the **Direct connection to the Internet** button and click **OK** to accept the settings.

Tip: The software upgrade files can be quite large. If you experience problems while uploading the software upgrade files to the cluster, you should disable proxies on the Web browser from where you will upload the file. This shortens the file upload time. If you disable proxies, you might not be able to connect to external Web sites. Therefore, you must make a record of your existing settings before you disable proxies in case you have to restore access to other Web sites.

Perform the following steps to upgrade the software:

1. Download the SAN Volume Controller code from the following IBM Web site:
<http://www.ibm.com/servers/storage/support/virtual/2145.html>
 - If you want to write the SAN Volume Controller code to a CD, you must download the CD image.
 - If you do not want to write the SAN Volume Controller code to a CD, you must download the install image.
2. Start a SAN Volume Controller Console session.
3. Launch the SAN Volume Controller application.
4. Click **Service and Maintenance** in the portfolio.
5. Click **Upgrade Software** to check the installed software level or to install a new level of software on the cluster. The Software Upgrade panel is displayed.

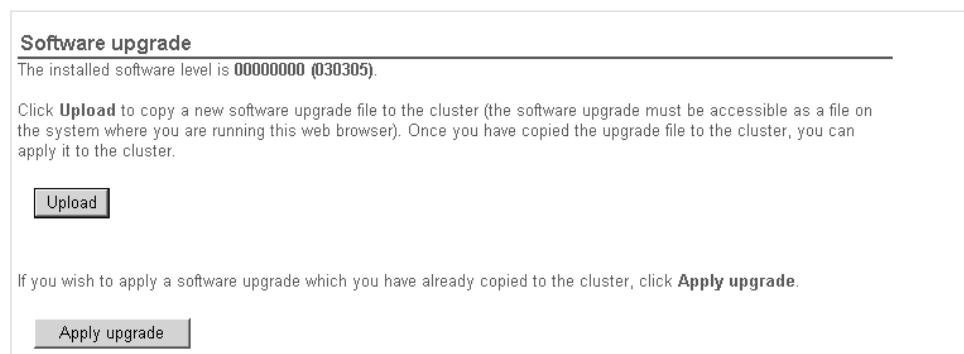


Figure 26. Software upgrade panel

6. Click **Upload**. The Software upgrade - file upload panel is displayed.

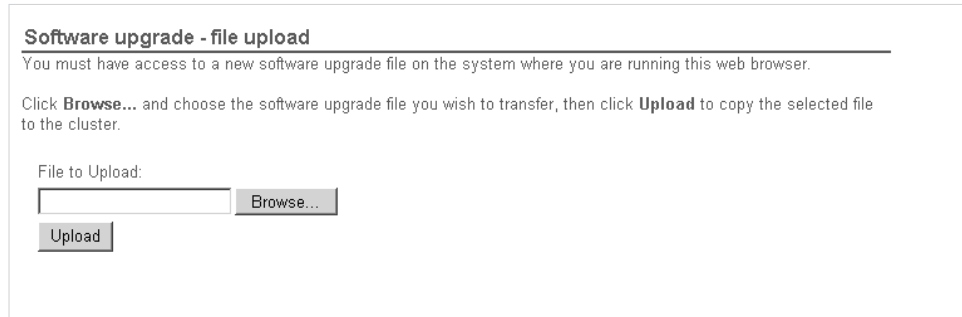


Figure 27. Software upgrade - file upload panel

7. Click **Browse** and select the SAN Volume Controller software file that you downloaded in step 1 on page 207.
8. Click **Upload** to copy the SAN Volume Controller software file to the cluster.

Before you begin the software upgrade, you must be aware of the following:

 - The install process fails if all the nodes that are configured in the cluster are not present. You cannot use the force flag to force the install process. If any node that is configured to be a member of the cluster is not present, the node must either be deleted from the cluster or be brought online before you can upgrade the software. Furthermore, if a node has been deleted from the cluster such that any I/O group has only one member, the software upgrade also fails. This is because the upgrade process causes a loss of access to data. The force flag can be used to override this restriction if you are prepared to loose access to data during the upgrade.
 - The software upgrade is distributed to all the nodes in the cluster using fibre-channel connections between the nodes.
 - Nodes are updated one at a time.
 - Nodes will run the new software, concurrently with normal cluster activity.
 - The procedure to update a single node takes approximately 5 minutes.
 - While the node is updated, it does not participate in I/O activity in the I/O group. As a result, all I/O activity for the VDisks in the I/O group is directed to the other node in the I/O group by the host multipathing software.
 - While the node is updated, the other node in the I/O group notices that its partner node is not participating in the cluster and attempts to flush the writeback cache and set it into write-through mode. This flush is not guaranteed to be successful or to complete and as a result concurrent software update creates a single point of data loss. If the remaining node in an I/O group experiences a failure during a software update of its partner, the only valid copy of dirty data in the writeback cache can be lost.
 - There is a 30 minute delay between node updates. The delay allows time for the host multipathing software to rediscover paths to the nodes which have been upgraded, so that there is no loss of access when another node in the I/O group is updated.
 - The software update is not committed until all nodes in the cluster have been successfully updated to the new software level. If all nodes successfully restart with the new software level, the new level is committed. When the new level is committed, the cluster vital product data (VPD) is updated to reflect the new software level. After the cluster VPD is updated, you can no longer downgrade to a software level with a lower major number.
 - You cannot invoke the new functions of the upgraded software until all member nodes are upgraded and the update has been committed.

- Because the software upgrade process takes some time, the install command completes as soon as the software level is verified by the cluster. To determine when the upgrade has completed, you must either display the software level in the cluster VPD or look for the Software upgrade complete event in the error/event log. If any node fails to restart with the new software level or fails at any other time during the process, the software level is backed-off.
 - During a software upgrade the version number of each node is updated when the software has been installed and the node has been restarted. The cluster software version number is updated when the new software level is committed.
 - When the software upgrade starts an entry is made in the error or event log and another entry is made when the upgrade completes or fails.
9. Click **Apply upgrade**. The Applying Software Upgrade panel is displayed. The Applying Software Upgrade panel enables you to select the upgrade and apply it to the cluster. A list of the software levels that you can apply to the cluster is displayed.

When a new software level is applied, it is automatically installed on all the nodes that are in the cluster.

Note: The software upgrade can take up to 30 minutes per node.

Related reference

Chapter 6, “Upgrading the SAN Volume Controller software,” on page 203
The SAN Volume Controller software can be upgraded while you run day-to-day operations.

Upgrading the SAN Volume Controller software using the CLI

You can use the command-line interface (CLI) to install software upgrades.

Perform the following steps to upgrade the software:

1. Download the SAN Volume Controller code from the following IBM Web site:
<http://www.ibm.com/servers/storage/support/virtual/2145.html>
 - If you want to write the SAN Volume Controller code to a CD, you must download the CD image.
 - If you do not want to write the SAN Volume Controller code to a CD, you must download the install image.
2. Use PuTTY scp (pscp) to copy the software upgrade files to the node.
3. Ensure that the software upgrade file has been successfully copied.

Before you begin the software upgrade, you must be aware of the following:

- The install process fails if all the nodes that are configured in the cluster are not present. You cannot use the force flag to force the install process. If any node that is configured to be a member of the cluster is not present, the node must either be deleted from the cluster or be brought online before you can upgrade the software. Furthermore, if a node has been deleted from the cluster such that any I/O group has only one member, the software upgrade also fails. This is because the upgrade process causes a loss of access to data. The force flag can be used to override this restriction if you are prepared to lose access to data during the upgrade.
- The software upgrade is distributed to all the nodes in the cluster using fibre-channel connections between the nodes.

- Nodes are updated one at a time.
 - Nodes will run the new software, concurrently with normal cluster activity.
 - The procedure to update a single node takes approximately 5 minutes.
 - While the node is updated, it does not participate in I/O activity in the I/O group. As a result, all I/O activity for the VDisks in the I/O group is directed to the other node in the I/O group by the host multipathing software.
 - While the node is updated, the other node in the I/O group notices that its partner node is not participating in the cluster and attempts to flush the writeback cache and set it into write-through mode. This flush is not guaranteed to be successful or to complete and as a result concurrent software update creates a single point of data loss. If the remaining node in an I/O group experiences a failure during a software update of its partner, the only valid copy of dirty data in the writeback cache can be lost.
 - There is a 30 minute delay between node updates. The delay allows time for the host multipathing software to rediscover paths to the nodes which have been upgraded, so that there is no loss of access when another node in the I/O group is updated.
 - The software update is not committed until all nodes in the cluster have been successfully updated to the new software level. If all nodes successfully restart with the new software level, the new level is committed. When the new level is committed, the cluster vital product data (VPD) is updated to reflect the new software level. After the cluster VPD is updated, you can no longer downgrade to a software level with a lower major number.
 - You cannot invoke the new functions of the upgraded software until all member nodes are upgraded and the update has been committed.
 - Because the software upgrade process takes some time, the install command completes as soon as the software level is verified by the cluster. To determine when the upgrade has completed, you must either display the software level in the cluster VPD or look for the Software upgrade complete event in the error/event log. If any node fails to restart with the new software level or fails at any other time during the process, the software level is backed-off.
 - During a software upgrade the version number of each node is updated when the software has been installed and the node has been restarted. The cluster software version number is updated when the new software level is committed.
 - When the software upgrade starts an entry is made in the error or event log and another entry is made when the upgrade completes or fails.
4. Issue the following CLI command to start the software upgrade process:
`svcservicetask applysoftware -file software_upgrade_file`
 Where *software_upgrade_file* is the name of the software upgrade file.
 5. Perform the following steps to verify that the software upgrade successfully completed:
 - a. Issue the **svctask dumperrlog** CLI command to send the contents of the error log to a text file.
 The following output is displayed in the text file if the software is successfully upgraded:

Upgrade completed successfully
 - b. Issue the **svcinfolnodevpd** CLI command for each node that is in the cluster. The software version field displays the new software level.

When a new software level is applied, it is automatically installed on all the nodes that are in the cluster.

Note: The software upgrade can take up to 30 minutes per node.

Related reference

Chapter 6, “Upgrading the SAN Volume Controller software,” on page 203
The SAN Volume Controller software can be upgraded while you run day-to-day operations.

“SAN Volume Controller library and related publications” on page xx

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

Performing a disruptive software upgrade using the CLI

You can use the command-line interface (CLI) to perform a disruptive software upgrade.

The SAN Volume Controller only supports concurrent software upgrades. To ensure that a software upgrade is coordinated across all nodes in the cluster, the nodes must be able to communicate with each other across the fibre-channel SAN. However, if this is not possible, you can perform a disruptive software upgrade.

Perform the following steps to complete the disruptive software upgrade process:

1. Stop any host applications and unmount the file systems that use storage that is managed by the SAN Volume Controller. If you are shutting down your hosts, this occurs as the host is shutdown. If you are not shutting down your hosts, you must manually stop host applications and unmount the file systems on each host. This step ensures that the hosts stop issuing I/O operations and that any data in the file system caches is flushed.
2. Issue the **svctask stopcluster** CLI command to shutdown the cluster. This CLI command stops the SAN Volume Controllers from issuing I/O to backend controllers and flushes data from the SAN Volume Controller cache.
3. Rezone the switch so that the SAN Volume Controller nodes are in one zone. Ensure that this zone does not include a host HBA or a backend controller (keep the old switch configuration so it can be restored during step 6). This step isolates the SAN Volume Controller from the rest of the SAN.
4. Power on all the SAN Volume Controller nodes and wait for them to reform a cluster.

Note: Because the SAN Volume Controller has been isolated from the backend storage, errors that indicate the backend storage is unavailable are logged.

5. Perform the software upgrade in the same manner as for a concurrent software upgrade.
6. Restore the original switch configuration.
7. Clear any error logs that were produced in step 4 indicating that backend storage is unavailable. Check that all backend storage is now online and accessible to the SAN Volume Controllers.
8. Remount file systems and start host applications.

Related tasks

“Shutting down a cluster using the CLI” on page 193

You can use the command-line interface (CLI) to shut down a cluster.

Related reference

Chapter 6, “Upgrading the SAN Volume Controller software,” on page 203
The SAN Volume Controller software can be upgraded while you run day-to-day operations.

Performing the node rescue

If it is necessary to replace the hard disk drive or if the software on the hard disk drive is corrupted, you can reinstall the software on the SAN Volume Controller by using the node rescue procedure.

Attention: If you recently replaced the service controller and the disk drive as part of the same repair operation, node rescue fails. See the related information about replacing a disk drive and a service controller to resolve this issue.

To provide an alternate boot device, a minimal operating system is also available in nonvolatile memory on the service controller. If it is necessary to replace the hard disk drive or the software on the hard disk drive has become corrupted, the SAN Volume Controller cannot boot and the hardware boot indicator remains on the front panel display or the boot operation does not progress.

If this occurs, you can reinstall the software on the SAN Volume Controller by using the node rescue procedure. Node rescue works by booting the operating system from the service controller and running a program that copies all the node software from any other SAN Volume Controller that can be found on the fibre-channel fabric.

Perform the following steps to complete the node rescue:

1. Ensure that the fibre-channel cables are connected.
2. Ensure that at least one other SAN Volume Controller node is connected to the fibre-channel fabric.
3. Turn off the SAN Volume Controller.
4. Press and hold the Left and Right buttons on the front panel.
5. Press the Power button.
6. Continue to hold the Left and Right buttons until the node-rescue-request symbol is displayed on the front panel (Figure 28).



Figure 28. Node-rescue-request display

The node rescue request symbol displays on the front panel display until the SAN Volume Controller starts to boot from the service controller. If the node rescue request symbol displays for more than two minutes, go to the hardware boot MAP to resolve the problem. When the node rescue starts, the service display shows the progress or failure of the node rescue operation.

Note: If the recovered node was part of a cluster, the node is now offline. Delete the offline node from the cluster and then add the node back into the cluster. If node recovery was used to recover a node that failed during a software

upgrade process, the automatic software downgrade process starts but might not continue until the failed node is deleted from the cluster. After the failed node is deleted, it is not possible to add the node back into the cluster until the downgrade process has completed. This may take up to four hours for an eight-node cluster.

If the cables are correctly located and the node rescue request symbol still displays, replace the field replaceable units (FRUs) in the following sequence:

SAN Volume Controller 2145-8F2	SAN Volume Controller 2145-4F2
1. Service controller 2. Frame assembly	1. Service controller 2. System board assembly

Related reference

Chapter 6, “Upgrading the SAN Volume Controller software,” on page 203
The SAN Volume Controller software can be upgraded while you run day-to-day operations.

Recovering from software upgrade problems automatically

The cluster automatically stops the software 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 are downgraded to the original software level. You must check the error log to determine the reason for the failure before you attempt to upgrade the cluster again.

Related reference

Chapter 6, “Upgrading the SAN Volume Controller software,” on page 203
The SAN Volume Controller software can be upgraded while you run day-to-day operations.

Recovering from software upgrade problems manually

When a new software level is committed, you might not be able to return to a previous software level because some data structures might have been changed such that they cannot be used with the previous software level. Therefore, if you have any problems, you must install the newest level of the software.

Attention: This procedure causes a loss of *all* data that is currently configured in the cluster. This procedure must only be used as a last resort and should only be done if you have recently backed-up your data.

In extreme conditions where you cannot wait for a software update and you need to return to the previous software level, you can use the following procedure.

Attention: This procedure causes the total loss of the SAN Volume Controller cluster. This procedure must only be used as a last resort.

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 service access mode.
3. Use the service access mode functions to force the download of the older software level.

4. Repeat the action for each of the failed nodes.
5. Use a node with a new software level to create a new cluster.

Related reference

Chapter 6, “Upgrading the SAN Volume Controller software,” on page 203
The SAN Volume Controller software can be upgraded while you run day-to-day operations.

Related information

“Resetting a refused SSH key” on page 144
You can reset a refused SSH key relationship between the SAN Volume Controller Console and the SAN Volume Controller cluster.

Chapter 7. Configuring and servicing storage subsystems

You must ensure that your disk controllers and switches are correctly configured to work with the SAN Volume Controller to avoid performance issues.

Virtualization provides many benefits over direct attached or direct SAN attached storage. However, virtualization is more susceptible to the creation of performance hot-spots than direct attached storage. Hot-spots can cause I/O errors on your hosts and can potentially cause a loss of access to data.

Related reference

“Configuration guidelines”

You must follow the guidelines and procedures for your storage subsystem to maximize performance and to avoid potential I/O problems.

Related information

“Identifying your storage subsystem”

The serial number that is presented by the command-line interface (CLI) and the SAN Volume Controller Console for the SAN Volume Controller is the serial number of the device.

Identifying your storage subsystem

The serial number that is presented by the command-line interface (CLI) and the SAN Volume Controller Console for the SAN Volume Controller is the serial number of the device.

The serial numbers can be viewed on your storage subsystem. If the serial numbers are not displayed, the worldwide node name (WWNN) or worldwide port name (WWPN) is displayed. The WWNN or WWPN can be used to identify the different subsystems.

Configuration guidelines

You must follow the guidelines and procedures for your storage subsystem to maximize performance and to avoid potential I/O problems.

Guidelines

- Avoid splitting arrays into multiple logical disks at the storage subsystem layer. Where possible, create a single logical disk from the entire capacity of the array.
- Depending on the redundancy that is required, RAID-5 arrays should be created using between 5 and 8 plus parity components. That is 5 + P, 6 + P, 7 + P or 8 + P.
- Ensure that managed disk (MDisk) groups contain MDisk with similar characteristics and approximately the same capacity. You must consider the following factors:
 - The underlying RAID type that the storage subsystem is using to implement the MDisk.
 - The number of physical disks in the RAID array and the physical disk type. For example: 10K/15K rpm, FC/SATA.
- If MDisk are not the same capacity, they can be specified multiple times when the MDisk group is created. For example, if you have two 400 MB disks and one

800 MB disk that are identified as MDisk 0, 1, and 2, you can create the MDisk group with the candidate IDs of 0:1:2:2. This doubles the number of extents on the 800 MB drive.

- Do not mix MDisks that greatly vary in performance in the same MDisk group. The overall MDisk group performance is limited by the slowest MDisk. Because some disk controllers can sustain much higher I/O bandwidths than others, do not mix MDisks that are provided by low-end subsystems with those that are provided by high-end subsystems.
- Avoid leaving virtual disks (VDisks) in image mode. Only use image mode to import existing data into the cluster. This data should be migrated across the other MDisks in the group as soon as possible to optimize the benefits of virtualization.
- Follow the FlashCopy requirements before you set up the storage. Balance the spread of the FlashCopy VDisks across the MDisk groups and then the storage subsystems. The I/O characteristics of the application that is writing to the source VDisk also effects the impact that FlashCopy operations have on your overall I/O throughput.
- Perform the appropriate calculations to ensure that your storage subsystems are configured correctly.

Related reference

“Storage subsystem logical disks”

Most storage subsystems provide some mechanism to create multiple logical disks from a single array. This is useful when the storage subsystem directly presents storage to the hosts.

“RAID array configuration” on page 217

When using virtualization, ensure that the storage devices are configured to provide some type of redundancy against hard disk failures.

“Optimal MDisk group configurations” on page 217

A managed disk (MDisk) group provides the pool of storage from which virtual disks (VDisks) are created. You must ensure that the entire pool of storage provides the same performance and reliability characteristics.

“Considerations for FlashCopy mappings” on page 218

Ensure that you have considered the type of I/O and frequency of update before you create the virtual disks (VDisks) that you want to use in FlashCopy mappings.

“Image mode and migrating existing data” on page 219

Image mode virtual disks (VDisks) enable you to import and then migrate existing data that is managed by the SAN Volume Controller.

Storage subsystem logical disks

Most storage subsystems provide some mechanism to create multiple logical disks from a single array. This is useful when the storage subsystem directly presents storage to the hosts.

However, in a virtualized SAN, use a one-to-one mapping between arrays and logical disks. For the arrays that are configured with a one-to-one mapping to logical disks, the subsequent load calculations and the managed disk (MDisk) and MDisk group configuration tasks are simplified.

Scenario: the logical disks are uneven

In this scenario, you have two RAID-5 arrays and both contain 5 + P components. Array A has a single logical disk that is presented to the SAN Volume Controller

cluster. This logical disk is seen by the cluster as mdisk0. Array B has three logical disks that are presented to the SAN Volume Controller cluster. These logical disks are seen by the cluster as mdisk1, mdisk2 and mdisk3. All four MDisks are assigned to the same MDisk group that is named mdisk_grp0. When a virtual disk (VDisk) is created by striping across this group, array A presents the first extent and array B presents the next 3 extents. As a result, when reading and writing to the VDisk, the loading is split 25% on the disks in array A and 75% on the disks in array B. The performance of the VDisk is about one third of what array B can sustain.

The uneven logical disks cause performance degradation and complexity in a simple configuration. You can avoid uneven logical disks by creating a single logical disk from each array.

RAID array configuration

When using virtualization, ensure that the storage devices are configured to provide some type of redundancy against hard disk failures.

A failure of a storage device can affect a larger amount of storage that is presented to the hosts. To provide redundancy, storage devices should be configured as RAID arrays that use either mirroring or parity to protect against single failures.

When creating RAID arrays with parity protection (for example, RAID-5 arrays) consider how many component disks you want to use in each array. If you use large amount of disks, you can reduce the number of disks that are required to provide availability for the same total capacity (1 per array). However, more disks means that a longer time is taken to rebuild a replacement disk after a disk failure, and during this period a second disk failure causes a loss of all array data. More data is affected by a disk failure for a larger number of member disks because performance is reduced while you rebuild onto a hot spare and more data is exposed if a second disk fails before the rebuild is complete. The smaller the number of disks, the more likely it is that write operations span an entire stripe (stripe size x number of members minus one). In this case, write performance is improved. The number of disk drives required to provide availability can be unacceptable if arrays are too small.

Note:

- For optimal performance, use arrays with between 6 and 8 member disks.
- When creating RAID arrays with mirroring, the number of component disks in each array does not affect redundancy or performance.

Optimal MDisk group configurations

A managed disk (MDisk) group provides the pool of storage from which virtual disks (VDisks) are created. You must ensure that the entire pool of storage provides the same performance and reliability characteristics.

- The performance of an MDisk group is generally governed by the slowest MDisk in the group.
- The reliability of an MDisk group is generally governed by the weakest MDisk in the group.
- If a single MDisk in a group fails, access to the entire group is lost.

Use the following guidelines when grouping similar disks:

- Group equally performing MDisks in a single group.

- Group similar arrays in a single group. For example, configure all 6 + P RAID-5 arrays in one group.
- Group MDisks from the same type of storage subsystem in a single group.
- Group MDisks that use the same type of underlying physical disk in a single group. For example, group MDisks by fibre-channel or SATA.
- Do not use single disks. Single disks do not provide redundancy. Failure of a single disk results in total data loss of the MDisk group to which it is assigned.

Scenario: Similar disks are not grouped together

You have two storage subsystems that are attached behind your SAN Volume Controller. One device is an IBM ESS, which contains ten 6 + P RAID-5 arrays and mdisks 0 through 9. The other device is an IBM FASTT200, which contains a single RAID-1 array, mdisk10, one single JBOD, mdisk11, and a large 15 + P RAID-5 array, mdisk12.

If you assigned mdisks 0 through 9 and mdisk11 into a single MDisk group, and the JBOD mdisk11 fails, you lose access to all of the ESS arrays, even though they are online. The performance is limited to the performance of the JBOD in the FASTT storage subsystem, thus slowing down the ESS arrays.

To fix this problem, you can create three groups. The first group must contain the ESS arrays, MDdisks 0 through 9, the second group must contain the RAID-1 array, and the third group must contain the large RAID-5 array.

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Considerations for FlashCopy mappings

Ensure that you have considered the type of I/O and frequency of update before you create the virtual disks (VDisks) that you want to use in FlashCopy mappings.

FlashCopy performs in direct proportion to the performance of the source and target disks. If you have a fast source disk and slow target disk, the performance of the source disk is reduced because it has to wait for the write to occur at the target before it can write to the source.

The FlashCopy implementation that is provided by the SAN Volume Controller copies at least 256 K every time a write is made to the source. This means that *any* write involves at minimum a read of 256 K from the source, write of the same 256 K at the target, and a write of the original change at the target. Therefore, when an application performs small 4 K writes, this is translated into 256 K.

Because of this overhead, consider the type of I/O your application performs during a FlashCopy. Ensure that you do not overload the storage. The calculations contain a heavy weighting when FlashCopy is active. The weighting depends on the type of I/O that is performed. Random writes have a much higher overhead than sequential writes. For example, the sequential write would have copied the entire 256 K.

You can spread the FlashCopy source VDisks and the FlashCopy destination VDisks between as many managed disk (MDisk) groups as possible. This limits the potential bottle-necking of a single storage subsystem, (assuming that the MDisk groups contain MDisks from different storage subsystems). However, this can still

result in potential bottle-necking if you want to maintain all your target VDisks on a single storage subsystem. You must ensure that you add the appropriate weighting to your calculations.

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Image mode and migrating existing data

Image mode virtual disks (VDisks) enable you to import and then migrate existing data that is managed by the SAN Volume Controller.

Ensure that you follow the guidelines for using image mode VDisks. This might be difficult because a configuration of logical disks and arrays that performs well in a direct SAN-attached environment can contain hot-spots or hot-component disks when connected through the SAN Volume Controller cluster.

If the existing storage subsystems are configured incorrectly with respect to the guidelines you might want to consider stopping I/O operations on the host systems while migrating the data into the cluster. If I/O operations are continued and the storage subsystem does not follow the guidelines, I/O operations can fail at the hosts and ultimately loss of access to the data occurs.

The procedure for importing managed disks (MDisks) that contain existing data depends on how much free capacity you have in the SAN Volume Controller cluster. You should have the same amount of free space in the cluster as the data that you want to migrate into the cluster. If you do not have this amount of capacity available, the migration causes the MDisk group to have an uneven distribution of data. Some MDisks are more heavily loaded than others. Further migration operations are required to ensure an even distribution of data and therefore, subsequent I/O loading.

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Migrating data with an equivalent amount of free capacity

To prevent managed disks (MDisks) from having an uneven distribution of data, ensure that the cluster has the same amount of free space as the data that you want to migrate.

Perform the following steps to migrate data:

1. Stop all I/O operations from the hosts. Un-map the logical disks that contain the data from the hosts.
2. Create one or more MDisk groups with free capacity. Ensure that the MDisk groups have enough free capacity to contain all of the migrating data and that they have balanced data distribution.
3. Create an empty MDisk group. This temporarily contains the data that is imported.
4. Perform the following steps to create an image mode virtual disk (VDisk) from the first unmanaged-mode MDisk that contains the data to be imported:
 - a. Map one logical disk from the storage subsystem to the SAN Volume Controller ports.

- b. Issue the **svctask detectmdisk** command-line interface (CLI) command on the cluster. The new unmanaged-mode MDisk that is found corresponds with the logical disk mapped in the previous step.
 - c. Create an image mode VDisk from this unmanaged-mode MDisk and assign it to the empty MDisk group that you just created.
 - d. Repeat for all logical disks as required.
5. If you have decided to continue the I/O operations while you migrate the data onto SAN Volume Controller, map all the image mode VDIs to the hosts using the SAN Volume Controller and continue to access the data through the SAN Volume Controller.
6. Perform the following steps to migrate the data to the MDisk groups that you created in step 2 on page 219:
 - a. Select the first image mode VDisk to be migrated.
 - b. Migrate this VDisk from its current MDisk group into one of the MDisk groups created in step 2 on page 219 above. This migrates all the data from the logical disk into the new free space.
 - c. Select the next image mode VDisk and repeat the previous step after the migration completes.
7. When all the VDIs have been migrated, the MDisk groups created in step 2 on page 219 contain the data that was on the image mode VDIs. The data is striped across the new groups and is virtualized.
8. Destroy the temporary MDisk group that contained the original image mode VDIs.
9. Go back to the storage subsystem and reconfigure the old arrays and logical disks according to the guidelines.
10. Add this storage back under the SAN Volume Controller and use the old storage to create new VDIs.

Related reference

“Configuration guidelines” on page 215

You must follow the guidelines and procedures for your storage subsystem to maximize performance and to avoid potential I/O problems.

Migrating data with a smaller amount of free capacity

If the free capacity in the SAN Volume Controller cluster is smaller than the capacity of the data that is imported, you can still migrate data.

The free capacity in the SAN Volume Controller cluster is smaller than the capacity of the data that is imported. The RAID arrays that contain the existing data still contain the data after importing the arrays to the cluster.

Scenario:

You have one managed disk (MDisk) in the destination MDisk group. You add image mode logical units from an array on the storage subsystem and migrate these logical units to the destination MDisk group. The logical units are then striped across the one managed-mode disk. Next, you add another logical unit to the destination MDisk group. The MDisk now contains two managed-mode disks, but all of the data is on the first managed-mode disk. As a result, some of the data must be migrated from the overloaded managed-mode disks to the under-utilized managed-mode disks.

Attention: The migration causes an uneven distribution of data across the MDisks in the MDisk group. The impact of this depends on the number of MDisks that are initially in the MDisks group and how many of these have free capacity.

This procedure might require subsequent migration of data within the MDisk group in order to balance the distribution of data across the MDisks in the group.

Perform the following steps to migrate data:

1. Select an MDisk group that contains enough free capacity to migrate *all* of the logical disks on the first array that you want to migrate to the cluster.
2. Create an empty MDisk group that can temporarily contain the data that is imported.
3. Stop all I/O operations to the logical disks that you want to migrate first, and unmap these disks from their hosts.
4. Perform the following steps to create an image mode virtual disk (VDisk) from the first unmanaged-mode MDisk that contains the data that you want to import:
 - a. Map one logical disk from the storage subsystem to the SAN Volume Controller ports.
 - b. Issue the **svctask detectmdisk** command-line interface (CLI) command on the cluster. The new unmanaged-mode MDisk that is found corresponds with the logical disk that was mapped in the previous step.
 - c. Create an image mode VDisk from this unmanaged-mode MDisk and assign it to use the empty MDisk group just created.
 - d. Repeat for all logical disks.
5. If you have decided to continue the I/O operations while you migrate the data to the SAN Volume Controller cluster, map all the image mode VDIs to the hosts using the SAN Volume Controller and continue to access the data through the SAN Volume Controller.
6. Perform the following steps to migrate the data into the MDisk groups that you created in step 1:
 - a. Select the first image mode VDisk that you want to migrate.
 - b. Migrate this VDisk from its current MDisk group into one of the MDisk groups created in step 1 above. This migrates all the data from the logical disk into the new free space.
 - c. Select the next image mode VDisk and repeat the previous step when the migration completes.
7. Perform the following steps to reconfigure the RAID array that contains the logical disks and add it to the MDisk group selected in step 1:
 - a. Remove the MDisks from the temporary MDisk group.
 - b. At the storage subsystem, the logical disks that have been migrated should be unmapped from the SAN Volume Controller cluster and deleted from the array (if more than one existed).
 - c. Create a single logical disk that uses the entire array capacity.
 - d. Map this new logical disk to the SAN Volume Controller ports.
 - e. Issue the **svctask detectmdisk** CLI command on the cluster. The new managed-mode MDisk that is found corresponds with the new logical disk that you created.
 - f. Add this managed-mode MDisk to the MDisk group selected in step 1.
8. Repeat steps 3 through 7 for the next array.

Related reference

“Configuration guidelines” on page 215

You must follow the guidelines and procedures for your storage subsystem to maximize performance and to avoid potential I/O problems.

Configuring a balanced storage subsystem

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

There are 2 major steps to attaching a storage subsystem to a SAN Volume Controller:

1. Setting the characteristics of the SAN Volume Controller to storage connections
2. Mapping logical units to these connections that allow the SAN Volume Controller to access them

The virtualization features of the SAN Volume Controller enable you to choose how your storage is divided and presented to hosts. While virtualization provides you with a great deal of flexibility, it also offers the potential to set up an overloaded storage subsystem. A storage subsystem is overloaded if the quantity of I/O transactions that are issued by the host systems exceeds the capability of the storage to process those transactions. If a storage subsystem is overloaded, it causes delays in the host systems and might cause I/O transactions to time out in the host. If I/O transactions time out, the host logs errors and I/Os fail back to the applications.

Example of an overloaded storage subsystem

You have used the SAN Volume Controller to virtualize a single RAID array and to divide the storage across 64 host systems. If all host systems attempt to access the storage at the same time, the single RAID array is overloaded.

Perform the following steps to configure a balanced storage subsystem:

1. Use Table 13 to calculate the I/O rate for each RAID array in the storage subsystem.

Note: The actual number of I/O operations per second that can be processed depends 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. For example, a RAID-5 array with eight component disks has an approximate I/O rate of $150 \times 7 = 1050$.

Table 13. 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	$150 \times N$
RAID-10, RAID 0+1, RAID 1+0 (striped + mirrored) arrays	N	$150 \times N$

2. Calculate the I/O rate for a managed disk (MDisk).
 - If there is a one-to-one relationship between backend arrays and MDisks, the I/O rate for an MDisk is the same as the I/O rate of the corresponding array.

- If an array is divided into multiple MDisks, the I/O rate per MDisk is the I/O rate of the array divided by the number of MDisks that are using the array.
3. Calculate the I/O rate for an MDisk group. The I/O rate for an MDisk group is the sum of the I/O rates of the MDisk that is in the MDisk group. For example, an MDisk group contains eight MDisks and each MDisk corresponds to a RAID-1 array. Using Table 13 on page 222, the I/O rate for each MDisk is calculated as 300. The I/O rate for the MDisk group is $300 \times 8 = 2400$.
 4. Use Table 14 to calculate the impact of FlashCopy relationships. If you are using the FlashCopy feature that is provided by the SAN Volume Controller, you must consider the additional amount of I/O that FlashCopy generates because it reduces the rate at which I/O from host systems can be processed. When a FlashCopy relationship copies write I/Os from the host systems to areas of the source or target virtual disk (VDisk) that are not yet copied, the SAN Volume Controller generates extra I/Os to copy the data before the write I/O is performed. The effect of using FlashCopy depends on the type of I/O workload that is generated by an application.

Table 14. Calculate the impact of FlashCopy relationships

Type of application	Impact to I/O rate	Additional weighting for FlashCopy
Application is not performing 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
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 VDisk that is the source or target of an active FlashCopy relationship consider the type of application that you want to use the VDisk and record the additional weighting for the VDisk.

Example

For example, a FlashCopy relationship is 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 VDisk. A second host application reads the target VDisk and writes the data to tape to create a backup. The additional weighting for the source VDisk is 14. The additional weighting for the destination VDisk is 0.

5. Calculate the I/O rate for VDIs in an MDisk group by performing the following steps:
 - a. Calculate the number of VDIs in the MDisk group.
 - b. Add the additional weighting for each VDisk that is the source or target of an active FlashCopy relationship.
 - c. Divide the I/O rate of the MDisk group by this number to calculate the I/O rate per VDisk.

Example 1

An MDisk group has an I/O rate of 2400 and contains 20 VDIs. There are no FlashCopy relationships. The I/O rate per VDisk is $2400 / 20 = 120$.

Example 2

An MDisk group has an I/O rate of 5000 and contains 20 VDIs. There are two active FlashCopy relationships that have source VDIs in the MDisk group. Both source VDIs are accessed by applications that issue random read and writes. As a result, the additional weighting for each VDI is 14. The I/O rate per VDI is $5000 / (20 + 14 + 14) = 104$.

6. Determine if the storage subsystem is overloaded. The figure that was determined in step 4 on page 223 provides some indication of how many I/O operations per second can be processed by each VDI in the MDisk 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, you can use the I/O statistics facilities that are provided by the SAN Volume Controller to measure the I/O rate of your virtual disks or use Table 15 as a guideline.

Table 15. Determine if the storage subsystem is overloaded

Type of Application	I/O rate per VDI
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

7. Interpret the result. If the I/O rate that is generated by the application exceeds the I/O rate per VDI that you calculated, you might be overloading your storage subsystem. You must carefully monitor the storage subsystem to determine if the backend storage limits the overall performance of the storage subsystem. It is also possible that the calculation above is too simplistic to model your storage use. For example, the calculation assumes that your applications generate the same I/O workload to all VDIs, which might not be the case.

You can use the I/O statistics facilities that are provided by the SAN Volume Controller to measure the I/O rate of your MDisks. You can also use the performance and I/O statistics facilities that are provided by your storage subsystems.

If your storage subsystem is overloaded there are several actions you can take to resolve the problem:

- Add more backend storage to the subsystem to increase the quantity of I/O that can be processed by the storage subsystem. The SAN Volume Controller provides virtualization and data migration facilities to redistribute the I/O workload of VDIs across a greater number of MDisks without having to take the storage offline.
- Stop unnecessary FlashCopy relationships to reduce the amount of I/O operations that are submitted to the backend storage. If you perform FlashCopys in parallel, consider reducing the amount of FlashCopy relationships that start in parallel.
- Adjust the queue depth to limit the I/O workload that is generated by a host. Depending on the type of host and type of host bus adapters (HBAs), it might be possible to limit the queue depth per VDI and/or limit the queue depth per HBA. The SAN Volume Controller also provides I/O governing features that can limit the I/O workload that is generated by hosts.

Note: Although these actions can be used to avoid I/O time-outs, performance of your storage subsystem is still limited by the amount of storage.

Discovering logical units

The SAN Volume Controller initialization includes a process called discovery.

The discovery process systematically explores all visible ports on the SAN for devices that identify themselves as storage subsystems. Each storage subsystem is interrogated to determine the number of logical units (LUs) that it exports. The LUs are interrogated to determine if new storage or a new path for previously discovered storage, is found. The set of LUs forms the SAN Volume Controller managed disk (MDisk) view.

The discovery process runs when ports are added to or deleted from the SAN and when certain error conditions occur. You can also manually run the discovery process using the **svctask detectmdisk** command-line interface (CLI) command or the **Discover MDisks** function from the SAN Volume Controller Console. The **svctask detectmdisk** CLI command and the **Discover MDisks** function have the cluster rescan the fibre-channel network. The rescan discovers any new MDisks that might have been added to the cluster and rebalances MDisk access across the available controller device ports.

Note: Some storage subsystems do not automatically export LUs to the SAN Volume Controller.

Guidelines for exporting LUs

Ensure that you are familiar with the following guidelines for exporting LUs to the SAN Volume Controller:

- When you define the SAN Volume Controller as a host object to the storage subsystem, you must include *all* ports on *all* nodes and candidate nodes.
- When you first create an LU, you *must* wait until it is initialized before you export it to the SAN Volume Controller.

Attention: Failure to wait for the LUs to initialize can result in excessive discovery times and an unstable view of the SAN.

- When you export an LU to the SAN Volume Controller, the LU *must* be accessible through all ports on the storage subsystem that is visible to the SAN Volume Controller.

Important: The LU *must* be identified by the same logical unit number (LUN) on all ports.

Expanding a logical unit using the CLI

You can use the command-line interface (CLI) to expand a logical unit.

Some storage subsystems enable you to expand the size of a logical unit (LU) using vendor-specific disk-configuration software that is provided. However, the SAN Volume Controller cannot use extra capacity that is provided in this way.

The LU has increased in size and this additional space must be made available for use.

Perform the following steps to ensure that this additional capacity is available to the SAN Volume Controller:

1. Issue the **svctask migrateexts** CLI command to migrate all the data from the managed disk (MDisk).

Note:

- For managed mode MDisks, issue the **svctask rmmdisk** CLI command to remove the MDisk from the MDisk group.
 - For image mode MDisks, issue the **svctask chmdisk** CLI command to change the mode of the image mode disk to unmanaged.
2. Issue the **svctask includemdisk** CLI command.
 3. Issue the **svctask detectmdisk** CLI command to rescan the fibre-channel network. The rescan discovers any new MDisks that have been added to the cluster and rebalances MDisk access across the available controller device ports. This can take a few minutes.
 4. Issue the **svctask lsmdisk** CLI command to display the additional capacity that has been expanded.

The extra capacity is available for use by the SAN Volume Controller.

Modifying a logical unit mapping using the CLI

You can modify a logical unit (LU) mapping using the command-line interface (CLI).

Perform the following steps to modify an LU mapping:

1. Migrate all of the data from the managed disk (MDisk) by performing the following steps:
 - a. If the MDisk is in managed mode or image mode and the virtual disk (VDisk) must be kept online, issue the following CLI command and then proceed to step 2 on page 227:

```
svctask rmmdisk -mdisk MDisk number -force MDisk group number
```

Where *MDisk number* is the number of the MDisk that you want to modify and *MDisk group number* is the number of the MDisk group for which you want to remove the MDisk.

Note:

- The VDisk becomes a striped MDisk *not* an image-mode VDisk.
 - All data that is stored on this MDisk is migrated to the other MDisks in the MDisk group.
 - This CLI command can fail if there are not enough free extents in the MDisk group.
- b. If the MDisk is in image mode and you do not want to convert the VDisk to a striped VDisk, stop all I/O to the image mode VDisk.
 - c. Issue the following CLI command to remove the host mapping:

```
svctask rmvdiskhostmap -host host name VDisk name
```

Where *host name* is the name of the host for which you want to remove the VDisk mapping and *VDisk name* is the name of the VDisk for which you want to remove mapping.

- d. Issue the following command to delete the VDisk:

```
svctask rmvdisk VDisk name
```

Where *VDisk name* is the name of the VDisk that you want to delete.

2. Remove the LU mapping on the storage subsystem so that the LUN is not visible to the SAN Volume Controller.
3. Issue the following CLI command to clear all error counters on the MDisk:

```
svctask includemdisk MDisk number
```

 Where *MDisk number* is the number of the MDisk that you want to modify.
4. Issue the following CLI command to rescan the fibre-channel network and detect that the LU is no longer there.

```
svctask detectmdisk MDisk number
```

 Where *MDisk number* is the number of the MDisk that you want to modify. The MDisk is removed from the configuration.
5. Issue the following CLI command to verify that the MDisk is removed:

```
svcinfolsmdisk MDisk number
```

 Where *MDisk number* is the number of the MDisk that you want to modify.
 - If the MDisk is still displayed, repeat steps 3 and 4.
6. Configure the mapping of the LU to the new LUN.
7. Issue the following CLI command:

```
svctask detectmdisk
```
8. Issue the following CLI command to check that the MDisk now has the correct LUN:

```
svcinfolsmdisk
```

The MDisk has the correct LUN.

Accessing controller devices with multiple remote ports

If a managed disk (MDisk) logical unit (LU) is accessible through multiple controller device ports, the SAN Volume Controller ensures that all nodes that access this LU coordinate their activity and access the LU through the same controller device port.

Monitoring LU access through multiple controller device ports

When the SAN Volume Controller can access an LU through multiple controller device ports, the SAN Volume Controller uses the following criteria to determine the accessibility of these controller device ports:

- The SAN Volume Controller node is a member of a cluster.
- The SAN Volume Controller node has fibre-channel connections to the controller device port.
- The SAN Volume Controller node has successfully discovered the LU.
- Slandering has not caused the SAN Volume Controller node to exclude access to the MDisk through the controller device port.

An MDisk path is presented to the cluster for all SAN Volume Controller nodes that meet these criteria.

Controller device port selection

When an MDisk is created, the SAN Volume Controller selects one of the controller device ports to access the MDisk.

Table 16 on page 228 describes the algorithm that the SAN Volume Controller uses to select the controller device port.

Table 16. Controller device port selection algorithm

Criteria	Description
Accessibility	Creates an initial set of candidate controller device ports. The set of candidate controller device ports include the ports that are accessible by the highest number of nodes.
Slandering	Reduces the set of candidate controller device ports to those with the lowest number of nodes.
Preference	Reduces the set of candidate controller device ports to those that the controller device uses as preferred ports.
Load balance	Selects the port from the set of candidate controller device ports that has the lowest MDisk access count.

After the initial device port selection is made for an MDisk, the following events can cause the selection algorithm to rerun:

- A new node joins the cluster and has a different view of the controller device than the other nodes in the cluster.
- The **svctask detectmdisk** command-line interface (CLI) command is run or the **Discover MDisks** SAN Volume Controller Console function is used. The **svctask detectmdisk** CLI command and the **Discover MDisks** function have the cluster rescan the fibre-channel network. The rescan discovers any new MDisks that might have been added to the cluster and rebalances MDisk access across the available controller device ports.
- Error recovery procedures (ERPs) are started because a controller device has changed its preferred port.
- New controller device ports are discovered for the controller device that is associated with the MDisk.
- The controller device port that is currently selected becomes inaccessible.
- Slandering has caused the SAN Volume Controller to exclude access to the MDisk through the controller device port.

Determining a storage subsystem name from its SAN Volume Controller name

You can determine a storage subsystem name from its SAN Volume Controller name.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to determine the name of the storage subsystem:

1. Click **Work with Managed Disks** → **Disk Controller Systems**. The Viewing Disk Controller Systems panel is displayed.
2. Select the link for the name of the storage subsystem for which you want to determine the name.
3. Record the Worldwide Node Name (WWNN). You can launch the native user interface or use the command-line tools to verify the name of the storage subsystem that uses this WWNN.

Determining a storage subsystem name from its SAN Volume Controller name using the CLI

You can determine a storage subsystem name from its SAN Volume Controller name using the command-line interface (CLI).

1. Issue the following CLI command to list the storage subsystem:

```
svcinfolsccontroller
```

2. Record the name or ID for the storage subsystem that you want to determine.
3. Issue the following CLI command:

```
svcinfolsccontroller controllername/id
```

where *controllername/id* is the name or ID that you recorded in step 2.

4. Record the worldwide node name (WWNN) for the device. 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.

Renaming a storage subsystem

You can rename a storage subsystem from the Renaming a Disk Controller System panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to rename a storage subsystem:

1. Click **Work with Managed Disks** → **Disk Controller Systems** in the portfolio. The Viewing Disk Controller Systems panel is displayed.
2. Select the storage subsystem that you want to rename and select **Rename a Disk Controller System** from the list. Click **Go**. The Renaming Disk Controller System panel is displayed.

Related concepts

“Storage subsystems” on page 23

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

Changing a configuration for an existing storage subsystem using the CLI

You can change a configuration for an existing storage subsystem using the command-line interface (CLI).

You must change the configuration for a storage subsystem in order to delete and replace logical units (LUs).

Perform the following steps to delete existing LUs and replace them with new LUs:

1. Issue the following CLI command to delete the managed disks (MDisks) that are associated with the LUs from their MDisk groups:

```
svctask rmmdisk -mdisk MDisk name1, MDisk name2 -force MDisk group name
```

Where *MDisk name1*, *MDisk name2* are the names of the MDisks that you want to delete.

2. Delete the existing LUs using the configuration software of the storage subsystem.
3. Issue the following command to delete the associated MDisks from the cluster:

```
svctask detectmdisk
```
4. Configure the new LUs using the configuration software of the storage subsystem.
5. Issue the following command to add the new LUs to the cluster:

```
svctask detectmdisk
```

Related concepts

“Storage subsystems” on page 23

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

“MDisk groups” on page 27

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

“MDisks” on page 24

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

Adding a new storage controller to a running configuration

You can add a new storage controller to your SAN at any time.

You must follow the zoning guidelines for your switch and also ensure that the controller is setup correctly for use with the SAN Volume Controller.

You must create one or more arrays on the new controller. For maximum redundancy and reliability, use RAID-5, RAID-1 or RAID-0+1 (sometimes called RAID-10). Generally 5+P arrays are recommended.

If your controller provides array partitioning, create a single partition from the entire capacity available in the array. You must record the LUN number that you assign to each partition. You must also follow the mapping guidelines (if your storage controller requires LUN mapping) to map the partitions or arrays to the SAN Volume Controller ports. You can determine the SAN Volume Controller ports by following the procedure for determining WWPNs.

Perform the following steps to add a new storage controller:

1. Ensure that the cluster has detected the new storage (MDisks).
 - a. Click **Work with Managed Disks** → **Managed Disks**. The Filtering Managed Disks panel is displayed.
 - b. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Managed Disks panel is displayed.
 - c. Select **Discover MDisks** from the task list and click **Go**.
2. Determine the storage controller name to validate that this is the correct controller. The controller will have automatically been assigned a default name.

- If you are unsure which controller is presenting the MDisks perform the following steps:
 - a. Click **Work with Managed Disks** → **Disk Controller Systems**. The Viewing Disk Controller Systems panel is displayed.
 - b. Find the new controller in the list. The new controller has the highest numbered default name.
- 3. Record the field controller LUN number. The controller LUN number corresponds with the LUN number that you assigned to each of the arrays or partitions.
 - a. Click **Work with Managed Disks** → **Managed Disks**. The Filtering Managed Disks panel is displayed.

Note: You might have to close the Viewing Managed Disks panel by clicking on the **X** to refresh the panel and display the Filtering Managed Disks panel.

- b. Select unmanaged from the **Mode** list and click **OK**. The Viewing Managed Disks panel is displayed. The MDisks shown on the Viewing Managed Disks panel should correspond with the RAID arrays or partitions that you created.
4. Create a new MDisk group and add only the RAID arrays that belong to the new controller to this MDisk group. To avoid mixing RAID types, create a new MDisk group for each set of RAID array types (for example, RAID-5, RAID-1).
 - a. Click **Work with Managed Disks** → **Managed Disk Groups**. The Filtering Managed Disk Groups panel is displayed.
 - b. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type.
 - c. Select **Create an MDisk Group** from the task list and click **Go**. The Create Managed Disk Group wizard begins.
 - d. Complete the wizard to create a new MDisk group.

Tip: Give each MDisk group that you create a descriptive name. For example, if your controller is named FAST650-fred, and the MDisk group contains RAID-5 arrays, name the MDisk Group F600-fred-R5.

Related concepts

Chapter 7, “Configuring and servicing storage subsystems,” on page 215
You must ensure that your disk controllers and switches are correctly configured to work with the SAN Volume Controller to avoid performance issues.

Related reference

“Switch zoning for the SAN Volume Controller” on page 65
Ensure that you are familiar with the constraints for zoning a switch.

Adding a new storage controller to a running configuration using the CLI

You can add a new disk controller system to your SAN at any time using the command-line interface (CLI).

You must follow the zoning guidelines for your switch and also ensure that the controller is setup correctly for use with the SAN Volume Controller.

You must create one or more arrays on the new controller. For maximum redundancy and reliability, use RAID-5, RAID-1 or RAID-0+1 (sometimes called RAID-10). Generally 5+P arrays are recommended.

If your controller provides array partitioning, create a single partition from the entire capacity available in the array. You must record the LUN number that you assign to each partition. You must also follow the mapping guidelines (if your storage controller requires LUN mapping) to map the partitions or arrays to the SAN Volume Controller ports. You can determine the SAN Volume Controller ports by following the procedure for determining WWPNs.

Perform the following steps to add a new storage controller:

1. Issue the following CLI command to ensure that the cluster has detected the new storage (MDisks):

```
svctask detectmdisk
```

2. Determine the storage controller name to validate that this is the correct controller. The controller is automatically assigned a default name.
 - If you are unsure which controller is presenting the MDisks, issue the following command to list the controllers:

```
svcinfolsccontroller
```

3. Find the new controller in the list. The new controller has the highest numbered default name.
4. Record the name of the controller and follow the instructions in the section about determining a disk controller system name.
5. Issue the following command to change the controller name to something that you can easily use to identify it:

```
svctask chcontroller -name newname oldname
```

Where *newname* is the name that you want to change the controller to and *oldname* is the name that you are changing.

6. Issue the following command to list the unmanaged MDisks:

```
svcinfolsmdisk -filtervalue mode=unmanaged:controller_name=new_name
```

These MDisks should correspond with the RAID arrays or partitions that you have created.

7. Record the field controller LUN number. This number corresponds with the LUN number that you assigned to each of the arrays or partitions.
8. Create a new MDisk group and add only the RAID arrays that belong to the new controller to this MDisk group. To avoid mixing RAID types, create a new MDisk group for each set of RAID array types (for example, RAID-5, RAID-1). Give each MDisk group that you create a descriptive name. For example, if your controller is named FAST650-fred, and the MDisk group contains RAID-5 arrays, name the MDisk Group F600-fred-R5.

```
svctask mkmdiskgrp -ext 16 -name mdisk_grp_name  
-mdisk colon separated list of RAID-x mdisks returned  
in step 4
```

This creates a new MDisk group with an extent size of 16MB.

Related concepts

Chapter 7, “Configuring and servicing storage subsystems,” on page 215
You must ensure that your disk controllers and switches are correctly configured to work with the SAN Volume Controller to avoid performance issues.

Related tasks

“Determining the WWPNs of a node using the CLI” on page 168
You can determine the worldwide port names (WWPNs) of a node using the command-line interface (CLI).

Related reference

“Switch zoning for the SAN Volume Controller” on page 65
Ensure that you are familiar with the constraints for zoning a switch.

Removing a storage subsystem

You can replace or decommission a storage subsystem.

This task assumes that you have already launched the SAN Volume Controller application.

During this procedure, you will add a new device, migrate data off of the storage subsystem and remove the old MDisks.

An alternative to following this procedure is to migrate all of the virtual disks (VDisks) that are using storage in this MDisk group to another MDisk group. This allows you to consolidate the VDisks in a single or new group. However, you can only migrate one VDisk at a time. The procedure outlined below migrates all the data through a single command.

You can also use this procedure to remove or replace a single MDisk in a group. If an MDisk experiences 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 replace just that MDisk. Steps 1 and 3 detail how you can add or remove a single MDisk rather than a list of MDisks.

Perform the following steps to remove a storage subsystem:

1. Add the new MDisks to the MDisk group by performing the following steps:
 - a. Click **Work with Managed Disks** → **Managed Disks**. The Filtering Managed Disk Groups panel is displayed.
 - b. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Managed Disk Groups panel is displayed.
 - c. Select the MDisk group that you want to add the new MDisks to and select **Add MDisks** from the task list. Click **Go**. The Adding Managed Disks to Managed Disk Group panel is displayed.
 - d. Select the new MDisks and click **OK**. The MDisk group should now contain both the old and new MDisks.
2. Ensure that the capacity of the new MDisks is the same or exceeds that of the old MDisks before proceeding to step 3.
3. Force the deletion of the old MDisks from the MDisk group to migrate all the data from the old MDisks to the new MDisks.
 - a. Click **Work with Managed Disks** → **Managed Disks**. The Filtering Managed Disk Groups panel is displayed.
 - b. Specify the filter criteria that you want to use and click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Managed Disk Groups panel is displayed.
 - c. Select the MDisk group that you want to add the new MDisks to and select **Remove MDisks** from the task list. Click **Go**. The Deleting Managed Disks from Managed Disk Group panel is displayed.

d. Select the old MDisks and click **OK**. The migration process begins.

Note: The amount of time this process runs depends on the number and size of MDisks and the number and size of the VDIsks that are using the MDisks.

4. Check the progress of the migration process by issuing the following command from the command-line interface (CLI): `svcinfolsmigrate`
5. When all the migration tasks are complete, for example, the command in step 4 returns no output, verify that the MDisks are unmanaged.
6. Access the storage subsystem and unmap the LUNs from the SAN Volume Controller ports.

Note: You can delete the LUNs if you no longer want to preserve the data that is on the LUNs.

7. Issue the following CLI command:
`svctask detectmdisk`
8. Verify that there are no MDisks for the storage subsystem that you want decommission.
9. Remove the storage subsystem from the SAN so the SAN Volume Controller ports can no longer access the storage subsystem.

Related tasks

“Adding a new storage controller to a running configuration using the CLI” on page 231

You can add a new disk controller system to your SAN at any time using the command-line interface (CLI).

Removing a storage subsystem using the CLI

You can replace or decommission a storage subsystem using the command-line interface (CLI).

During this procedure, you will add a new device, migrate data off of the storage subsystem and remove the old MDisks.

An alternative to following this procedure is to migrate all of the virtual disks (VDIsks) that are using storage in this MDisk group to another MDisk group. This allows you to consolidate the VDIsks in a single or new group. However, you can only migrate one VDisk at a time. The procedure outlined below migrates all the data through a single command.

You can also use this procedure to remove or replace a single MDisk in a group. If an MDisk experiences 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 replace just that MDisk.

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 mdiskx:mdisky:mdiskz... mdisk_grp_name
```

Where *mdiskx:mdisky:mdiskz...* are the names of new MDisks that have a total capacity that is larger than the decommissioned MDisks and *<mdisk_grp_name>* is the name of the MDisk group that contains the MDisks that you want to decommission.

You should now have an MDisk group that you want to decommission and the new MDisks.

3. Ensure that the capacity of the new MDisks is the same or exceeds that of the old MDisks before you proceed to step 4.
4. Issue the following command to force delete the old MDisks from the group:

```
svctask rmmdisk -force -mdisk mdiskx:mdisky:mdiskz... mdisk_grp_name
```

Where *mdiskx:mdisky:mdiskz...>* are the old MDisks that you want to delete and *mdisk_grp_name>* is the name of the MDisk group that contains the MDisks that you want to 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 takes some time to complete, even though the command returns immediately.

5. Check the progress of the migration process by issuing the following command:

```
svcinfo lsmigrate
```
6. When all the migration tasks are complete, for example, the command in step 5 returns no output, verify that the MDisks are unmanaged.
7. Access the storage subsystem and unmap the LUNs from the SAN Volume Controller ports.

Note: You can delete the LUNs if you no longer want to preserve the data that is on the LUNs.

8. Issue the following CLI command:

```
svctask detectmdisk
```
9. Verify that there are no MDisks for the storage subsystem that you want decommission.
10. Remove the storage subsystem from the SAN so that the SAN Volume Controller ports can no longer access the storage subsystem.

Related tasks

“Removing a storage subsystem using the CLI” on page 234

You can replace or decommission a storage subsystem using the command-line interface (CLI).

“Determining the relationship between VDIsks and MDisks using the CLI” on page 169

You can determine the relationship between virtual disks (VDIsks) and managed disks (MDisks) using the command-line interface (CLI).

“Migrating VDIsks between MDisk groups using the CLI” on page 184

You can migrate virtual disks (VDIsks) between managed disk (MDisk) groups using the command-line interface (CLI).

Removing MDisks that represent unconfigured LUs using the CLI

You can use the command-line interface (CLI) to remove MDisks from the cluster.

When you remove LUs from your storage subsystem, the managed disks (MDisks) that represent those LUs might still exist in the cluster. However, the cluster cannot access these MDisks because the LUs that these MDisks represent have been unconfigured or removed from the storage subsystem. You must remove these MDisks.

Perform the following steps to remove MDisks:

1. Run the **svctask includemdisk** CLI command on all the affected MDisks.
2. Run the **svctask rmmdisk** CLI command on all affected MDisks. This puts the MDisks into the unmanaged mode.
3. Run the **svctask detectmdisk** CLI command. The cluster detects that the MDisks no longer exist in the storage subsystem.

All of the MDisks that represent unconfigured LUs are removed from the cluster.

Related tasks

“Discovering MDisks using the CLI” on page 157

You can use the command-line interface (CLI) to discover managed disks (MDisks).

Creating a quorum disk

A quorum disk is used to resolve tie-break situations when the voting set of nodes disagree on the current state of the cluster.

Quorum disk creation and extent allocation

The use of a quorum disk prevents the cluster from being split exactly in half. If the cluster is split in half, both halves can either continue operating or stop operating.

During quorum disk discovery, the system assesses each logical unit (LU) to determine its potential use as a quorum disk. From the set of eligible LUs, the system nominates three quorum candidate disks.

An LU must meet the following criteria to be considered a candidate for a quorum disk:

- It must be in managed space mode.
- It must be visible to all nodes in the cluster.
- It must be presented by a storage subsystem that is an approved host for quorum disks.
- It must have sufficient free extents to hold the cluster state and the configuration metadata.

If possible, the quorum disk candidates are presented by different devices. After the quorum candidate disks are selected, the cluster selects one of the candidate quorum disks to become the quorum disk. After the quorum disk is selected, the cluster does not attempt to ensure that the candidate quorum disks are presented by different devices. The quorum disk candidates can be updated by configuration activity if other eligible LUs are available.

If no quorum disk candidates are found after the discovery, one of the following situations has occurred:

- No LUs exist in managed space mode. An error is logged when this situation occurs.

- LUs exist in managed space mode, but they do not meet the eligibility criteria. An error is logged when this situation occurs.

Manual discovery

When you create or remove LUNs on a storage subsystem, the managed disk (MDisk) view is not automatically updated.

You must issue the **svctask detectmdisk** command-line interface (CLI) command or use the **Discover MDisks** function from the SAN Volume Controller Console to have the cluster rescan the fibre-channel network. The rescan discovers any new MDisks that might have been added to the cluster and rebalances MDisk access across the available controller device ports.

Servicing storage subsystems

Storage subsystems that are supported for attachment to the SAN Volume Controller are designed with redundant components and access paths to allow concurrent maintenance. Hosts have continuous access to their data during component failure and replacement.

The following guidelines apply to all storage subsystems that are attached to the SAN Volume Controller:

- Always follow the service instructions that are provided in the documentation for your storage subsystem.
- Ensure that there are no unfixed errors in the SAN Volume Controller error log before you perform any maintenance procedures.
- After you perform a maintenance procedure, check the SAN Volume Controller error log and fix any errors. Expect to see the following types of errors:
 - MDisk error recovery procedures (ERPs)
 - Reduced paths

The following are the two categories of service actions for storage subsystems:

- Controller code upgrade
- Field replaceable unit (FRU) replacement

Controller code upgrade

Ensure that you are familiar with the following guidelines for upgrading controller code:

- Check to see if the SAN Volume Controller supports concurrent maintenance for your storage subsystem.
- Allow the storage subsystem to coordinate the entire upgrade process.
- If it is not possible to allow the storage subsystem to coordinate the entire upgrade process, perform the following steps:
 1. Reduce the storage subsystem workload by 50%.
 2. Use the service interface to manually failover all logical units (LUs) from the controller that you want to upgrade.
 3. Upgrade the controller code.
 4. Restart the controller.
 5. Manually failback the LUs to their original controller.
 6. Repeat for all controllers.

FRU replacement

Ensure that you are familiar with the following guidelines for replacing FRUs:

- If the component you want to replace is directly in the host-side data path (for example, cable, fibre-channel port, or controller), disable the external data paths to prepare for upgrade. To disable external data paths, disconnect or disable the appropriate ports on the fabric switch. The SAN Volume Controller ERPs reroute access over the alternate path.
- If the component you want to replace is in the internal data path (for example, cache or disk drive) and did not completely fail, ensure that the data is backed up before you attempt to replace the component.
- If the component you want to replace is not in the data path, (for example, uninterruptible power supplies, fans or batteries) the component is generally dual redundant and can be replaced without additional steps.

Configuring the EMC CLARiiON subsystem

This section provides information about configuring the EMC CLARiiON storage system for its attachment to a SAN Volume Controller.

Related concepts

“MDisk groups” on page 27

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

Related tasks

“Configuring the EMC CLARiiON controller with Access Logix installed” on page 239

The SAN Volume Controller does not have access to the storage controller logical units (LUs) if Access Logix is installed on the EMC CLARiiON controller. You must use the EMC CLARiiON configuration tools to associate the SAN Volume Controller and LU.

“Configuring the EMC CLARiiON controller without Access Logix installed” on page 242

If Access Logix is not installed on an EMC CLARiiON controller, all logical units (LUs) that were created on the controller can be used by the SAN Volume Controller.

“Creating MDisk groups” on page 112

You can create a new managed disk (MDisk) group using the Create a Managed Disk Group wizard.

Related reference

“Access Logix” on page 239

Access Logix is an optional feature of the firmware code that provides the functionality that is known as LUN Mapping or LUN Virtualization.

“Supported models of the EMC CLARiiON” on page 242

The SAN Volume Controller supports models of the EMC CLARiiON.

“Supported firmware levels for the EMC CLARiiON” on page 243

The SAN Volume Controller supports EMC CLARiiON.

“Concurrent maintenance on the EMC CLARiiON” on page 243

Concurrent maintenance is the ability to perform I/O operations to a controller while simultaneously performing maintenance on it.

“User interface on EMC CLARiiON” on page 244

Ensure that you are familiar with the user interface that supports the EMC CLARiiON subsystem.

“Sharing the EMC CLARiiON between a host and the SAN Volume Controller” on page 244

The EMC CLARiiON can be shared between a host and a SAN Volume Controller.

“Switch zoning limitations for the EMC CLARiiON” on page 244

There are limitations in switch zoning for the SAN Volume Controller and EMC CLARiiON.

“Quorum disks on the EMC CLARiiON” on page 245

The EMC CLARiiON supports quorum disks.

“Advanced functions for the EMC CLARiiON” on page 246

Some advanced functions of the EMC CLARiiON are not supported by the SAN Volume Controller.

“Logical unit creation and deletion on the EMC CLARiiON” on page 246

Binding an LU to a RAID group can take a significant amount of time on the EMC CLARiiON.

“Configuring settings for the EMC CLARiiON” on page 246

A number of settings and options are available through the EMC CLARiiON configuration interface.

Access Logix

Access Logix is an optional feature of the firmware code that provides the functionality that is known as LUN Mapping or LUN Virtualization.

You can use the software tab in the storage subsystems properties page of the SAN Volume Controller Console to determine if Access Logix is installed. After Access Logix is installed it can be disabled but not removed. The following are the two modes of operation for Access Logix:

- **Access Logix not installed:** In this mode of operation, all LUNs are accessible from all target ports by any host. Therefore, the SAN fabric must be zoned to ensure that only the SAN Volume Controller can access the target ports.
- **Access Logix enabled:** In this mode of operation, a storage group can be formed from a set of LUNs. Only the hosts that are assigned to the storage group are allowed to access these LUNs.

Configuring the EMC CLARiiON controller with Access Logix installed

The SAN Volume Controller does not have access to the storage controller logical units (LUs) if Access Logix is installed on the EMC CLARiiON controller. You must use the EMC CLARiiON configuration tools to associate the SAN Volume Controller and LU.

The following prerequisites must be met before you can configure an EMC CLARiiON controller with Access Logix installed:

- The EMC CLARiiON controller is not connected to the SAN Volume Controller
- You have a RAID controller with LUs and you have identified which LUs you want to present to the SAN Volume Controller

You must complete the following tasks to configure an EMC CLARiiON controller with Access Logix installed:

- Register the SAN Volume Controller ports with the EMC CLARiiON
- Configure storage groups

The association between the SAN Volume Controller and the LU is formed when you create a storage group that contains both the LU and the SAN Volume Controller.

Related tasks

“Configuring your storage groups” on page 241

Storage groups can only be configured if Access Logix is installed and enabled.

Registering the SAN Volume Controller ports with the EMC CLARiiON

You must register the SAN Volume Controller ports with an EMC CLARiiON controller if Access Logix is installed.

The following prerequisites must be met before you can register the SAN Volume Controller ports with an EMC CLARiiON controller that has Access Logix installed:

- The EMC CLARiiON controller is not connected to the SAN Volume Controller
- You have a RAID controller with LUs and you have identified which LUs you want to present to the SAN Volume Controller

Each initiator port (WWPN) must be registered against a host name and against a target port to which access is granted. If a host has multiple initiator ports, multiple table entries with the same host name are listed. If a host is allowed access using multiple target ports, multiple table entries are listed. For SAN Volume Controller hosts, all WWPN entries should carry the same host name.

The following table lists the associations:

Option	EMC CLARiiON default setting	SAN Volume Controller required setting
WWPN	N/A	Any
WWN	N/A	Any
Host name	N/A	Any
SP port	N/A	Any
Initiator type	3	3
ArrayCommPath	Enable	Disable
Failover mode	0	2
Unit Serial Number	Array	Any

1. Connect the fibre channel and zone the fabric as required.
2. Issue the **svctask detectmdisk** command-line interface (CLI) command.
3. Right-click on the storage subsystem from the Enterprise Storage window.
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 SAN Volume Controller ports to appear in the dialog box. Use the WWPN to identify them. This can take several minutes.
7. Click **Group Edit**.
8. Select all instances of all the SAN Volume Controller ports in the Available dialog box.
9. Click the right arrow to move them to the selected box.
10. Fill in the **HBA WWN** field. You must know the following information:
 - WWNN of each SAN Volume Controller 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 port. The following is an example of the output:

```
50:05:07:68:01:00:8B:D8:50:05:07:68:01:20:8B:D8
```

11. Select **A** in the field marked SP™ and **0** in the SP Port field.
12. Select **CLARiiON Open** in the drop down list of the **Initiator Type** field.
13. Deselect the ArrayCommPath checkbox if it has been selected.
14. Select **2** in the drop down list of the **Failover Mode** field.

Attention: Failure to select failover mode 2 prevents the SAN Volume Controller from being able to failover I/O. Your data might become unavailable in the event of a single failure.

 - a. If this is the first time that a port has been registered, ensure that you select the New Host option. Otherwise, select Existing Host.
 - b. Ensure that the same host name is entered for each port that is registered.
15. Assign a host name in the Host Name field.
16. Click **OK**.
17. Specify the IP address of your switch. The EMC CLARiiON does not use this IP address. However it must be unique (within the EMC CLARiiON) to prevent errant behavior by Navisphere.
18. Repeat step 11 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
19. Repeat steps 1 on page 240 to 18 to register the rest of your SAN Volume Controller WWPNs.

All your WWPNs are registered against the host name that you specified.

Related tasks

“Adding nodes to a cluster using the CLI” on page 153

You can use the command-line interface (CLI) to add nodes to a cluster.

Configuring your storage groups

Storage groups can only be configured if Access Logix is installed and enabled.

Access Logix provides the following LUN mapping:

Note:

- A subset of logical units (LUs) can form a storage group.
 - An LU can be in multiple storage groups.
 - A host can be added to a storage group. This host has access to all LUs in the storage group.
 - A host *cannot* be added to a second storage group.
1. Right-click on the storage subsystem from the Enterprise Storage window.
 2. Select **Create Storage Group**. The Create Storage Group window is displayed.
 3. Enter a name for your storage group in the **Storage Group Name** field.
 4. Select **Dedicated** in the **Sharing State** field.
 5. Click **OK**. The storage group is created.

6. Right-click the storage group in the Enterprise Storage window.
7. Select **Properties**. The Storage Group Properties window is displayed.
8. Perform the following steps from the Storage Group Properties window:
 - a. Select the **LUNs** tab.
 - b. Select the LUNs that you want the SAN Volume Controller to manage in the Available LUNs table.

Attention: Ensure that the LUs 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 that you created when you registered the SAN Volume Controller ports with the EMC CLARiiON.

Attention: Ensure that only SAN Volume Controller 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 without Access Logix installed

If Access Logix is not installed on an EMC CLARiiON controller, all logical units (LUs) that were created on the controller can be used by the SAN Volume Controller.

No further configuration of the EMC CLARiiON controller is necessary.

Configure the switch zoning such that no hosts can access these LUs.

Related reference

“Switch zoning for the SAN Volume Controller” on page 65
Ensure that you are familiar with the constraints for zoning a switch.

Supported models of the EMC CLARiiON

The SAN Volume Controller supports models of the EMC CLARiiON.

Table 17 lists the supported models of the EMC CLARiiON. See the following Web site for the latest supported models:

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

Table 17. Supported models of the EMC CLARiiON

Model
FC4700-1
FC4700-2
CX200

Table 17. Supported models of the EMC CLARiiON (continued)

Model
CX300
CX400
CX500
CX600
CX700

Supported firmware levels for the EMC CLARiiON

The SAN Volume Controller supports EMC CLARiiON.

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

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

Concurrent maintenance on the EMC CLARiiON

Concurrent maintenance is the ability to perform I/O operations to a controller while simultaneously performing maintenance on it.

Important: An EMC Field Engineer must perform all maintenance procedures.

The EMC CLARiiON FC series and the SAN Volume Controller allow concurrent replacement of the following components:

- Disk drives
- Controller fans (fans must be replaced within 2 minutes or controllers are shut down.)
- Disk enclosure fans (fans must be replaced within 2 minutes or controllers are shut down.)
- Controller (service processor: you must first disable cache)
- Fibre Channel Bypass cards (LCC)
- Power supplies (you must first remove fans.)
- UPS battery (SPS)

EMC CLARiiON FC devices require that the I/O is quiesced during code upgrade. Consequently, the SAN Volume Controller does not support concurrent upgrade of the FC controller code.

The EMC CLARiiON CX series and the SAN Volume Controller allow 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 are shut down.)
- UPS battery (SPS)

The SAN Volume Controller and EMC CLARiiON CX devices support concurrent code upgrade of the CX controllers.

Note:

- EMC CLARiiON procedures for concurrent upgrade must be followed in all cases.
- 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* a concurrent operation.

User interface on EMC CLARiiON

Ensure that you are familiar with the user interface that supports the EMC CLARiiON subsystem.

NaviSphere or Navicli

The following user interface applications are available with the EMC CLARiiON:

- The Web-based application NaviSphere can be accessed from any Web browser.
- The command-line interface (CLI) Navicli is installed as part of the NaviSphere Agent software (the host software).

Note: Some options and features are only accessible through the CLI.

Communication with the controller in both cases is out-of-band. Therefore, the host does not need to be connected to the storage over fibre-channel and cannot be connected without Access Logix.

Web Server

A Web server is running on each of the controllers on the subsystem. During normal operation, the user interface only allows basic monitoring of the subsystem and displays an error log. If you press the reset button on the controller to put the controller in diagnostic mode, the user interface allows firmware upgrades and subsystem configuration resets.

Sharing the EMC CLARiiON between a host and the SAN Volume Controller

The EMC CLARiiON can be shared between a host and a SAN Volume Controller.

- Split controller access is only supported when Access Logix is installed and enabled.
- A host cannot be connected to both the SAN Volume Controller and EMC CLARiiON at the same time.
- LUs must not be shared between a host and a SAN Volume Controller.
- Partitions in a RAID group must not be shared between a host and a SAN Volume Controller.

Related concepts

“Storage subsystems” on page 55

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

Switch zoning limitations for the EMC CLARiiON

There are limitations in switch zoning for the SAN Volume Controller and EMC CLARiiON.

The EMC CLARiiON must be configured to present logical units (LUs) to all SAN Volume Controller initiator ports that are in the fabric zone.

Only SAN Volume Controller initiator ports that are LUN masked on the EMC CLARiiON controller should be present in the fabric zone.

You must consider the number of connections (process logins) that are consumed by the SAN Volume Controller cluster and the EMC CLARiiON. Use the following calculation to determine the number of connections for a single fabric:

- Number of SAN Volume Controller nodes \times the number of initiator ports \times the number of target ports

If this exceeds the subsystem capabilities, reduce the number of initiator or target ports in the configuration without introducing a single point of failure.

- To reduce the number of *initiator* ports, use only two of the four ports on each SAN Volume Controller node (one per HBA) and configure two fabrics, or fabric zones, such that these are the only initiator ports that are visible to each target port.
- To reduce the number of *target* ports, use ports from more than one controller.

The EMC CLARiiON CX200 provides 2 ports and supports 30 connections. Using a single SAN fabric, a 4-node cluster requires 32 connections ($4 \times 4 \times 2$). This exceeds the CX200 capability and exposes the SAN Volume Controller cluster integrity. Because there are only two target ports available, you must reduce the number of initiator ports. This consumes only 16 of the available 30 connections.

Note: The EMC CLARiiON CX200 cannot be used in an 8-node cluster configuration because the number of initiator ports cannot be fewer than 16 (2 per node) and the number of target ports cannot be fewer than 2. This consumes 32 connections and still exceeds the subsystem limit.

EMC CLARiiON FC4700 and CX400 systems provide 4 target ports and support 64 connections. Using a single SAN fabric, a 4-node cluster requires 64 connections ($4 \times 4 \times 4$). Therefore, this equals the EMC CLARiiON capabilities and is only a problem if split support with other hosts is required. Reducing either the number of initiator ports or target ports consumes 32 of the available 64 connections.

EMC CLARiiON CX600 provides 8 target ports and supports 128 connections. A 4-node cluster consumes all 128 connections ($4 \times 4 \times 8$). An eight-node cluster exceeds the connection limit and none of the reduction schemes must be used.

Related reference

“Switch zoning for the SAN Volume Controller” on page 65
Ensure that you are familiar with the constraints for zoning a switch.
Switch zoning for the SAN Volume Controller
Ensure that you are familiar with the constraints for zoning a switch.

Quorum disks on the EMC CLARiiON

The EMC CLARiiON supports quorum disks.

A SAN Volume Controller configuration that only includes the EMC CLARiiON is permitted.

Related concepts

Chapter 7, “Configuring and servicing storage subsystems,” on page 215
You must ensure that your disk controllers and switches are correctly configured to work with the SAN Volume Controller to avoid performance issues.

Related information

“Creating a quorum disk” on page 236

A quorum disk is used to resolve tie-break situations when the voting set of nodes disagree on the current state of the cluster.

Advanced functions for the EMC CLARiiON

Some advanced functions of the EMC CLARiiON are not supported by the SAN Volume Controller.

Advanced copy functions

Advanced copy functions for EMC CLARiiON (for example, SnapView, MirrorView and SANcopy) are not supported for disks that are managed by the SAN Volume Controller because the copy function does *not* extend to the SAN Volume Controller cache.

MetaLUN

MetaLUN allows a logical unit (LU) to be expanded using LUs in other RAID groups. The SAN Volume Controller only supports MetaLUN for the migration of image mode virtual disks.

Logical unit creation and deletion on the EMC CLARiiON

Binding an LU to a RAID group can take a significant amount of time on the EMC CLARiiON.

The LU must not be added to a storage group until binding is complete. As a safeguard, the SAN Volume Controller will not discover the LU if binding is in progress. A subsequent manual discovery is required.

Related tasks

“Discovering MDisks using the CLI” on page 157

You can use the command-line interface (CLI) to discover managed disks (MDisks).

Configuring settings for the EMC CLARiiON

A number of settings and options are available through the EMC CLARiiON configuration interface.

The following settings and options are supported by the SAN Volume Controller:

- Subsystem
- Port
- Logical unit

Related reference

“Global settings for the EMC CLARiiON” on page 247

Global settings apply across an EMC CLARiiON subsystem.

“Controller settings for the EMC CLARiiON” on page 247

The controller settings for the EMC CLARiiON are the settings that apply across one EMC CLARiiON subsystem.

“Port settings for the EMC CLARiiON” on page 248

Port settings are configurable at the port level.

“Logical unit settings for the EMC CLARiiON” on page 248

Logical unit (LU) settings are configurable at the LU level.

Global settings for the EMC CLARiiON

Global settings apply across an EMC CLARiiON subsystem.

Table 18 lists the global settings that are supported by the SAN Volume Controller.

Table 18. EMC CLARiiON global settings supported by the SAN Volume Controller

Option	EMC CLARiiON default setting	SAN Volume Controller 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
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 concepts

“Storage subsystems” on page 23

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

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Controller settings for the EMC CLARiiON

The controller settings for the EMC CLARiiON are the settings that apply across one EMC CLARiiON subsystem.

Table 19 lists the options that can be set by the EMC CLARiiON.

Table 19. EMC CLARiiON controller settings supported by the SAN Volume Controller

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 cannot obtain or change the configuration options that are listed above. You must configure the options that are listed above.

Related concepts

“Storage subsystems” on page 23

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

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Port settings for the EMC CLARiiON

Port settings are configurable at the port level.

Table 20 lists port settings, the EMC CLARiiON defaults, and the required settings for the SAN Volume Controller.

Table 20. EMC CLARiiON port settings supported by the SAN Volume Controller

Option	EMC CLARiiON default setting	SAN Volume Controller required setting
Port speed	2 GB	Either 1 or 2 GB

Note: The SAN Volume Controller cannot obtain or change the configuration options that are listed above. You must configure the options that are listed above.

Related concepts

“Storage subsystems” on page 23

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

“VDisk-to-host mapping” on page 33

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

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Logical unit settings for the EMC CLARiiON

Logical unit (LU) settings are configurable at the LU level.

Table 21 lists the options that must be set for each LU that is accessed by the SAN Volume Controller. LUs that are accessed by hosts can be configured differently.

Table 21. EMC CLARiiON LU settings supported by the SAN Volume Controller

Option	EMC CLARiiON default setting	SAN Volume Controller 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 cannot obtain or change the configuration options that are listed above. You must configure the options that are listed above.

Related concepts

“Storage subsystems” on page 23

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

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Configuring the EMC Symmetrix subsystem

This section provides information about configuring the EMC Symmetrix for attachment to a SAN Volume Controller. Except where explicitly noted, all information in this section also applies to EMC DMX subsystems.

Related reference

“Supported models of the EMC Symmetrix controller”

The SAN Volume Controller supports models of the EMC Symmetrix controller.

“Supported firmware levels for the EMC Symmetrix controller” on page 251

The SAN Volume Controller supports EMC Symmetrix.

“Concurrent maintenance on the EMC Symmetrix” on page 251

Concurrent maintenance is the capability to perform I/O operations to an EMC Symmetrix while simultaneously performing maintenance operations on it.

“User interface on EMC Symmetrix” on page 252

Ensure that you are familiar with the user interface application that supports the EMC Symmetrix subsystem.

“Sharing the EMC Symmetrix controller between a host and a SAN Volume Controller” on page 252

The EMC Symmetrix can be shared between a host and a SAN Volume Controller.

“Switch zoning limitations for the EMC Symmetrix” on page 253

Consider the following limitations when you plan switch zoning and connect to the SAN.

“Quorum disks on EMC Symmetrix” on page 254

The SAN Volume Controller chooses managed disks (MDisks) that are presented by the EMC Symmetrix as quorum disks.

“Advanced functions for EMC Symmetrix” on page 254

Advanced copy functions for Symmetrix (for example, SRDF and TimeFinder) are *not* supported for disks that are managed by the SAN Volume Controller because the copy does not extend to the SAN Volume Controller cache.

“Logical unit creation and deletion on EMC Symmetrix” on page 254

A logical unit (LU) that is exported by an EMC Symmetrix, meaning it is visible to a host, is either a *Symmetrix device* or a *Meta device*.

“Configuring settings for the EMC Symmetrix” on page 255

A number of settings and options are available through the EMC Symmetrix configuration interface.

Supported models of the EMC Symmetrix controller

The SAN Volume Controller supports models of the EMC Symmetrix controller.

Table 22 lists the supported models of the EMC Symmetrix controller. See the following Web site for the latest supported models:

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

Table 22. Supported models of the EMC Symmetrix

Model
DMX 800
DMX 1000
DMX 2000
DMX 3000

Table 22. Supported models of the EMC Symmetrix (continued)

Model
Symmetrix 8000 (Symm 5)-8130
Symmetrix 8000 (Symm 5)-8230
Symmetrix 8000 (Symm 5)-8430
Symmetrix 8000 (Symm 5)-8530
Symmetrix 8000 (Symm 5)-8730
Symmetrix 8000 (Symm 5)-8830

Related reference

“Configuring the EMC Symmetrix subsystem” on page 250

This section provides information about configuring the EMC Symmetrix for attachment to a SAN Volume Controller. Except where explicitly noted, all information in this section also applies to EMC DMX subsystems.

Supported firmware levels for the EMC Symmetrix controller

The SAN Volume Controller supports EMC Symmetrix.

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

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

Related reference

“Configuring the EMC Symmetrix subsystem” on page 250

This section provides information about configuring the EMC Symmetrix for attachment to a SAN Volume Controller. Except where explicitly noted, all information in this section also applies to EMC DMX subsystems.

Concurrent maintenance on the EMC Symmetrix

Concurrent maintenance is the capability to perform I/O operations to an EMC Symmetrix while simultaneously performing maintenance operations on it.

Important: Service actions and upgrade procedures can only be performed by an EMC Field Engineer.

The EMC Symmetrix is an Enterprise class device that supports nondisruptive replacement of the following components:

- Channel Director
- Disk Director
- Cache card
- Disk drive
- Cooling fan
- Comms card
- EPO card
- Operator panel
- PSU
- Service Processor
- Batteries
- Ethernet hub

The SAN Volume Controller and EMC Symmetrix support concurrent upgrade of the EMC Symmetrix firmware.

Related reference

“Configuring the EMC Symmetrix subsystem” on page 250

This section provides information about configuring the EMC Symmetrix for attachment to a SAN Volume Controller. Except where explicitly noted, all information in this section also applies to EMC DMX subsystems.

User interface on EMC Symmetrix

Ensure that you are familiar with the user interface application that supports the EMC Symmetrix subsystem.

EMC Control Center

A basic EMC Symmetrix configuration is performed by an EMC Field Engineer (FE) using the EMC Symmetrix service processor. After the initial configuration, you can configure and control the exported storage. The FE defines the storage device types and sets the configurable options.

You can configure and control the exported storage as described below.

You can use the EMC Control Center to manage and monitor the EMC Symmetrix storage system.

You can use Volume Logix for volume configuration management. Volume Logix allows you to control access rights to the storage when multiple hosts share target ports.

SYMCLI

The EMC Symmetrix Command Line Interface (SYMCLI) allows the server to monitor and control EMC Symmetrix.

Related reference

“Configuring the EMC Symmetrix subsystem” on page 250

This section provides information about configuring the EMC Symmetrix for attachment to a SAN Volume Controller. Except where explicitly noted, all information in this section also applies to EMC DMX subsystems.

Sharing the EMC Symmetrix controller between a host and a SAN Volume Controller

The EMC Symmetrix can be shared between a host and a SAN Volume Controller.

The following restrictions apply when you share the EMC Symmetrix controller between a host and a SAN Volume Controller:

- Target ports must not be shared between the SAN Volume Controller and other hosts.
- A single host must not be connected to a SAN Volume Controller and an EMC Symmetrix because the multipathing drivers (for example, Subsystem Device Driver (SDD) and PowerPath) cannot coexist.
- Other hosts can be directly connected to an EMC Symmetrix at the same time as a SAN Volume Controller, under the following conditions:
 - The fabric must be zoned such that other hosts cannot access the target ports that are used by the SAN Volume Controller.

- The EMC Symmetrix must be configured such that other hosts cannot access the LUs that are managed by the SAN Volume Controller.

Related concepts

“Storage subsystems” on page 55

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

Related reference

“Configuring the EMC Symmetrix subsystem” on page 250

This section provides information about configuring the EMC Symmetrix for attachment to a SAN Volume Controller. Except where explicitly noted, all information in this section also applies to EMC DMX subsystems.

Switch zoning limitations for the EMC Symmetrix

Consider the following limitations when you plan switch zoning and connect to the SAN.

Switch zoning

The SAN Volume Controller switch zone must include at least one target port on two or more fibre-channel adapters (FAs) to avoid a single point of failure.

The EMC Symmetrix must be configured to present logical units (LUs) to all SAN Volume Controller initiator ports that are in the fabric zone.

Only SAN Volume Controller initiator ports that are LUN masked on the EMC Symmetrix controller should be present in the fabric zone.

Connecting to the SAN

The EMC Symmetrix connects to the SAN through a fibre-channel director. Directors are installed in pairs and each consists of two boards, one of which is a fibre-channel adapter (FA). The FA provides 2 - 12 target ports. The EMC Symmetrix assigns a worldwide node name (WWNN) per target port and the SAN Volume Controller can resolve up to four WWNN's per subsystem. If you want to connect more than four target ports to a SAN Volume Controller, the following procedure must be performed:

1. Divide the set of target ports into groups of 2 - 4.
2. Define a discrete set of logical units (LUs) for each group.
3. Map the LUs to each target port in their group.

The SAN Volume Controller views each group of target ports as a separate subsystem. Ensure that no LUs are a member of more than one group.

Related reference

“Configuring the EMC Symmetrix subsystem” on page 250

This section provides information about configuring the EMC Symmetrix for attachment to a SAN Volume Controller. Except where explicitly noted, all information in this section also applies to EMC DMX subsystems.

“Switch zoning for the SAN Volume Controller” on page 65

Ensure that you are familiar with the constraints for zoning a switch.

Quorum disks on EMC Symmetrix

The SAN Volume Controller chooses managed disks (MDisks) that are presented by the EMC Symmetrix as quorum disks.

The SAN Volume Controller uses a logical unit (LU) that is presented by an EMC Symmetrix as a quorum disk. The SAN Volume Controller provides a quorum disk even if the connection is through a single port.

Related concepts

Chapter 7, “Configuring and servicing storage subsystems,” on page 215
You must ensure that your disk controllers and switches are correctly configured to work with the SAN Volume Controller to avoid performance issues.

Related reference

“Configuring the EMC Symmetrix subsystem” on page 250
This section provides information about configuring the EMC Symmetrix for attachment to a SAN Volume Controller. Except where explicitly noted, all information in this section also applies to EMC DMX subsystems.

Advanced functions for EMC Symmetrix

Advanced copy functions for Symmetrix (for example, SRDF and TimeFinder) are *not* supported for disks that are managed by the SAN Volume Controller because the copy does not extend to the SAN Volume Controller cache.

Related reference

“Configuring the EMC Symmetrix subsystem” on page 250
This section provides information about configuring the EMC Symmetrix for attachment to a SAN Volume Controller. Except where explicitly noted, all information in this section also applies to EMC DMX subsystems.

Logical unit creation and deletion on EMC Symmetrix

A logical unit (LU) that is exported by an EMC Symmetrix, meaning it is visible to a host, is either a *Symmetrix device* or a *Meta device*.

Symmetrix device

Symmetrix device is an EMC term for an LU that is hosted by an EMC Symmetrix. These are all emulated devices and have exactly the same characteristics. The following are the characteristics of a Symmetrix device:

- N cylinders
- 15 tracks per cylinder
- 64 logical blocks per track
- 512 bytes per logical block

Symmetrix devices can be created using the **create dev** command from the EMC Symmetrix Command Line Interface (SYMCLI). The configuration of an LU can be changed using the **convert dev** command from the SYMCLI. Each physical storage device in an EMC Symmetrix is partitioned into 1 - 128 hyper-volumes (hypers). Each hyper can be up to 16 GB. A Symmetrix device maps to one or more hypers, depending on how it is configured. The following are examples of hyper configurations:

- Hypers can be mirrored (2-way, 3-way, 4-way)
- Hypers can be formed into RAID-S groups

Meta device

Meta device is an EMC term for a concatenated chain of EMC Symmetrix devices. This enables the EMC Symmetrix to provide LUs that are larger than a hyper. Up to 255 hypers can be concatenated to form a single meta device. Meta devices can be created using the **form meta** and **add dev** commands from the SYMCLI.

Related concepts

“Storage subsystems” on page 23

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

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Related reference

“Configuring the EMC Symmetrix subsystem” on page 250

This section provides information about configuring the EMC Symmetrix for attachment to a SAN Volume Controller. Except where explicitly noted, all information in this section also applies to EMC DMX subsystems.

Configuring settings for the EMC Symmetrix

A number of settings and options are available through the EMC Symmetrix configuration interface.

The settings and options can have a scope of the following:

- Subsystem
- Port
- Logical unit (LU)

Related concepts

“Storage subsystems” on page 23

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

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Related reference

“Configuring the EMC Symmetrix subsystem” on page 250

This section provides information about configuring the EMC Symmetrix for attachment to a SAN Volume Controller. Except where explicitly noted, all information in this section also applies to EMC DMX subsystems.

“Global settings for the EMC Symmetrix” on page 256

Global settings apply across an EMC Symmetrix subsystem.

“Port settings for the EMC Symmetrix” on page 256

Target port characteristics can be set using the **set port** command from the Symmetrix Command Line Interface (SYMCLI).

“Logical unit settings for the EMC Symmetrix” on page 257

Logical unit (LU) settings are configurable at the LU level.

“Mapping and virtualization settings for the EMC Symmetrix” on page 258
 Mapping a logical unit (LU) to a host is a function of the EMC Control Center.

Related information

“Servicing storage subsystems” on page 237

Storage subsystems that are supported for attachment to the SAN Volume Controller are designed with redundant components and access paths to allow concurrent maintenance. Hosts have continuous access to their data during component failure and replacement.

Global settings for the EMC Symmetrix

Global settings apply across an EMC Symmetrix subsystem.

EMC Symmetrix characteristics can be set using the **set Symmetrix** command from the Symmetrix Command Line Interface (SYMCLI). The characteristics can be viewed using the **symconfigure** command from the SYMCLI.

Table 23 lists the EMC Symmetrix global settings that are supported by the SAN Volume Controller.

Table 23. EMC Symmetrix global settings supported by the SAN Volume Controller

Option	EMC Symmetrix default setting	SAN Volume Controller required setting
max_hypers_per_disk	-	Any
dynamic_rdf	Disable	Disable
fba_multi_access_cache	Disable	N/A
Raid_s_support	Disable	Enable or Disable

Related concepts

“Storage subsystems” on page 23

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

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Port settings for the EMC Symmetrix

Target port characteristics can be set using the **set port** command from the Symmetrix Command Line Interface (SYMCLI).

The target port characteristics can be viewed using the **symcfg** command from the SYMCLI.

Table 24 lists the EMC Symmetrix port settings that are supported by the SAN Volume Controller.

Table 24. EMC Symmetrix port settings supported by the SAN Volume Controller

Option	EMC Symmetrix default setting	SAN Volume Controller required setting
Disk_Array	Enabled	Disabled
Volume_Set_Addresssing	Enabled	Disabled

Table 24. EMC Symmetrix port settings supported by the SAN Volume Controller (continued)

Option	EMC Symmetrix default setting	SAN Volume Controller required setting
Hard_Addressing	Enabled	Enabled
Non_Participating	Disabled	Disabled
Global_3rdParty_Logout	Enabled	Enabled
Tagged_Commands	Enabled	Enabled
Common_Serial_Number	-	Enabled
Disable_Q_Reset_on_UA	Disabled	Disabled
Return_busy_for_abort	Disabled	Disabled
SCSI-3	Disabled	Disabled
Environ_Set	Disabled	Disabled
Unique_WWN	Enabled	Enabled
Point_to_Point	Disabled	Enabled
VCM_State	Disabled	Either
OpenVMS	Disabled	Disabled

Related concepts

“Storage subsystems” on page 23

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

“VDisk-to-host mapping” on page 33

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

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Logical unit settings for the EMC Symmetrix

Logical unit (LU) settings are configurable at the LU level.

LU characteristics can be set using the **set device** command from the Symmetrix Command Line Interface (SYMCLI).

Table 25 lists the options that must be set for each LU that is accessed by the SAN Volume Controller.

Table 25. EMC Symmetrix LU settings supported by the SAN Volume Controller

Option	EMC Symmetrix default setting	SAN Volume Controller required setting
emulation	-	FBA
attribute	-	RAD

Related concepts

“Storage subsystems” on page 23

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

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Mapping and virtualization settings for the EMC Symmetrix

Mapping a logical unit (LU) to a host is a function of the EMC Control Center.

LUs can be mapped to a particular director or target port using the **map dev** command from the Symmetrix Command Line Interface (SYMCLI). LUs can be unmapped using the **unmap dev** command from the SYMCLI.

Related reference

“Switch zoning for the SAN Volume Controller” on page 65

Ensure that you are familiar with the constraints for zoning a switch.

Configuring the IBM TotalStorage Enterprise Storage Server

This section provides information about configuring the IBM TotalStorage Enterprise Storage Server (ESS) so it can attach to a SAN Volume Controller.

Related tasks

“Configuring the IBM ESS” on page 259

The IBM Enterprise Storage Server (ESS) provides functionality that is compatible with the SAN Volume Controller.

Related reference

“Supported models of the IBM ESS” on page 260

The SAN Volume Controller supports models of the IBM Enterprise Storage Server (ESS).

“Supported firmware levels for the IBM ESS” on page 260

The SAN Volume Controller supports the IBM Enterprise Storage Server (ESS).

“Concurrent maintenance on the IBM ESS” on page 261

Concurrent maintenance is the capability to perform I/O operations to an IBM Enterprise Storage Server (ESS) while simultaneously performing maintenance operations on it.

“User interface on the IBM ESS” on page 261

Ensure that you are familiar with the user interface application that supports the IBM Enterprise Storage Server (ESS) subsystem.

“Sharing the IBM ESS between a host and the SAN Volume Controller” on page 261

The IBM Enterprise Storage Server (ESS) can be shared between a host and a SAN Volume Controller.

“Switch zoning limitations for the IBM ESS” on page 262

Consider the following limitations when you zone the IBM Enterprise Storage Server (ESS) to the SAN Volume Controller.

“Quorum disks on the IBM ESS” on page 262

The SAN Volume Controller can choose managed disks (MDisks) that are presented by the IBM Enterprise Storage Server (ESS) controller as quorum disks.

“Advanced functions for the IBM ESS” on page 262

SAN Volume Controller cache-disabled virtual disk (VDisks) can be used as the source or target for IBM Enterprise Storage Server (ESS) advanced copy functions (for example, FlashCopy, MetroMirror, GlobalCopy).

“Logical unit creation and deletion on the IBM ESS” on page 263

Certain IBM Enterprise Storage Server (ESS) types are supported for use with the SAN Volume Controller.

Configuring the IBM ESS

The IBM Enterprise Storage Server (ESS) provides functionality that is compatible with the SAN Volume Controller.

Perform the following steps to configure the IBM ESS:

1. Enter the IP address of the IBM ESS in a Web browser to access the ESS Specialist.
2. Login with your 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 each initiator port on each SAN Volume Controller node in your cluster. Complete the following fields:
 - a. Enter a unique name for each port in the **Nickname** field. For example, enter knode or lnode.
 - b. Select **IBM SAN Volume Controller** in the **Host Type** field. If IBM SAN Volume Controller is not available, select **RS/6000**.
 - c. Select **Fibre Channel attached** in the **Host Attachment** field.
 - d. Leave the **Hostname/IP address** field blank.
 - e. Select the WWPN from the list or enter it manually into the **WWPN** field. A configuration command fails if you use WWPN 0 in the command string.
8. Click **Perform Configuration Update** after you are finished adding all of the ports.
9. Click **Add Volumes** to add the volumes that you want the SAN Volume Controller to use. The Add Volumes panel is displayed.
10. Perform the following steps in the Add Volumes panel:
 - a. Select any of the SAN Volume Controller 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. Click **Perform Configuration Update** after you have created all the volumes.
11. Perform the following steps to map the volumes to all of your SAN Volume Controller ports:
 - 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 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 concepts

Chapter 7, “Configuring and servicing storage subsystems,” on page 215
 You must ensure that your disk controllers and switches are correctly configured to work with the SAN Volume Controller to avoid performance issues.

Related tasks

“Configuring a balanced storage subsystem” on page 222
 The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.
 “Expanding a logical unit using the CLI” on page 225
 You can use the command-line interface (CLI) to expand a logical unit.

Related information

“Configuring the IBM TotalStorage Enterprise Storage Server” on page 258
 This section provides information about configuring the IBM TotalStorage Enterprise Storage Server (ESS) so it can attach to a SAN Volume Controller.

Supported models of the IBM ESS

The SAN Volume Controller supports models of the IBM Enterprise Storage Server (ESS).

Table 26 lists the supported models of the IBM ESS. See the following Web site for the latest supported models:

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

Table 26. Supported models of the IBM ESS

Model
2105-E10
2105-E20
2105-F10
2105-F20
2105-750
2105-800
2107-900
1750-500
Note: Support for these models is dependent on the product availability date.

Related information

“Configuring the IBM TotalStorage Enterprise Storage Server” on page 258
 This section provides information about configuring the IBM TotalStorage Enterprise Storage Server (ESS) so it can attach to a SAN Volume Controller.

Supported firmware levels for the IBM ESS

The SAN Volume Controller supports the IBM Enterprise Storage Server (ESS).

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

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

Related information

“Configuring the IBM TotalStorage Enterprise Storage Server” on page 258
This section provides information about configuring the IBM TotalStorage Enterprise Storage Server (ESS) so it can attach to a SAN Volume Controller.

Concurrent maintenance on the IBM ESS

Concurrent maintenance is the capability to perform I/O operations to an IBM Enterprise Storage Server (ESS) while simultaneously performing maintenance operations on it.

All IBM ESS concurrent maintenance procedures are supported.

Related information

“Configuring the IBM TotalStorage Enterprise Storage Server” on page 258
This section provides information about configuring the IBM TotalStorage Enterprise Storage Server (ESS) so it can attach to a SAN Volume Controller.

User interface on the IBM ESS

Ensure that you are familiar with the user interface application that supports the IBM Enterprise Storage Server (ESS) subsystem.

Web Server

A Web server runs on each of the controllers on the subsystem. During normal operation, the user interface application allows only basic monitoring of the subsystem and displays an error log. If you press the reset button on the controller to put the controller into diagnostic mode, the user interface application allows firmware upgrades and subsystem configuration resets.

Related information

“Configuring the IBM TotalStorage Enterprise Storage Server” on page 258
This section provides information about configuring the IBM TotalStorage Enterprise Storage Server (ESS) so it can attach to a SAN Volume Controller.

Target port groups for the IBM ESS

The IBM Enterprise Storage Server (ESS) model 1750-500 uses the SCSI Target Port Groups feature to indicate a preferred path for each logical unit (LU).

Sharing the IBM ESS between a host and the SAN Volume Controller

The IBM Enterprise Storage Server (ESS) can be shared between a host and a SAN Volume Controller.

The following restrictions apply when you share the IBM ESS between a host and a SAN Volume Controller:

- If an IBM ESS port is in the same zone as a SAN Volume Controller port, that same IBM ESS port should not be in the same zone as another host.
- A single host can have both IBM ESS direct-attached and SAN Volume Controller virtualized disks configured to it.
- If a LUN is managed by the SAN Volume Controller, it *cannot* be mapped to another host.

See the following Web site for the latest supported configurations:

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

Related concepts

“Storage subsystems” on page 55

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

Related information

“Configuring the IBM TotalStorage Enterprise Storage Server” on page 258

This section provides information about configuring the IBM TotalStorage Enterprise Storage Server (ESS) so it can attach to a SAN Volume Controller.

Switch zoning limitations for the IBM ESS

Consider the following limitations when you zone the IBM Enterprise Storage Server (ESS) to the SAN Volume Controller.

To avoid a single point of failure on the IBM ESS, you must have a minimum of two SAN connections from two separate adapter bays. The maximum number of IBM ESS SAN connections in the SAN Volume Controller switch zone is 16.

Note: The IBM ESS provides ESCON[®], FICON[®] and Ultra SCSI connectivity; however, only a 1 or 2 Gb fibre-channel SAN attachment is supported by the SAN Volume Controller.

Related reference

“Switch zoning for the SAN Volume Controller” on page 65

Ensure that you are familiar with the constraints for zoning a switch.

Related information

“Configuring the IBM TotalStorage Enterprise Storage Server” on page 258

This section provides information about configuring the IBM TotalStorage Enterprise Storage Server (ESS) so it can attach to a SAN Volume Controller.

Quorum disks on the IBM ESS

The SAN Volume Controller can choose managed disks (MDisks) that are presented by the IBM Enterprise Storage Server (ESS) controller as quorum disks.

Related concepts

Chapter 7, “Configuring and servicing storage subsystems,” on page 215

You must ensure that your disk controllers and switches are correctly configured to work with the SAN Volume Controller to avoid performance issues.

Related information

“Configuring the IBM TotalStorage Enterprise Storage Server” on page 258

This section provides information about configuring the IBM TotalStorage Enterprise Storage Server (ESS) so it can attach to a SAN Volume Controller.

“Creating a quorum disk” on page 236

A quorum disk is used to resolve tie-break situations when the voting set of nodes disagree on the current state of the cluster.

Advanced functions for the IBM ESS

SAN Volume Controller cache-disabled virtual disk (VDisks) can be used as the source or target for IBM Enterprise Storage Server (ESS) advanced copy functions (for example, FlashCopy, MetroMirror, GlobalCopy).

Related information

“Configuring the IBM TotalStorage Enterprise Storage Server” on page 258

This section provides information about configuring the IBM TotalStorage Enterprise Storage Server (ESS) so it can attach to a SAN Volume Controller.

Logical unit creation and deletion on the IBM ESS

Certain IBM Enterprise Storage Server (ESS) types are supported for use with the SAN Volume Controller.

Before you delete or unmap a logical unit (LU) from the SAN Volume Controller, remove the LU from the managed disk (MDisk) group. The following is supported:

- LU size of 1 GB to 2 TB.
- RAID 5 and RAID 10 LUs.
- LUs can be added dynamically.

Attention: When adding a new LU, you *must* select the **Use same ID/LUN in source and target** checkbox. Failure to select the **Use same ID/LUN in source and target** checkbox can cause loss in redundancy or a loss of data integrity. The detect MDisks task in the SAN Volume Controller Console or the **lsinfo detectmdisks** command-line interface (CLI) command must be run for the SAN Volume Controller to detect the new disks.

Related concepts

“Storage subsystems” on page 23

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

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Related information

“Configuring the IBM TotalStorage Enterprise Storage Server” on page 258

This section provides information about configuring the IBM TotalStorage Enterprise Storage Server (ESS) so it can attach to a SAN Volume Controller.

Configuring the IBM TotalStorage DS4000 (formerly FAStT) series subsystem

This section provides information about configuring the IBM TotalStorage DS4000 series subsystem so it can attach to a SAN Volume Controller. Certain models of the IBM TotalStorage DS4000 series of controllers are equivalent to StorageTek models; therefore, the SAN Volume Controller also supports models of the StorageTek FlexLine series and StorageTek D series.

The information in this section also applies to the supported models of the StorageTek FlexLine series and StorageTek D series.

Related tasks

“Configuring IBM DS4000 series disk controllers for the storage server” on page 264

The IBM DS4000 series of disk controllers provide functionality that is compatible with the SAN Volume Controller.

Related reference

“Supported options of the IBM DS4000 series controller” on page 266

The IBM DS4000 series disk controllers provide functionality that can be used with the SAN Volume Controller.

“Supported models of the IBM DS4000 series of controllers” on page 267
The SAN Volume Controller supports models of the IBM DS4000 series. Certain models of the IBM DS4000 series of controllers are equivalent to StorageTek models; therefore, the SAN Volume Controller also supports models of StorageTek’s FlexLine series and D series.

“Supported firmware levels for the IBM DS4000 series” on page 268
The SAN Volume Controller supports the IBM DS4000 series.

“Concurrent maintenance on the IBM DS4000 series” on page 269
Concurrent maintenance is the capability to perform I/O operations to an IBM DS4000 series controller while simultaneously performing maintenance operations on it. Refer to your IBM DS4000 series documentation for information about concurrent maintenance.

“User interface on the IBM DS4000 series” on page 269
Ensure that you are familiar with the user interface that supports the IBM DS4000 series subsystem.

“Sharing the IBM DS4000 series controller between a host and the SAN Volume Controller” on page 269
The IBM DS4000 series controller can be shared between a host and a SAN Volume Controller.

“Quorum disks on the IBM DS4000 series” on page 270
The SAN Volume Controller can choose managed disks (MDisks) that are presented by the IBM DS4000 series controller as quorum disks.

“Advanced functions for the IBM DS4000 series” on page 270
SAN Volume Controller cache-disabled virtual disks (VDisks) can be used as the source or target for IBM DS4000 advanced copy functions (for example, FlashCopy and MetroMirror).

“Logical unit creation and deletion on the IBM DS4000 series” on page 272
You can create or delete logical units on the IBM DS4000 series.

“Configuration interface for the IBM DS4000 series” on page 272
The IBM DS4000 series provides a configuration application.

“Controller settings for the IBM DS4000 series” on page 273
Controller settings are the settings that apply across one IBM DS4000 series controller.

Configuring IBM DS4000 series disk controllers for the storage server

The IBM DS4000 series of disk controllers provide functionality that is compatible with the SAN Volume Controller.

Attention: The SAN Volume Controller 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 use storage that is provided by the IBM DS4000 series of controllers that you want to update before you install new ESM firmware.

The IBM DS4000 storage servers have several options. The following steps provide the supported options and impact on the SAN Volume Controller:

1. Perform the following steps for the host type option:
 - a. You must set either the default host type of your IBM DS4000 series or the host type of the chosen partition to the following:
IBM TS SAN VCE
 - 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. Perform the following steps for the worldwide node name (WWNN) option:
 - a. Set the subsystem so that both controllers have the same WWNN.
 - b. See the following Web site for the scripts that are available to change the setup of the IBM DS4000:
<http://www.ibm.com/storage/support/>
3. Perform the following steps for the auto volume transfer (AVT) option:
 - a. Ensure that the AVT option is enabled. The host type selection should have already enabled the AVT option.
 - b. View the storage subsystem profile data to confirm that you have the AVT option enabled. This storage profile is presented as a text view in a separate window.
 - c. See the following Web site for the scripts that are available to enable the AVT option:
<http://www.ibm.com/storage/support/>

The following limitations apply to partitions:

- Only one IBM DS4000 series storage partition that contains any of the ports of any of the nodes in a single SAN Volume Controller cluster can be created.
- Only map one partition to any of the ports on any of the nodes that are in the SAN Volume Controller cluster to avoid unexpected behavior. For example, you can lose access to your storage or you might not receive warning messages, even if there are errors logged in the SAN Volume Controller error log.

The following limitation applies to the IBM DS4000 series Copy Services:

- The IBM DS4000 series Copy Services must not be used when the SAN Volume Controller is attached to the IBM DS4000 series.
- You can use partitioning to allow IBM DS4000 series Copy Services to be used on other hosts.

The following information applies to the access LUN (also known as the Universal Transport Mechanism (UTM) LUN):

- The access/UTM LUN is a special LUN that allows the SAN Volume Controller to be configured through software over the fibre-channel connection. The access/UTM LUN does not have to be in the partition that contains the SAN Volume Controller ports because the access/UTM LUN is not required by the SAN Volume Controller. No errors are generated if the access/UTM LUN is not in the partition.

The following information applies to the logical unit (LU):

- The SAN Volume Controller attempts to follow the preferred ownership that is specified by the IBM DS4000 series. You can specify which controller (A or B) is used for I/O operations to an LU.
- If the SAN Volume Controller can see the ports of the preferred controller and error conditions do not exist, the SAN Volume Controller accesses the LU through one of the ports on the preferred controller.
- If error conditions exist, the SAN Volume Controller ignores the preferred ownership of the IBM DS4000 series.

Related concepts

Chapter 7, “Configuring and servicing storage subsystems,” on page 215
You must ensure that your disk controllers and switches are correctly configured to work with the SAN Volume Controller to avoid performance issues.

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

“Configuring the IBM ESS” on page 259

The IBM Enterprise Storage Server (ESS) provides functionality that is compatible with the SAN Volume Controller.

“Expanding a logical unit using the CLI” on page 225

You can use the command-line interface (CLI) to expand a logical unit.

Related information

“Configuring the IBM TotalStorage DS4000 (formerly FAStT) series subsystem” on page 263

This section provides information about configuring the IBM TotalStorage DS4000 series subsystem so it can attach to a SAN Volume Controller. Certain models of the IBM TotalStorage DS4000 series of controllers are equivalent to StorageTek models; therefore, the SAN Volume Controller also supports models of the StorageTek FlexLine series and StorageTek D series.

Supported options of the IBM DS4000 series controller

The IBM DS4000 series disk controllers provide functionality that can be used with the SAN Volume Controller.

The IBM DS4000 series storage manager has several options and actions that you can perform.

Controller run diagnostics

The diagnostics should be automatically recovered by the SAN Volume Controller software. After the controller run diagnostics option is used, check your managed disks (MDisks) to ensure that they have not been set to degraded mode.

Controller disable data transfer

The controller disable data transfer option is not supported when a SAN Volume Controller is attached to the IBM DS4000 series. Loss of availability and redundancy can occur if data transfer is disabled.

Setting an array Offline

Do not set an array offline because you can lose access to the MDisk group.

Array increase capacity

The array increase capacity option is supported but the new capacity is not usable until the MDisk is removed from the MDisk group and re-added to the MDisk group. You might have to migrate data to increase the capacity.

Redistribute logical drives or change ownership of the preferred path

These options are supported; however, they might not take effect until a rediscovery is started on the SAN Volume Controller cluster. You can use the **svctask detectmdisk** command-line interface (CLI) command to start a cluster rediscovery. The discovery process rescans the fibre-channel network to discover any new MDisks that might have been added to the cluster and to rebalance MDisk access across the available controller device ports.

Controller reset

You should only use the controller reset option if you are directed to do so by IBM Service and the alternate controller is functional and available to the SAN. The SAN Volume Controller reset should be automatically recovered by the SAN Volume Controller software.

Check your MDisks to ensure that they have not been set to degraded state during the controller reset process. You can issue the **svctask includemdisk** CLI command to repair degraded MDisks.

Related concepts

Chapter 7, “Configuring and servicing storage subsystems,” on page 215
You must ensure that your disk controllers and switches are correctly configured to work with the SAN Volume Controller to avoid performance issues.

Related tasks

“Configuring a balanced storage subsystem” on page 222
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

“Configuring the IBM ESS” on page 259
The IBM Enterprise Storage Server (ESS) provides functionality that is compatible with the SAN Volume Controller.

Related information

“Configuring the IBM TotalStorage DS4000 (formerly FASTT) series subsystem” on page 263
This section provides information about configuring the IBM TotalStorage DS4000 series subsystem so it can attach to a SAN Volume Controller. Certain models of the IBM TotalStorage DS4000 series of controllers are equivalent to StorageTek models; therefore, the SAN Volume Controller also supports models of the StorageTek FlexLine series and StorageTek D series.

Supported models of the IBM DS4000 series of controllers

The SAN Volume Controller supports models of the IBM DS4000 series. Certain models of the IBM DS4000 series of controllers are equivalent to StorageTek models; therefore, the SAN Volume Controller also supports models of StorageTek’s FlexLine series and D series.

Table 27 on page 268 lists the supported models of the IBM DS4000 series, StorageTek Flexline series and StorageTek D series of controllers. See the following Web site for the latest supported models:

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

Table 27. Supported models of the IBM DS4000 series, StorageTek FlexLine series and StorageTek D series of controllers

IBM		StorageTek	
DS4000 series	FAStT series	D-series	FlexLine series
DS4100	100	-	FLX210
DS4200	200	D173	-
DS4400	500	-	-
DS4300	600/600-Turbo	D240	FLX240
DS4600	700	D178	-
DS4500	900	D280	FLX280
DS4800	-	-	FLX380

Note: Some levels of FAStT microcode support a maximum of 32 LUNs per host partition, newer versions allow up to 256 LUNs per host partition.

The following expansion units are also supported by the SAN Volume Controller:

- EXP500
- DS4000 EXP100 R2
- DS4000 EXP700
- EXP710

Related information

“Configuring the IBM TotalStorage DS4000 (formerly FAStT) series subsystem” on page 263

This section provides information about configuring the IBM TotalStorage DS4000 series subsystem so it can attach to a SAN Volume Controller. Certain models of the IBM TotalStorage DS4000 series of controllers are equivalent to StorageTek models; therefore, the SAN Volume Controller also supports models of the StorageTek FlexLine series and StorageTek D series.

Supported firmware levels for the IBM DS4000 series

The SAN Volume Controller supports the IBM DS4000 series.

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

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

The Web site includes the maximum number of LUNs per partition that are supported by the firmware level.

Related information

“Configuring the IBM TotalStorage DS4000 (formerly FAStT) series subsystem” on page 263

This section provides information about configuring the IBM TotalStorage DS4000 series subsystem so it can attach to a SAN Volume Controller. Certain models of the IBM TotalStorage DS4000 series of controllers are equivalent to StorageTek models; therefore, the SAN Volume Controller also supports models of the StorageTek FlexLine series and StorageTek D series.

Concurrent maintenance on the IBM DS4000 series

Concurrent maintenance is the capability to perform I/O operations to an IBM DS4000 series controller while simultaneously performing maintenance operations on it. Refer to your IBM DS4000 series documentation for information about concurrent maintenance.

Related information

“Configuring the IBM TotalStorage DS4000 (formerly FAStT) series subsystem” on page 263

This section provides information about configuring the IBM TotalStorage DS4000 series subsystem so it can attach to a SAN Volume Controller. Certain models of the IBM TotalStorage DS4000 series of controllers are equivalent to StorageTek models; therefore, the SAN Volume Controller also supports models of the StorageTek FlexLine series and StorageTek D series.

User interface on the IBM DS4000 series

Ensure that you are familiar with the user interface that supports the IBM DS4000 series subsystem.

Web Server

A Web server is running on each of the controllers in the subsystem. During normal operation, the user interface only allows basic monitoring of the subsystem and displays an error log. If you press the reset button to put a controller in diagnostic mode, the user interface allows firmware upgrades and subsystem configuration resets.

Related information

“Configuring the IBM TotalStorage DS4000 (formerly FAStT) series subsystem” on page 263

This section provides information about configuring the IBM TotalStorage DS4000 series subsystem so it can attach to a SAN Volume Controller. Certain models of the IBM TotalStorage DS4000 series of controllers are equivalent to StorageTek models; therefore, the SAN Volume Controller also supports models of the StorageTek FlexLine series and StorageTek D series.

Sharing the IBM DS4000 series controller between a host and the SAN Volume Controller

The IBM DS4000 series controller can be shared between a host and a SAN Volume Controller.

Attention: The IBM DS4000 series partitioning function does not have the same meaning as used by IBM.

The IBM DS4000 series 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 accessed logical units.

Note: The SAN Volume Controller partition must either contain all the ports of the SAN Volume Controller cluster that are connected to the SAN, or are zoned to have access to the IBM DS4000 series ports. At least one port from each IBM DS4000 series controller must be visible by the SAN Volume Controller cluster.

Related concepts

“Storage subsystems” on page 55

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

Related information

“Configuring the IBM TotalStorage DS4000 (formerly FAStT) series subsystem” on page 263

This section provides information about configuring the IBM TotalStorage DS4000 series subsystem so it can attach to a SAN Volume Controller. Certain models of the IBM TotalStorage DS4000 series of controllers are equivalent to StorageTek models; therefore, the SAN Volume Controller also supports models of the StorageTek FlexLine series and StorageTek D series.

Quorum disks on the IBM DS4000 series

The SAN Volume Controller can choose managed disks (MDisks) that are presented by the IBM DS4000 series controller as quorum disks.

Note: The FASsT series 200 does not support quorum disks.

Related concepts

Chapter 7, “Configuring and servicing storage subsystems,” on page 215

You must ensure that your disk controllers and switches are correctly configured to work with the SAN Volume Controller to avoid performance issues.

Related information

“Configuring the IBM TotalStorage DS4000 (formerly FAStT) series subsystem” on page 263

This section provides information about configuring the IBM TotalStorage DS4000 series subsystem so it can attach to a SAN Volume Controller. Certain models of the IBM TotalStorage DS4000 series of controllers are equivalent to StorageTek models; therefore, the SAN Volume Controller also supports models of the StorageTek FlexLine series and StorageTek D series.

“Creating a quorum disk” on page 236

A quorum disk is used to resolve tie-break situations when the voting set of nodes disagree on the current state of the cluster.

Advanced functions for the IBM DS4000 series

SAN Volume Controller cache-disabled virtual disks (VDisks) can be used as the source or target for IBM DS4000 advanced copy functions (for example, FlashCopy and MetroMirror).

Related concepts

“Metro Mirror” on page 47

Metro Mirror enables you to set up a relationship between two virtual disks (VDisks), so that updates that are made by an application to one VDisk are mirrored on the other VDisk.

Related reference

“Data migration on an existing IBM DS4000 series installation that contains partitions” on page 271

You can migrate data on an existing IBM DS4000 series installation that contains partitions.

Related information

“Configuring the IBM TotalStorage DS4000 (formerly FAStT) series subsystem” on page 263

This section provides information about configuring the IBM TotalStorage DS4000 series subsystem so it can attach to a SAN Volume Controller. Certain

models of the IBM TotalStorage DS4000 series of controllers are equivalent to StorageTek models; therefore, the SAN Volume Controller also supports models of the StorageTek FlexLine series and StorageTek D series.

Data migration on an existing IBM DS4000 series installation that contains partitions

You can migrate data on an existing IBM DS4000 series installation that contains partitions.

You can enable the SAN Volume Controller 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 IBM DS4000 series partition can contain up to 32 LUNs. Each partition can only access a unique set of HBA ports, as defined by the worldwide port names (WWPNs). For a single host to access multiple partitions, unique host fibre ports (WWPNs) must be assigned to each partition. All LUNs within a partition are surfaced to the assigned host fibre ports (no sub-partition LUN mapping).

Host A is mapped to LUN 0, 1, 2 in Partition 0
Host B is mapped to LUN 0, 1, 2, 3, 4, 5 in Partition 1
Host C is mapped to LUN 0, 1, 2 in Partition 2

To allow Host A to access the LUNs in partition B, you must remove one of the HBAs (for example, A1) from the access list for partition 0 and add it to partition 1. A1 cannot be on the access list for more than one partition.

To add a SAN Volume Controller into this configuration without backup and restore cycles requires a set of unique SAN Volume Controller HBA port WWPNs for each partition. This allows the IBM DS4000 series to surface the LUNs to the SAN Volume Controller, which then configures these LUNs as image-mode LUNs and surfaces them to the required hosts. This violates a requirement that all SAN Volume Controller nodes must be able to see all back-end storage. To fix this problem, change the IBM DS4000 series to allow more than 32 LUNs in 1 storage partition, so that you can move all the LUNs from all the other partitions into 1 partition and map to the SAN Volume Controller cluster.

Scenario: the SAN Volume Controller nodes cannot see all back-end storage

The IBM DS4000 series has 8 partitions with 30 LUNs in each.

Perform the following steps to allow the SAN Volume Controller nodes to see all back-end storage:

1. Change the mappings for the first 4 partitions on the IBM DS4000 series such that each partition is mapped to 1 port on each node. This maintains redundancy across the cluster.
2. Create a new partition on the IBM DS4000 series that is mapped to all 4 ports on all the nodes.
3. Gradually migrate the data into the managed disks (MDisks) in the target partition. As storage is freed from the source partitions, it can be reused as new storage in the target partition. As partitions are deleted, new partitions that must be migrated can be mapped and migrated in the same way. The host side data access and integrity is maintained throughout this process.

Logical unit creation and deletion on the IBM DS4000 series

You can create or delete logical units on the IBM DS4000 series.

Certain IBM DS4000 series controller types are supported for use with the SAN Volume Controller.

To create a logical disk, you must set either the default host type of your IBM DS4000 series or the host type of the chosen partition to the following:

IBM TS SAN VCE

Perform one of the following tasks to set the host type:

- Click **Storage Subsystem** → **Change** → **Default Host Type**
- For each host port, specify the host type of that port or modify existing ports.

Related concepts

“Storage subsystems” on page 23

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

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Related information

“Configuring the IBM TotalStorage DS4000 (formerly FASTT) series subsystem” on page 263

This section provides information about configuring the IBM TotalStorage DS4000 series subsystem so it can attach to a SAN Volume Controller. Certain models of the IBM TotalStorage DS4000 series of controllers are equivalent to StorageTek models; therefore, the SAN Volume Controller also supports models of the StorageTek FlexLine series and StorageTek D series.

Configuration interface for the IBM DS4000 series

The IBM DS4000 series provides a configuration application.

The access LUN, also known as the Universal Transport Mechanism (UTM) LUN, is the configuration interface for the IBM DS4000 series controller.

The access LUN might not be in a partition that contains the SAN Volume Controller ports because it is not required by the SAN Volume Controller. The UTM LUN is a special LUN that allows the IBM DS4000 series to be configured through suitable software over the fibre-channel connection. Because the SAN Volume Controller does not require the UTM LUN, it does not generate errors either way. The IBM DS4000 series *must not* have the Access UTM LUN that is presented as LUN 0 (zero).

It is possible to use in-band (over fibre channel) and out-of-band (over Ethernet) to allow the IBM DS4000 series configuration software to communicate with more than one IBM DS4000 series. If using in-band configuration, the Access UTM LUN must be configured in a partition that does not include any logical units that are accessed by the SAN Volume Controller cluster.

Note: In-band is not supported for access to the LUN while in the SAN Volume Controller partition.

Related information

“Configuring the IBM TotalStorage DS4000 (formerly FASTT) series subsystem” on page 263

This section provides information about configuring the IBM TotalStorage DS4000 series subsystem so it can attach to a SAN Volume Controller. Certain models of the IBM TotalStorage DS4000 series of controllers are equivalent to StorageTek models; therefore, the SAN Volume Controller also supports models of the StorageTek FlexLine series and StorageTek D series.

Controller settings for the IBM DS4000 series

Controller settings are the settings that apply across one IBM DS4000 series controller.

You must configure the following settings for the IBM DS4000 series:

- You must set either the default host type of your IBM DS4000 series or the host type of the chosen partition to the following:

IBM TS SAN VCE

Perform one of the following tasks to set the host type:

- Click **Storage Subsystem** → **Change** → **Default Host Type**
- For each host port, specify the host type of that port or modify existing ports.
- Set the subsystem so that both controllers have the same worldwide node name (WWNN). See the following Web site for the scripts that are available to change the setup of the IBM DS4000:
<http://www.ibm.com/storage/support/>
- Ensure that the AVT option is enabled. The host type selection should have already enabled the AVT option. View the storage subsystem profile data to confirm that you have the AVT option enabled. This storage profile is presented as a text view in a separate window. See the following Web site for the scripts that are available to enable the AVT option:
<http://www.ibm.com/storage/support/>
- You must have the following enabled on any logical units that are mapped to the IBM DS4000 series:
 - read caching
 - write caching
 - write cache mirroring
- You must *not* have caching without batteries enabled.

Related concepts

“Storage subsystems” on page 23

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

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Related reference

“Configuring settings for the IBM DS4000 series” on page 274

The IBM DS4000 series controller configuration interface provides configuration settings and options that are supported with the SAN Volume Controller.

“Global settings for the IBM DS4000 series” on page 275

Global settings apply across an IBM DS4000 series controller.

“Logical unit settings for the IBM DS4000 series” on page 275

Logical unit (LU) settings are configurable at the LU level.

“Miscellaneous settings for the IBM DS4000 series” on page 276

Refer to your IBM DS4000 series documentation for information about other settings.

Related information

“Configuring the IBM TotalStorage DS4000 (formerly FAStT) series subsystem” on page 263

This section provides information about configuring the IBM TotalStorage DS4000 series subsystem so it can attach to a SAN Volume Controller. Certain models of the IBM TotalStorage DS4000 series of controllers are equivalent to StorageTek models; therefore, the SAN Volume Controller also supports models of the StorageTek FlexLine series and StorageTek D series.

Configuring settings for the IBM DS4000 series

The IBM DS4000 series controller configuration interface provides configuration settings and options that are supported with the SAN Volume Controller.

These settings and options can have a scope of the following:

- Subsystem
- Logical unit (LU)
 - The SAN Volume Controller attempts to follow the IBM DS4000 series specified preferred ownership. You can specify which controller (A or B) is used to perform I/O operations to a given LU. If the SAN Volume Controller can see the ports of the preferred controller and no error conditions exist, the SAN Volume Controller accesses that LU through one of the ports on that controller. Under error conditions, the ownership is ignored.
 - Ensure that you have the following enabled on any LUs that are mapped to the SAN Volume Controller:
 - read caching
 - write caching
 - write cache mirroring
 - Ensure that caching without batteries is *not* enabled.

Related concepts

“Storage subsystems” on page 23

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

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

“Configuring IBM DS4000 series disk controllers for the storage server” on page 264

The IBM DS4000 series of disk controllers provide functionality that is compatible with the SAN Volume Controller.

Related information

“Servicing storage subsystems” on page 237

Storage subsystems that are supported for attachment to the SAN Volume

Controller are designed with redundant components and access paths to allow concurrent maintenance. Hosts have continuous access to their data during component failure and replacement.

Global settings for the IBM DS4000 series

Global settings apply across an IBM DS4000 series controller.

Table 28 lists the IBM DS4000 series global settings that are supported by the SAN Volume Controller.

Table 28. IBM DS4000 series controller global settings supported by the SAN Volume Controller

Option	DS4000 series default setting
Start flushing	80%
Stop flushing	80%
Cache block size	4 Kb

These settings can be adjusted depending on the performance requirements. Do not modify these settings unless you are directed by IBM Support.

A host type of IBM TS SAN VCE must be used to establish the correct global settings for the SAN Volume Controller. Either set this as the system default host type or, if partitioning is enabled, associate each SAN Volume Controller port with this host type.

Related concepts

“Storage subsystems” on page 23

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

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Logical unit settings for the IBM DS4000 series

Logical unit (LU) settings are configurable at the LU level.

LUs that are accessed by hosts can be configured differently.

The read ahead cache multiplier is typically set to 0 or 1. Do not modify this setting unless you are directed to do so by IBM Support.

The following must be enabled on any LUs that are mapped to the SAN Volume Controller:

- read caching
- write caching
- write cache mirroring

You must not have caching without batteries enabled.

When you create a new LU, set the host type for that LU to the following:

IBM TS SAN VCE

Note: IBM TS SAN VCE is set as the default if the default type was already displayed.

Related concepts

“Storage subsystems” on page 23

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

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Miscellaneous settings for the IBM DS4000 series

Refer to your IBM DS4000 series documentation for information about other settings.

Related concepts

“Storage subsystems” on page 23

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

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Configuring the HDS Lightning series subsystem

This section provides information about configuring the Hitachi Data Systems (HDS) Lightning series subsystem so it can attach to a SAN Volume Controller. Certain models of the HDS Lightning series are equivalent to HP and Sun models; therefore, the SAN Volume Controller also supports models of the Sun StorEdge series and the HP XP series.

The information in this section also applies to the supported models of the Sun StorEdge series and the HP XP series

Related reference

“Supported models of the HDS Lightning” on page 277

The SAN Volume Controller supports models of the HDS Lightning. Certain models of the HDS Lightning are equivalent to Sun StorEdge and HPQ XP models; therefore, the SAN Volume Controller also supports models of the Sun StorEdge and the HPQ XP.

“Supported firmware levels for HDS Lightning” on page 277

The SAN Volume Controller supports the HDS Lightning.

“Concurrent maintenance on the HDS Lightning” on page 278

Concurrent maintenance is the capability to perform I/O operations to an HDS Lightning while simultaneously performing maintenance operations on it.

“User interface on HDS Lightning” on page 278

Ensure that you are familiar with the user interface application that supports the HDS Lightning subsystem.

“Sharing the HDS Lightning 99xxV between a host and the SAN Volume Controller” on page 278
 The HDS Lightning 99xxV can be shared between a host and a SAN Volume Controller with certain restrictions.

“Quorum disks on HDS Lightning 99xxV” on page 280
 HDS Lightning 99xxV is not an approved host for quorum disks. Therefore, configurations with only HDS Lightning are not possible.

“Advanced functions for HDS Lightning” on page 280
 Some advanced functions of the HDS Lightning are not supported by the SAN Volume Controller.

“Logical unit configuration for HDS Lightning” on page 281
 Logical unit (LU) configuration for HDS Lightning supports both RAID 1 and RAID 5 arrays.

“Configuring settings for HDS Lightning” on page 282
 The Lightning configuration interface provides functionality for configuration.

Supported models of the HDS Lightning

The SAN Volume Controller supports models of the HDS Lightning. Certain models of the HDS Lightning are equivalent to Sun StorEdge and HPQ XP models; therefore, the SAN Volume Controller also supports models of the Sun StorEdge and the HPQ XP.

Table 29 lists the models of the HDS Lightning, Sun StorEdge, and HPQ XP that are supported by the SAN Volume Controller.

See the following Web site for the latest supported models:

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

Table 29. Supported models of the HDS Lightning, Sun StorEdge and HPQ XP

HDS Lightning models	Sun StorEdge models	HPQ XP models
Lightning 9910	StorEdge 9910	XP48
Lightning 9960	StorEdge 9960	XP512
Lightning 9970V	StorEdge 9970	XP128
Lightning 9980V	StorEdge 9980	XP1024

Related information

“Configuring the HDS Lightning series subsystem” on page 276
 This section provides information about configuring the Hitachi Data Systems (HDS) Lightning series subsystem so it can attach to a SAN Volume Controller. Certain models of the HDS Lightning series are equivalent to HP and Sun models; therefore, the SAN Volume Controller also supports models of the Sun StorEdge series and the HP XP series.

Supported firmware levels for HDS Lightning

The SAN Volume Controller supports the HDS Lightning.

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

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

Note: Concurrent upgrade of the controller firmware is *not* supported with the SAN Volume Controller.

Related information

“Configuring the HDS Lightning series subsystem” on page 276

This section provides information about configuring the Hitachi Data Systems (HDS) Lightning series subsystem so it can attach to a SAN Volume Controller. Certain models of the HDS Lightning series are equivalent to HP and Sun models; therefore, the SAN Volume Controller also supports models of the Sun StorEdge series and the HP XP series.

Concurrent maintenance on the HDS Lightning

Concurrent maintenance is the capability to perform I/O operations to an HDS Lightning while simultaneously performing maintenance operations on it.

Important: An HDS Field Engineer must perform all maintenance procedures.

Related information

“Configuring the HDS Lightning series subsystem” on page 276

This section provides information about configuring the Hitachi Data Systems (HDS) Lightning series subsystem so it can attach to a SAN Volume Controller. Certain models of the HDS Lightning series are equivalent to HP and Sun models; therefore, the SAN Volume Controller also supports models of the Sun StorEdge series and the HP XP series.

User interface on HDS Lightning

Ensure that you are familiar with the user interface application that supports the HDS Lightning subsystem.

SVP

HDS Lightning has a laptop in the controller frame. The laptop runs SVP as the primary configuration user interface. You can use SVP to perform most configuration tasks and to monitor the controller.

HiCommand

The HiCommand is a graphical user interface that allows basic creation of storage and system monitoring. The HiCommand communicates with HDS Lightning through Ethernet.

Related information

“Configuring the HDS Lightning series subsystem” on page 276

This section provides information about configuring the Hitachi Data Systems (HDS) Lightning series subsystem so it can attach to a SAN Volume Controller. Certain models of the HDS Lightning series are equivalent to HP and Sun models; therefore, the SAN Volume Controller also supports models of the Sun StorEdge series and the HP XP series.

Sharing the HDS Lightning 99xxV between a host and the SAN Volume Controller

The HDS Lightning 99xxV can be shared between a host and a SAN Volume Controller with certain restrictions.

Sharing ports

The HDS Lightning 99xxV can be shared between a host and a SAN Volume Controller with certain restrictions. The following restrictions apply:

- The same host cannot be connected to both a SAN Volume Controller and an HDS Lightning at the same time because HDLM and the subsystem device driver (SDD) cannot coexist.
- A controller port cannot be shared between a host and a SAN Volume Controller. If a controller port is used by a SAN Volume Controller 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.

Supported Topologies

The SAN Volume Controller supports connection to the HDS Lightning with the following restrictions:

- The SAN Volume Controller resolves up to four worldwide node names (WWNNs) per subsystem and allows up to 512 LUs per WWNN. The HDS Lightning assigns a WWNN per port, therefore the SAN Volume Controller can be a limitation to both capacity (2048 LUs) and bandwidth (4 ports). You can use the following procedure for HDS 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 LUs to each group.

The SAN Volume Controller interprets each group as a separate subsystem.

- If an LU is mapped to the SAN Volume Controller port as LUNx, the LU 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 that is managed by the SAN Volume Controller. LUNs that are 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 HDS Lightning.

Related concepts

“Storage subsystems” on page 55

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

Related information

“Configuring the HDS Lightning series subsystem” on page 276

This section provides information about configuring the Hitachi Data Systems (HDS) Lightning series subsystem so it can attach to a SAN Volume Controller. Certain models of the HDS Lightning series are equivalent to HP and Sun models; therefore, the SAN Volume Controller also supports models of the Sun StorEdge series and the HP XP series.

Quorum disks on HDS Lightning 99xxV

HDS Lightning 99xxV is not an approved host for quorum disks. Therefore, configurations with only HDS Lightning are not possible.

Related information

“Configuring the HDS Lightning series subsystem” on page 276

This section provides information about configuring the Hitachi Data Systems (HDS) Lightning series subsystem so it can attach to a SAN Volume Controller. Certain models of the HDS Lightning series are equivalent to HP and Sun models; therefore, the SAN Volume Controller also supports models of the Sun StorEdge series and the HP XP series.

“Creating a quorum disk” on page 236

A quorum disk is used to resolve tie-break situations when the voting set of nodes disagree on the current state of the cluster.

Advanced functions for HDS Lightning

Some advanced functions of the HDS Lightning are not supported by the SAN Volume Controller.

Advanced copy functions

Advanced copy functions for HDS Lightning (for example, ShadowImage, Remote Copy, Data Migration) are not supported for disks that are managed by the SAN Volume Controller because the copy function does not extend to the SAN Volume Controller cache.

LU Expansion

The HDS Lightning 99xxV supports Logical Unit Expansion (LUSE). LUSE is *not* a concurrent operation. LUSE is accomplished by concatenating between 2 and 26 existing logical units (LUs) together. Before LUSE can be performed on an LU, the LU must be removed from the managed disk (MDisk) group and unmapped from the SAN Volume Controller.

Attention: LUSE destroys all data that exists on the LU, except on a Windows system.

TrueCopy

TrueCopy is functionally similar to Metro Mirror. TrueCopy is not supported when the disk controller system is used with the SAN Volume Controller. Even when an HDS Lightning 99xxV is shared between a host and a SAN Volume Controller, TrueCopy is not supported on the ports that are zoned directly with the host.

Virtual LVI/LUNs

The HDS Lightning 99xxV supports Virtual LVI/LUNs. Virtual LVI/LUNs is *not* a concurrent operation. Virtual LVI/LUNs allows you to divide LUNs into several smaller virtual LUNs for use by the HDS Lightning. You must first create existing LUNs into free space and then define their own LUNs using that free space. Virtual LVI/LUNs must *not* be managed or mapped to a SAN Volume Controller.

LUNs that are set up using either LUSE or Virtual LVI/LUNs appear as normal LUNs after they are created. Therefore, LUNs that are set up using LUSE or Virtual LVI/LUNs can be used by the SAN Volume Controller after they are created.

Write protect

LUs cannot be explicitly set to write-protected. However, some of the advanced features, such as Metro Mirror, can be used to write-protect an LU as part of the function. Metro Mirror must not be used for LUs that are in use by a SAN Volume Controller.

Related concepts

“FlashCopy” on page 35

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

“Metro Mirror” on page 47

Metro Mirror enables you to set up a relationship between two virtual disks (VDisks), so that updates that are made by an application to one VDisk are mirrored on the other VDisk.

Related information

“Configuring the HDS Lightning series subsystem” on page 276

This section provides information about configuring the Hitachi Data Systems (HDS) Lightning series subsystem so it can attach to a SAN Volume Controller. Certain models of the HDS Lightning series are equivalent to HP and Sun models; therefore, the SAN Volume Controller also supports models of the Sun StorEdge series and the HP XP series.

Logical unit configuration for HDS Lightning

Logical unit (LU) configuration for HDS Lightning supports both RAID 1 and RAID 5 arrays.

The HDS Lightning subsystem can have up to 8192 LUs defined; however, only 256 LUs can be mapped to a single port. Report LUNs is supported by LUN 0, so the SAN Volume Controller can detect all LUNs.

In the event that a LUN 0 is not configured, the HDS Lightning subsystem presents a pseudo LUN at LUN 0. The inquiry data for this pseudo LUN slightly differs from the inquiry data of normal LUNs. The difference allows the SAN Volume Controller to recognize the pseudo LUN and exclude it from I/O. The pseudo LUN can accept the report LUNs command.

The HDS Lightning subsystem supports both open-mode attachment and S/390 attachment. The emulation mode is set when the LU is defined. All LUNs that are presented to a SAN Volume Controller must use open emulation. All LUNs with open emulation use a standard 512 byte block size.

The HDS Lightning subsystem can only have certain sized LUs defined. These LUs can be expanded by merging 2 - 36 of these LUs together using the Logical Unit Size Expansion (LUSE) feature. They can also be made into several, smaller virtual LUNs by using the Virtual LVI/LUN feature.

Special LUs

When an LU is mapped to a host, you have the option to make it a *command LUN*. Command LUNs support in-band configuration commands, but not I/O. Therefore, you cannot map command LUNs to the SAN Volume Controller.

Related reference

“Logical unit creation and deletion on HDS Lightning” on page 282

The SAN Volume Controller supports logical unit expansion (LUSE) with certain restrictions.

Related information

“Configuring the HDS Lightning series subsystem” on page 276

This section provides information about configuring the Hitachi Data Systems (HDS) Lightning series subsystem so it can attach to a SAN Volume Controller. Certain models of the HDS Lightning series are equivalent to HP and Sun models; therefore, the SAN Volume Controller also supports models of the Sun StorEdge series and the HP XP series.

Logical unit creation and deletion on HDS Lightning

The SAN Volume Controller supports logical unit expansion (LUSE) with certain restrictions.

The following restrictions apply:

- Before LUSE can be performed on an LU, the LU must be unmounted from a host, and have no available paths. The LUSE function destroys all data that exists on the LU, except for LUs on a Windows operating system.
- LUSE must not be performed on any disk that is managed by the SAN Volume Controller.
- If data exists on a disk and you want to use image mode to import the data, LUSE must not be used on the disk before you import the data.

Related concepts

“Storage subsystems” on page 23

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

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

“Configuring IBM DS4000 series disk controllers for the storage server” on page 264

The IBM DS4000 series of disk controllers provide functionality that is compatible with the SAN Volume Controller.

Configuring settings for HDS Lightning

The Lightning configuration interface provides functionality for configuration.

These options and settings can have a scope of the following:

- Subsystem
- Port
- Logical unit (LU)

Related reference

“Global settings for HDS Lightning” on page 283

Global settings apply across an HDS Lightning disk controller system.

“Controller settings for HDS Lightning” on page 283

Controller settings are settings that apply across the entire HDS Lightning controller.

“Port settings for HDS Lightning” on page 283

Port settings are configurable at the port level.

“Logical unit settings for HDS Lightning” on page 284
 Logical unit (LU) settings apply to individual LUs that are configured in the HDS Lightning controller.

Related information

“Configuring the HDS Lightning series subsystem” on page 276
 This section provides information about configuring the Hitachi Data Systems (HDS) Lightning series subsystem so it can attach to a SAN Volume Controller. Certain models of the HDS Lightning series are equivalent to HP and Sun models; therefore, the SAN Volume Controller also supports models of the Sun StorEdge series and the HP XP series.

Global settings for HDS Lightning

Global settings apply across an HDS Lightning disk controller system.

Table 30 lists the global settings for HDS Lightning.

Table 30. HDS Lightning global settings supported by the SAN Volume Controller

Option	Lightning default setting	SAN Volume Controller Required setting
Spare disk recover	Interleave	Interleave
Disk copy place	Medium	Medium
Copy operation	Correction copy and dynamic sparing	Correction copy and dynamic sparing
Read configuration data mode	Selected	Selected
PS off timer	Not selected	Not selected

Related reference

“Configuring settings for HDS Lightning” on page 282
 The Lightning configuration interface provides functionality for configuration.

Controller settings for HDS Lightning

Controller settings are settings that apply across the entire HDS Lightning controller.

Table 31 lists the HDS Lightning controller settings that are supported by the SAN Volume Controller.

Table 31. HDS Lightning controller settings supported by the SAN Volume Controller

Option	HDS Lightning default setting	SAN Volume Controller required setting
PCB mode	Standard	Standard

Related reference

“Configuring settings for HDS Lightning” on page 282
 The Lightning configuration interface provides functionality for configuration.

Port settings for HDS Lightning

Port settings are configurable at the port level.

There are no options available with the scope of a single controller.

- The ports are included in switch zones.
- The switch zones only present the ports directly to the hosts and not to a SAN Volume Controller.

Table 32 lists the HDS Lightning port settings that are supported by the SAN Volume Controller.

Table 32. HDS Lightning port settings supported by the SAN Volume Controller

Option	HDS Lightning default setting	SAN Volume Controller required setting
Address	AL/PA	AL/PA
Fabric	On	On
Connection	Point-to-Point	Point-to-Point
Security switch	On	On or off
Host type	Default	Windows

Related reference

“Configuring settings for HDS Lightning” on page 282

The Lightning configuration interface provides functionality for configuration.

Logical unit settings for HDS Lightning

Logical unit (LU) settings apply to individual LUs that are configured in the HDS Lightning controller.

HDS Lightning LUs must be configured as described in Table 33 if the LUN is associated with ports in a switch zone that is accessible to the SAN Volume Controller.

Table 33. HDS Lightning LU settings for the SAN Volume Controller

Option	HDS Lightning default setting	SAN Volume Controller required setting
Command device	Off	Off
Command security	Off	Off

Note: These settings only apply to LUs that are accessible by the SAN Volume Controller.

Related reference

“Configuring settings for HDS Lightning” on page 282

The Lightning configuration interface provides functionality for configuration.

Configuring the HDS Thunder subsystem

This section provides information about configuring the HDS Thunder subsystem so it can attach to a SAN Volume Controller.

Related tasks

“Setting up an HDS Thunder with more than four ports” on page 288

You can set up an HDS Thunder with more than four ports.

Related reference

“Supported models of the HDS Thunder” on page 285

The SAN Volume Controller can be used with models of the HDS Thunder.

“Supported firmware levels for HDS Thunder” on page 286

The SAN Volume Controller supports the HDS Thunder.

“Concurrent maintenance on the HDS Thunder” on page 286

Concurrent maintenance is the capability to perform I/O operations to an HDS Thunder while simultaneously performing maintenance operations on it.

“User interface on the HDS Thunder” on page 286

Ensure that you are familiar with the user interface applications that support the HDS Thunder subsystem.

“Sharing the HDS Thunder between host and the SAN Volume Controller” on page 287

The HDS Thunder 9200 and 95xxV can be shared between a host and a SAN Volume Controller with certain restrictions.

“Quorum disks on HDS Thunder” on page 288

Managed disks presented by the Thunder 9200 and 95xxV may be chosen by the SAN Volume Controller as quorum disks.

“Advanced functions for HDS Thunder” on page 289

Some advanced functions of the HDS Thunder are not supported by the SAN Volume Controller.

“Logical unit creation and deletion on HDS Thunder” on page 290

Before you create or delete a logical unit on HDS Thunder, consider the following constraints.

“Configuring settings for HDS Thunder” on page 291

The Thunder configuration interface provides functionality for configuration.

Supported models of the HDS Thunder

The SAN Volume Controller can be used with models of the HDS Thunder.

Tables Table 34 and Table 35 list the models of the HDS Thunder that are supported by the SAN Volume Controller.

See the following Web site for the latest supported models:

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

Table 34. Supported Thunder 9200 models

Model	Description
Thunder 9200 Rackmount	Up to 100 disks
Thunder 9200 Deskside 20	Maximum of 20 disks
Thunder 9200 Deskside 10	Maximum of 10 disks

Table 35. Supported Thunder 95xxV models

Model	Description
Thunder 9520V Deskside	-
Thunder 9530V Deskside/Rackmount	Supports 2 - 14 disks
Thunder 9531V Deskside	Pre-configured with 5 disks
Thunder 9532V Deskside	Pre-configured with 9 disks
Thunder 9533V Deskside	Pre-configured with 13 disks
Thunder 9570V Deskside/Rackmount	Supports 2 - 224 disks
Thunder 9580V Rackmount	Supports 5 - 449 disks
Thunder 9585V Rackmount	Supports 5 - 449 disks

Related information

“Configuring the HDS Thunder subsystem” on page 284
This section provides information about configuring the HDS Thunder subsystem so it can attach to a SAN Volume Controller.

Supported firmware levels for HDS Thunder

The SAN Volume Controller supports the HDS Thunder.

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

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

Note: Concurrent upgrade of the controller firmware is supported with the SAN Volume Controller.

Related information

“Configuring the HDS Thunder subsystem” on page 284
This section provides information about configuring the HDS Thunder subsystem so it can attach to a SAN Volume Controller.

Concurrent maintenance on the HDS Thunder

Concurrent maintenance is the capability to perform I/O operations to an HDS Thunder while simultaneously performing maintenance operations on it.

Important: An HDS Field Engineer must perform all maintenance operations.

The SAN Volume Controller and Thunder support concurrent hardware maintenance and firmware upgrade operations.

Related information

“Configuring the HDS Thunder subsystem” on page 284
This section provides information about configuring the HDS Thunder subsystem so it can attach to a SAN Volume Controller.

User interface on the HDS Thunder

Ensure that you are familiar with the user interface applications that support the HDS Thunder subsystem.

In-band configuration

Disable the HDS Thunder command LUN when you use the user interface applications.

DAMP

The Disk Array Management Program (DAMP) is the primary user interface application for configuring the HDS Thunder. Use DAMP to upgrade firmware, change settings, and create and monitor storage.

DAMP supports an Ethernet connection to the HDS Thunder. An out-of-band command-line interface is available with DAMP that supports the majority of the functions that are provided in DAMP.

HiCommand

HiCommand is another configuration user interface that is available for the HDS Thunder. You must have access to DAMP to use HiCommand to configure settings. HiCommand is more restricted than DAMP. It allows basic creation of storage and provides some monitoring features. HiCommand works for both HDS Thunder and HDS Lightning subsystems.

HiCommand uses Ethernet to connect to the HDS Thunder.

Web Server

A Web server runs on each of the controllers on the subsystem. During normal operation, the user interface only allows basic monitoring of the subsystem and displays an error log. If a controller is put into diagnostic mode by pressing the reset button on the controller, the user interface allows for firmware upgrades and subsystem configuration resets.

Related information

“Configuring the HDS Thunder subsystem” on page 284

This section provides information about configuring the HDS Thunder subsystem so it can attach to a SAN Volume Controller.

Sharing the HDS Thunder between host and the SAN Volume Controller

The HDS Thunder 9200 and 95xxV can be shared between a host and a SAN Volume Controller with certain restrictions.

The following restrictions apply when you share the HDS Thunder between a host and a SAN Volume Controller:

- The same host cannot be connected to both a SAN Volume Controller and a Thunder at the same time because Hitachi Dynamic Link Manager (HDLM) and the subsystem device driver (SDD) do not coexist.
- For Thunder 9200 only, a target port cannot be shared between a host and a SAN Volume Controller. In other words, if a target port is used by a SAN Volume Controller 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. Thus, Thunder 9200 must be set into M-TID M-LUN mode and Mapping Mode enabled on Thunder 95xx. No LU can have a LUN Number associated with a port which is zoned for host use while also having a LUN Number associated with a port which is zoned for a SAN Volume Controller.

Related concepts

“Storage subsystems” on page 55

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

Related information

“Configuring the HDS Thunder subsystem” on page 284

This section provides information about configuring the HDS Thunder subsystem so it can attach to a SAN Volume Controller.

Setting up an HDS Thunder with more than four ports

You can set up an HDS Thunder with more than four ports.

Perform the following steps to set up an HDS Thunder with more than four ports:

1. Set the Mapping Mode to **Enabled**.
2. Divide the ports into groups of four (or two). For redundancy, at least one port from each controller should be in each group.
3. Make a note of all of the LUNs currently on the array. Each LUN that you want managed by the San Volume Controller should be in one group.
4. Divide the LUNs into groups: You should have one group of LUNs for each group of ports.
5. From the **Host Groups** view:

- a. Select the first port in the first port group.

- b. Select **Option**

Set the port options.

Select **Logical Unit**.

Select **Modify Mapping** from the list.

Perform the following steps from the Modify Mapping panel:

- 1) Select a LUN from the first LUN group from the "LUN" column

- 2) Select "Host LUN" 0, and click **Add**.

This repositions the mapping to the "reserved configuration" column.

- 3) Select the next LUN from the first group

- 4) Select "Host LUN" 1, and click **Add**.

Repeat the previous step for all ports in the first port group. Ensure that the LUN and Host LUN IDs are identical for all ports. Failure to make these identical causes I/O failures.

- 5) Repeat the previous two steps for all port groups.

Related information

"Configuring the HDS Thunder subsystem" on page 284

This section provides information about configuring the HDS Thunder subsystem so it can attach to a SAN Volume Controller.

Quorum disks on HDS Thunder

Managed disks presented by the Thunder 9200 and 95xxV may be chosen by the SAN Volume Controller as quorum disks.

Managed disks presented by the Thunder 9200 and 95xxV may be chosen by the SAN Volume Controller as quorum disks during initialization of the cluster. The selection made can be changed by the following methods:

- **Set quorum disk** command
- Setting a Quorum Disk panel

Related information

"Configuring the HDS Thunder subsystem" on page 284

This section provides information about configuring the HDS Thunder subsystem so it can attach to a SAN Volume Controller.

"Creating a quorum disk" on page 236

A quorum disk is used to resolve tie-break situations when the voting set of nodes disagree on the current state of the cluster.

Advanced functions for HDS Thunder

Some advanced functions of the HDS Thunder are not supported by the SAN Volume Controller.

Advanced copy functions

Advanced copy functions for HDS Thunder (for example: ShadowImage, TrueCopy, HiCopy) are not supported for disks that are managed by the SAN Volume Controller because the copy function does not extend to the SAN Volume Controller cache.

LUN Security

LUN Security enables LUN masking by the worldwide name (WWN) of the initiator port. This function is not supported for logical units (LUs) that are used by the SAN Volume Controller.

Partitioning

Partitioning splits a RAID array into up to 128 smaller LUs, each of which serves as an independent disk like entity. The SAN Volume Controller and HDS Thunder supports the partitioning function.

Dynamic array expansion

The HDS Thunder allows the last LU that is defined in a RAID group to be expanded. This function is not supported with the SAN Volume Controller attachment. Do *not* perform dynamic array expansion on LUs that are in use by a SAN Volume Controller.

Note: Use in this context means that the LU has a LUN number that is associated with a fibre-channel port, and this fibre-channel port is contained in a switch zone that also contains SAN Volume Controller fibre-channel ports.

Host storage domains and virtual fibre-channel ports

The HDS Thunder 95xxV supports host storage domains (HSD) and virtual fibre-channel ports. Each fibre-channel port can support multiple HSDs. Each host in a given HSD is presented with a virtual target port and a unique set of LUNs.

The Thunder 9200 does not support HSD and virtual fibre-channel ports.

Related concepts

“FlashCopy” on page 35

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

“Metro Mirror” on page 47

Metro Mirror enables you to set up a relationship between two virtual disks (VDisks), so that updates that are made by an application to one VDisk are mirrored on the other VDisk.

Related information

“Configuring the HDS Thunder subsystem” on page 284

This section provides information about configuring the HDS Thunder subsystem so it can attach to a SAN Volume Controller.

Logical unit creation and deletion on HDS Thunder

Before you create or delete a logical unit on HDS Thunder, consider the following constraints.

The Thunder configuration interface enables you to create and delete logical unit number (LUNs). You must avoid certain creation and deletion scenarios to prevent data corruption. This topic discusses those scenarios.

Creation and deletion scenarios

The Thunder configuration interface enables you to create and delete LUNs. Certain creation and deletion scenarios must be avoided to prevent data corruption. For example, the configuration interface enables you to create LUN A, delete LUN A, and then create LUN B with the same unique ID as LUN A. Doing this with a SAN Volume Controller attached could cause data corruption because the SAN Volume Controller might not realize that LUN B is different than LUN A.

Attention: Before you delete a LUN using the Thunder configuration interface, the LUN must first be removed from the managed disk group that contains it.

Dynamic addition of LUNs

Perform the following procedure to add LUNs dynamically. Using this procedure prevents the existing LUNs from rejecting I/O and returning a status of unavailable during dynamic addition of LUNs.

1. Create the new LUNs using the Disk Array Management Program (DAMP), which is the Thunder configuration tool.
2. Quiesce all I/O.
3. Perform either an offline format or an online format of all new LUNs on the controller using DAMP. Wait for the format to complete.
4. Go into the LUN mapping function of DAMP. Add mapping for the new LUN to all of the controller ports that are available to the SAN Volume Controller on the fabric.
5. Restart the controller. (Model 9200 only)
6. After the controller has restarted, restart I/O.

LUN mapping considerations

If LUN mapping is used as described in the LUN mapping topic, the controller must be restarted to pick up the new LUN mapping configuration. For each managed disk group (MDisk) group that contains an MDisk that is supported by an LU on the Thunder disk controller, all virtual disks in those MDisk groups will go offline.

Related concepts

“Storage subsystems” on page 23

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

“MDisks” on page 24

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

Related reference

“Mapping and virtualization settings for HDS Thunder” on page 295
Thunder supports different modes of operation. These modes affect LUN mapping or masking and virtualization.

Related information

“Configuring the HDS Thunder subsystem” on page 284

This section provides information about configuring the HDS Thunder subsystem so it can attach to a SAN Volume Controller.

Configuring settings for HDS Thunder

The Thunder configuration interface provides functionality for configuration.

These options and settings can have a scope of a:

- Subsystem
- Port
- Logical unit

Related concepts

“Storage subsystems” on page 23

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

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Related reference

“Global settings for the HDS Thunder”

Global settings apply across a Thunder disk controller system.

“Controller settings for HDS Thunder” on page 293

Controller settings are settings that apply across the entire HDS Thunder controller. There are no options available with the scope of a single controller.

“Port settings for the HDS Thunder” on page 293

Port settings are configurable at the port level.

“Logical unit settings for the HDS Thunder” on page 294

Logical unit (LU) settings apply to individual LUs configured in the Thunder controller.

“Mapping and virtualization settings for HDS Thunder” on page 295

Thunder supports different modes of operation. These modes affect LUN mapping or masking and virtualization.

Related information

“Configuring the HDS Thunder subsystem” on page 284

This section provides information about configuring the HDS Thunder subsystem so it can attach to a SAN Volume Controller.

“Servicing storage subsystems” on page 237

Storage subsystems that are supported for attachment to the SAN Volume Controller are designed with redundant components and access paths to allow concurrent maintenance. Hosts have continuous access to their data during component failure and replacement.

Global settings for the HDS Thunder

Global settings apply across a Thunder disk controller system.

The following table lists the global settings for HDS Thunder.

Table 36. Thunder global settings supported by the SAN Volume Controller

Option	Thunder Default Setting	SAN Volume Controller Required Setting
Start attribute	Dual active mode	Dual active mode
SCSI ID/Port takeover mode	N/A	N/A
Default controller	N/A	N/A
Data-share mode	Used	Used
Serial number		Same as the Thunder default setting
Delay planned shutdown	0	0
Drive detach mode	False	False
Multipath controller (Thunder 9200 only)	False	False
PROCOM mode	False	False
Report status	False	False
Multipath (Array unit)	False	False
Turbo LU warning	False	False
NX mode	False	False
Auto reconstruction mode	False	False
Forced write-through mode	False	False
Changing logical unit mode 1	False	False
Multiple stream mode (Thunder 9200 only)	False	False
Multiple stream mode (write) (Thunder 95xxV only)	False	False
Multiple stream mode (read) (Thunder 95xxV only)	False	False
RAID 3 mode (Thunder 9200 only)	False	False
Target ID (9200 only) Mapping mode on 95xx	S-TID, M-LUN	M-TID, M-LUN (if sharing controller, otherwise S-TID, M-LUN)
Data striping size	16K; 32K; 64K	Any (Thunder 9200) 64K (Thunder 95xxV)
Operation if processor failure occurs	Reset the fault	Reset the fault
Command queuing	True	True
ANSI Version	N/A	N/A
Vendor ID	HITACHI	HITACHI
Product ID (Thunder 9200)	DF500F	DF500F
Product ID (Thunder 95xxV)	DF500F	DF600F
ROM microprogram version	<Empty>	<Empty>
RAM microprogram version	<Empty>	<Empty>

Table 36. Thunder global settings supported by the SAN Volume Controller (continued)

Option	Thunder Default Setting	SAN Volume Controller Required Setting
Web title	<Empty>	Any setting supported
Cache mode (Thunder 9200 only)	All off	All off
Link separation (Thunder 9200 only)	False	False
ROM Pseudo-response command processing (Thunder 9200 only)	N/A	N/A
Save data pointer response (Thunder 9200 only)	N/A	N/A
Controller identifier	False	False
RS232C error information outflow mode	Off	Any
Execute write and verify mode	True	True

Controller settings for HDS Thunder

Controller settings are settings that apply across the entire HDS Thunder controller. There are no options available with the scope of a single controller.

Port settings for the HDS Tunder

Port settings are configurable at the port level.

This topic lists those per port settings, what the Thunder defaults are, and what the required settings are for the SAN Volume Controller. Port settings for clearing LUN reservations are also listed.

The settings listed in the table Table 37 apply to those HDS Thunder 9200 disk controllers that are in a switch zone that contains SAN Volume Controllers. If the Thunder disk controller is shared between a SAN Volume Controller and another host, you can configure with different settings than shown if both of the following conditions are true:

There are no options available with the scope of a single controller.

- The ports are included in switch zones
- The switch zones only present the ports directly to the hosts and not to a SAN Volume Controller

Table 37. Thunder port settings supported by the SAN Volume Controller

Option	Thunder Default Setting	SAN Volume Controller Required Setting
Host connection mode 1	Standard	Standard
VxVM DMP mode (Thunder 9200 only)	False	False
HP connection mode	False	False
Report inquiry page 83H (Thunder 9200 only)	False	True
UA (06/2A00) suppress mode	False	False

Table 37. Thunder port settings supported by the SAN Volume Controller (continued)

Option	Thunder Default Setting	SAN Volume Controller Required Setting
HISUP mode	False	False
CCHS mode	False	False
Standard inquiry data expand (Thunder 9200 only)	False	False
Host connection mode 2	False	False
Product ID DF400 mode	False	False
HBA WWN report mode (Thunder 9200 only)	False	False
NACA mode	False	False
SUN cluster connection mode	False	False
Persistent RSV cluster mode	False	False
ftServer connection mode 1 (Thunder 9200 only)	False	False
ftServer connection mode 2	False	False
SRC Read Command reject	False	False
Reset/LIP mode (signal)	False	False
Reset/LIP mode (progress)	False	False
Reset ALL LIP port mode	False	False
Reset target (reset bus device mode)	False	True
Reserve mode	False	True
Reset logical unit mode	False	True
Reset logout of third party process mode	False	False
Read Frame minimum 128 byte mode (Thunder 950xxV only)	False	False
Topology	Point-to-point	Fabric

Logical unit settings for the HDS Thunder

Logical unit (LU) settings apply to individual LUs configured in the Thunder controller.

Thunder LUs must be configured as described in Table 38 if the logical unit number (LUN) is associated with ports in a switch zone that is accessible to the SAN Volume Controller.

Table 38. Thunder LU settings for the SAN Volume Controller

Option	Required Values	Default Setting
LUN default controller	Controller 0 or Controller 1	Any

Note: These settings only apply to LUs accessible by SAN Volume Controller.

Data corruption scenarios to avoid

Scenario 1: The configuration application enables you to change the serial number for an LU. Changing the serial number also changes the unique user identifier (UID) for the LU. Since the serial number is also used to determine the WWPN of the controller ports, two LUNs cannot have the same unique ID on the same SAN because two controllers cannot have the same WWPN on the same SAN.

Scenario 2: The serial number is also used to determine the WWPN of the controller ports. Therefore, two LUNs must not have the same ID on the same SAN, because this results in two controllers having the same WWPN on the same SAN. This is not a valid configuration.

Attention: Do not change the serial number for an LU that is managed by a SAN Volume Controller because this could result in data loss or undetected data corruption.

Scenario 3: The configuration application enables you to create LUN A, delete LUN A, and create LUN B with the same unique ID as LUN A. Doing this with a LUN managed by a SAN Volume Controller could result in data corruption because the SAN Volume Controller might not recognize that LUN B is different than LUN A.

Related concepts

“Storage subsystems” on page 23

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

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Mapping and virtualization settings for HDS Thunder

Thunder supports different modes of operation. These modes affect LUN mapping or masking and virtualization.

SAN Volume Controller supports the S-TID M-LUN and M-TID M-LUN modes on Thunder 9200, and Mapping Mode enabled or disabled on Thunder 95xx.

Attention: Thunder does not provide an interface that enables a SAN Volume Controller to detect and ensure that the mapping or masking and virtualization options are set properly. Therefore, you or someone else in your organization must ensure that these options are set as outlined in this topic.

S-TID M-LUN modes

In S-TID M-LUN mode all LUs are accessible via all ports on Thunder with the same LUN Number on each port. This is the simplest mode and it should be used for all situations, except, where a Thunder subsystem is being shared between a host and a SAN Volume Controller.

M-TID M-LUN modes

If a Thunder is being shared between a host and a SAN Volume Controller, then you must use M-TID M-LUN mode. Configure your Thunder so that all SAN Volume

Controllers accessible LUs have the same LUN Number on all ports through which they can be accessed.

A SAN Volume Controller can access controller ports x and y. The SAN Volume Controller also sees an LU on port x that has LUN number p. In this situation the following conditions must be met:

- The SAN Volume Controller must see either the same LU on port y with LUN number p or it must not see the LU at all on port y.
- The LU should not appear as any other LUN number on port y.
- The LU must not be mapped to any Thunder port which is zoned for use directly by a host in a configuration where the Thunder is shared between a host and a SAN Volume Controller.

M-TID M-LUN mode enables LU virtualization by target port. In this mode, a single LU can be seen as different LUN numbers across all of the controller ports. For example, LU A may be LUN 0 on controller port 1, LUN 3 on controller port 2, and not visible at all on controller ports 3 and 4.

Note: The SAN Volume Controller does not support this.

In addition, M-TID M-LUN mode enables a single LU to be seen as multiple LUN numbers on the same controller port. For example, LU B may be LUN 1 and LUN 2 on controller port 1.

Note: The SAN Volume Controller does not support this.

Note: Thunder 9200 controllers require a reboot in order for changes to LUN mapping to take effect.

Related concepts

“Storage subsystems” on page 23

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

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Configuring the HP StorageWorks subsystem

This section provides information about configuring the HP StorageWorks subsystem, which uses the HSG80 controller, so it can attach to a SAN Volume Controller.

Support for HSG80 is limited in SAN Volume Controller version 1.1.1 to a single port connection. Version 1.2.0 enables multiport connection but this introduces some restrictions on LUN partitioning.

MDisk Groups and MDisks

Attention: This topic only applies to SAN Volume Controller version 1.1.1 when only a single port connection was supported.

A managed disk (MDisk) group should contain either no HSG80 LUNs or LUNs that are only from a single HSG80 subsystem. *No other configuration is supported.* An MDisk group that consists of LUNs from HSG80 storage and non-HSG80 storage would potentially contain a single point of failure, if the HSG80 subsystem was connected to the cluster by a single port. Consequently, any virtual disks (VDisks) that are created from such an MDisk group can potentially contain a single point of failure.

Related concepts

“MDisk groups” on page 27

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

Related tasks

“Configuring the HP StorageWorks controller” on page 300

To configure the HP StorageWorks HSG80 controller, complete these steps.

“Creating MDisk groups” on page 112

You can create a new managed disk (MDisk) group using the Create a Managed Disk Group wizard.

Related reference

“Supported models of the HP StorageWorks controller” on page 303

The following table lists the models of the HP StorageWorks controller that are supported by the SAN Volume Controller. See the following Web site for the latest supported models.

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

“Supported firmware levels for the HP StorageWorks controller” on page 304

See the following Web site for specific HP StorageWorks firmware levels and the latest supported hardware.

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

“Concurrent maintenance on the HP StorageWorks” on page 304

Concurrent maintenance is the capability to perform I/O operations to a HP StorageWorks controller while simultaneously performing maintenance operations on it.

“Configuration interface for the HP StorageWorks” on page 305

The Command Console configuration and service utility is the configuration interface for the HSG80 subsystem.

“Sharing the HP StorageWorks controller between a host and the SAN Volume Controller” on page 305

The HP StorageWorks controller can be shared between a host and a SAN Volume Controller.

“Switch zoning limitations for the HP StorageWorks subsystem” on page 306

Consider the following limitations when planning switch zoning and connection to the SAN.

“Quorum disks on HP StorageWorks” on page 306

Managed disks presented by the HP StorageWorks controllers will be chosen by the SAN Volume Controller as quorum disks.

“Advanced functions for HP StorageWorks” on page 307

Some advanced functions of the HP StorageWorks are not supported by the SAN Volume Controller.

“HP StorageWorks advanced functions” on page 308

VDisks that are created from MDisks presented by an HSG80 controller may be used in SAN Volume Controller FlashCopy mappings or SAN Volume Controller Metro Mirror relationships.

“Logical unit creation and deletion on the HP StorageWorks” on page 308
 Ensure you are familiar with the HSG80 container types for logical unit configuration.

“Configuring settings for the HP StorageWorks” on page 308
 The HP StorageWorks HSG80 configuration interface provides configuration settings and options that are supported with the SAN Volume Controller.

Related information

“HP StorageWorks definitions”

The following terms are used in the IBM and HP documentation and have different meanings.

HP StorageWorks definitions

The following terms are used in the IBM and HP documentation and have different meanings.

IBM term	IBM definition	HP term	HP definition
container	A visual user-interface component that holds objects.	container	(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.
device	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.	device	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.
just a bunch of disks (JBOD)	See <i>non-RAID</i> .	just a bunch of disks (JBOD)	A group of single-device logical units not configured into any other container type.

IBM term	IBM definition	HP term	HP definition
mirrorset	See <i>RAID 1</i> .	mirrorset	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.
non-RAID	Disks that are not in a redundant array of independent disks (RAID).	non-RAID	See <i>just a bunch of disks</i> .
RAID 0	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.	RAID 0	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.
RAID 1	A form of storage array in which two or more identical copies of data are maintained on separate media. Also known as mirrorset.	RAID 1	See <i>mirrorset</i> .
RAID 5	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.	RAID 5	See <i>RAIDset</i> .

IBM term	IBM definition	HP term	HP definition
RAIDset	See <i>RAID 5</i> .	RAIDset	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.
partition	A logical division of storage on a fixed disk.	partition	A logical division of a container represented to the host as a logical unit.
stripeset	See <i>RAID 0</i> .	stripeset	See <i>RAID 0</i> .

Configuring the HP StorageWorks controller

To configure the HP StorageWorks HSG80 controller, complete these steps.

Assume that the HP StorageWorks subsystem is not in use.

Note: When you configure a SAN Volume Controller cluster to work with an HSG80 controller, you must not exceed the limit of 96 process logins.

1. Verify that the SAN Volume Controller front panel is clear of 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:
 - a. That the controller software is at the supported level. See <http://www.ibm.com/servers/storage/support/virtual/2145.html>.
 - b. Ensure that the controllers are configured for MULTIBUS FAILOVER with each other.
 - c. Ensure that the controllers are running in SCSI-3 mode.
 - d. Ensure that MIRRORED_CACHE is enabled.
 - e. Ensure that the Host Connection Table is *not* locked.
5. Issue the **SHOW DEVICES FULL** command to verify the following:
 - a. Ensure that none of the LUNs are TRANSPORTABLE.
 - b. 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:
 - a. Ensure that all LUNs are set to RUN and NOWRITEPROTECT
 - b. Ensure that all LUNs are ONLINE to either THIS or OTHER controller.
 - c. Ensure that all LUNs that are to be made available to the SAN Volume Controller have ALL access.
 - d. Ensure that all LUNs have Host Based Logging NOT specified.
8. If you have partitioned LUNs, refer to the HP StorageWorks controllers topic.
9. Issue the **SHOW CONNECTIONS FULL** command to verify that you have enough spare entries for all combinations of SAN Volume Controller ports and HP StorageWorks ports.
10. Connect up to four known good fibre-channel cables between your fibre-channel switches and your HP StorageWorks subsystem.
11. Ensure that your fibre-channel switches are zoned such that the SAN Volume Controller and the HP StorageWorks subsystem are in a zone. Refer to zoning a switch.
12. 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).
13. Issue the **SHOW CONNECTIONS FULL** command to verify that the new connections have appeared for each SAN Volume Controller port and HP StorageWorks port combination.
14. Verify that No rejected hosts is displayed at the end of the SHOW CONNECTIONS output.
15. On the SAN Volume Controller, issue the **svctask detectmdisk** command to discover the controller.
16. Issue the **svcinfolcontroller** command to verify that the two HSG80 serial numbers appear under the `ctrl s/n`.
17. Issue the **svcinfolmdisk** command to verify that the additional MDisks that correspond to the UNITS shown in the HP StorageWorks subsystem.

You can now use the SAN Volume Controller commands to create an MDisk group. You can also create and map VDIs from these MDisk groups. Check your SAN Volume Controller front panel has no 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 tasks

“Creating MDisk groups using the CLI” on page 158

You can use the command-line interface (CLI) to create a managed disk (MDisk) group.

“Creating VDIs” on page 161

You can use the command-line interface (CLI) to create a virtual disk (VDisk).

Related reference

“HP StorageWorks controllers” on page 302

The SAN Volume Controller supports HP StorageWorks Controllers.

“Switch zoning for the SAN Volume Controller” on page 65
Ensure that you are familiar with the constraints for zoning a switch.

HP StorageWorks controllers

The SAN Volume Controller supports HP StorageWorks Controllers.

A SAN Volume Controller with level 1.2.0 and later supports up to four fibre-channel connections per HP StorageWorks subsystem. The support for partitioned LUNs is restricted to a single fibre-channel connection.

Attention: Before you make any changes to your system, back up important application data.

You can use the HP StorageWorks command **SHOW UNITS** to display all units that are partitioned. Table 39 provides an example of the information that is provided by the **SHOW UNITS** command.

Table 39. Determining partition usage

HSG80 "SHOW UNITS" LUN	Uses	Used by
D1	R50	-
D2	R52	-
D3	R53	(partition)
D4	R54	-
D5	DISK50000	(partition)
D6	D51	-
D7	DISK30300	(partition)
D8	DISK10000	(partition)
D9	R55	-

Here *D3*, *D5*, *D7* and *D8* are partitioned units.

Scenario 1

This scenario assumes that you have no partitioned units on any HP StorageWorks controllers.

If there are no partitions on any of the HP StorageWorks controllers that are or will be connected to the SAN Volume Controller, ensure that the SAN Volume Controller has level 1.2.0 or later installed on each of the SAN Volume Controller clusters. If level 1.2.0 or later is installed, additional fibre-channel connections can be zoned and physically connected.

Scenario 2

This scenario assumes that you are using HP StorageWorks controllers with a SAN Volume Controller that has level 1.1.1 installed, and uses a single fibre-channel attachment or zone. If partitions are present on the HP StorageWorks controllers, there are two options available:

Option 1: Migrate data from the partitioned units

Migrate the data that resides on partitioned units and then delete the partitioned units. Perform the following steps to migrate your data:

1. Perform a concurrent code load to install level 1.2.0 or later.
2. Migrate the data that resides on partitioned units by performing one of the following actions:
 - Use the **svctask migratevdisk** command-line interface (CLI) command to migrate all virtual disks (VDisks) that are in groups that include at least one partitioned unit to groups that contain no partitioned units. You can use the **svcinfo lsmdisk** CLI command and the **SHOW UNITS FULL** command to correlate which HP StorageWorks units correspond with which managed disks (MDisks) on the SAN Volume Controller by comparing the unit identifiers (UIDs).
 - Ensure that the MDisk groups have enough unused space on the MDisks that correspond to unpartitioned units for a copy of all the data on MDisks that correspond to the partitioned units. Delete the MDisks using the **svctask rmmdisk** CLI command.
3. Rezone to utilize the extra ports on the HP StorageWorks controller.

Option 2: Retain partitioned units

Perform a concurrent code load to upgrade the SAN Volume Controller to level 1.2.0. or later. Retain the partitioned units and continue to use a single fibre-channel attachment.

Note: You must not zone in any additional fibre-channel ports on the HP StorageWorks controller because MDisks that are based on partitioned units are taken offline. If you have partitioned LUNs that are not allocated to unit numbers and you subsequently add these to your configuration, these units must be online to the controller that has the zoned in fibre-channel port. Press the reset button on the other controller to bring the units online. This is only necessary for unmanaged MDisks.

Scenario 3

This scenario assumes that you have partitions present on an HP StorageWorks controller that you want to connect to a SAN Volume Controller that already has level 1.2.0 or later installed.

You must initially zone in a single fibre-channel connection to one of the HP StorageWorks controllers, and ensure that all the units are online. Press the reset button on the other controller to bring the units online. You can then select one of the options that are described in Scenario 2. You do not need to perform a concurrent code load because the correct level is already installed.

Supported models of the HP StorageWorks controller

The following table lists the models of the HP StorageWorks controller that are supported by the SAN Volume Controller. See the following Web site for the latest supported models. <http://www.ibm.com/servers/storage/support/virtual/2145.html>

Attention: The SAN Volume Controller 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 40. 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 HP StorageWorks firmware levels and the latest supported hardware.

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

Note: Concurrent upgrade of the controller firmware is also not supported with the SAN Volume Controller.

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.

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 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

Configuration interface for the HP StorageWorks

The Command Console configuration and service utility is the configuration interface for the HSG80 subsystem.

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 reference

“Connection settings for the HP StorageWorks” on page 312

The HP StorageWorks HSG80 includes options that are configurable at the connection level.

Sharing the HP StorageWorks controller between a host and the SAN Volume Controller

The HP StorageWorks controller can be shared between a host and a SAN Volume Controller.

- A host must not be connected to both a SAN Volume Controller and a HP StorageWorks HSG80 subsystem at the same time.
- Target ports must not be shared between a host and a SAN Volume Controller. Specifically, if a controller port is being used by a SAN Volume Controller it must not be present in a switch zone which enables a host to access the port.
- LU's and RAID arrays must not be shared between a host and a SAN Volume Controller.
- Partitions in the same container must all be either on the SAN Volume Controller or on hosts.

Related concepts

“Storage subsystems” on page 55

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

Switch zoning limitations for the HP StorageWorks subsystem

Consider the following limitations when planning switch zoning and connection to the SAN.

Attention: HSG80 based subsystems are supported with a single controller or dual controllers in the subsystem. Since the SAN Volume Controller 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

For code version 1.1.1, regardless of whether the subsystem uses one or two HSG80 controllers, only a single fibre-channel (FC) port attached to the subsystem can be present in a switch zone that contains SAN Volume Controller FC ports. This guarantees that the nodes in the cluster can access at most one port on an HSG80 subsystem.

For code version 1.2.0, switches should be zoned so that ports on the HSG80 are in the switch zone which contains all ports of the SAN Volume Controller nodes.

Connecting to the SAN

Multiple ports from an HSG80 should be physically connected to the fibre-channel SAN in order to enable servicing of the subsystem. However, switch zoning must be used in the manner outlined in this topic.

Note: If the HPQ Command Console is not able to access an FC port on each of the controllers in a two-controller subsystem, there is a risk of an undetected single point of failure.

If you're using partitioned containers then there are restrictions on the way the HSG80 should be connected to the SAN Volume Controller nodes.

Related reference

“Switch zoning for the SAN Volume Controller” on page 65

Ensure that you are familiar with the constraints for zoning a switch.

Quorum disks on HP StorageWorks

Managed disks presented by the HP StorageWorks controllers will be chosen by the SAN Volume Controller as quorum disks.

The SAN Volume Controller uses a logical unit (LU) presented by an HSG80 controller as a quorum disk. It will provide a quorum disk even if the connection is by a single port, although this is not recommended. If you are connecting the HSG80 subsystem with a single fibre, then you should ensure that you have another subsystem on which to put your quorum disk. Use the command line **svctask setquorum** to move quorum disks to another subsystem.

Managed disks provided by HSG80 may be chosen by the SAN Volume Controller software as quorum disks and can be set as quorum disks using the command-line interface. This means that clusters attached only to the HSG80 controllers are supported.

Related concepts

Chapter 7, “Configuring and servicing storage subsystems,” on page 215

You must ensure that your disk controllers and switches are correctly configured to work with the SAN Volume Controller to avoid performance issues.

Related information

“Creating a quorum disk” on page 236

A quorum disk is used to resolve tie-break situations when the voting set of nodes disagree on the current state of the cluster.

Advanced functions for HP StorageWorks

Some advanced functions of the HP StorageWorks are not supported by the SAN Volume Controller.

Advanced copy functions

Advanced copy functions for HP StorageWorks systems (for example, SnapShot and RemoteCopy) are not supported for disks that are managed by the SAN Volume Controller because the copy function does not extend to the SAN Volume Controller cache.

Partitioning

HSG80 supports partitioning. A partition is a logical division of a container that is represented to the host as a logical unit (LU). A container can be a RAID array or a JBOD (just a bunch of disks). All container types are candidates for partitions. Any non-transportable disk or storage set can be divided into a maximum of 8 partitions.

The following restrictions apply to partitioning:

- Partitioned containers are fully supported if the HSG80 subsystem is connected to the SAN by a single port.
- Partitioned containers are not configured by the SAN Volume Controller if the HSG80 subsystem is connected to the SAN by multiple ports.
- Partitioned containers are removed from the configuration if a single port connection becomes a multiport connection.
- Partitioned containers are configured if a multiport connection becomes a single port connection.

You must partition containers such that no spare capacity exists because there is no way to detect unused partitions. With a multiport connection, subsequent attempts to use this capacity removes all partitions on the container from the configuration.

Dynamic array expansion (LU expansion)

HSG80 does not provide dynamic array expansion.

Write protection of LUNs

Write protection of LUNs is not supported for use with the SAN Volume Controller.

HP StorageWorks advanced functions

VDisks that are created from MDisks presented by an HSG80 controller may be used in SAN Volume Controller FlashCopy mappings or SAN Volume Controller Metro Mirror relationships.

That is, SAN Volume Controller copy services fully supports the use of MDisks presented by an HSG80 controller.

Related concepts

“Metro Mirror” on page 47

Metro Mirror enables you to set up a relationship between two virtual disks (VDisks), so that updates that are made by an application to one VDisk are mirrored on the other VDisk.

Logical unit creation and deletion on the HP StorageWorks

Ensure you are familiar with the HSG80 container types for logical unit configuration.

Table 41 lists the valid container types.

Table 41. HSG80 container types for logical unit configuration

Container	Number of Members	Maximum Size
JBOD - non-transportable Attention: 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 concepts

“Storage subsystems” on page 23

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

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Configuring settings for the HP StorageWorks

The HP StorageWorks HSG80 configuration interface provides configuration settings and options 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 concepts

“Storage subsystems” on page 23

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

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Related reference

“Global settings for the HP StorageWorks”

Global settings apply across a HP StorageWorks HSG80 subsystem.

“Controller settings for the HP StorageWorks” on page 310

Controller settings are settings that apply across one HSG80 controller.

“Port settings for the HP StorageWorks” on page 310

Port settings are configurable at the port level.

“Logical unit settings for the HP StorageWorks” on page 311

Logical unit (LU) settings are configurable at the LU level.

“Connection settings for the HP StorageWorks” on page 312

The HP StorageWorks HSG80 includes options that are configurable at the connection level.

“Mapping and virtualization settings for the HP StorageWorks” on page 313

Consider these constraints for LUN mapping or masking and virtualization from the point of view of the HSG80 controller and their use in a SAN Volume Controller environment.

Related information

“Servicing storage subsystems” on page 237

Storage subsystems that are supported for attachment to the SAN Volume Controller are designed with redundant components and access paths to allow concurrent maintenance. Hosts have continuous access to their data during component failure and replacement.

Global settings for the HP StorageWorks

Global settings apply across a HP StorageWorks HSG80 subsystem.

Table 42. HSG80 global settings supported by the SAN Volume Controller

Option	HSG80 Default Setting	SAN Volume Controller Required Setting
DRIVE_ERROR_THRESHOLD	800	Default
FAILEDSET	Not defined	n/a

Related concepts

“Storage subsystems” on page 23

A *storage subsystem* is a device that coordinates and controls the operation of

one or more disk drives. A storage subsystem also synchronizes the operation of the drives with the operation of the system as a whole.

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Controller settings for the HP StorageWorks

Controller settings are settings that apply across one HSG80 controller.

Table 43 describes the options that can be set by HSG80 command-line interface (CLI) commands for each controller.

Table 43. HSG80 controller settings supported by the SAN Volume Controller

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
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 concepts

“Storage subsystems” on page 23

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

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Port settings for the HP StorageWorks

Port settings are configurable at the port level.

Only one port per HSG80 pair may be used with the SAN Volume Controller.

Note: The 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 44. HSG80 port settings supported by the

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, 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. The access rights can be checked using the "SHOW CONNECTIONS FULL" and any UNIT_OFFSETs using the "SHOW CONNECTIONS FULL" command.

Related concepts

"Storage subsystems" on page 23

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

"VDisk-to-host mapping" on page 33

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

Related tasks

"Configuring a balanced storage subsystem" on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Logical unit settings for the HP StorageWorks

Logical unit (LU) settings are configurable at the LU level.

Table 45 describes the options that must be set for each LU that is accessed by the SAN Volume Controller. LUs that are accessed by hosts can be configured differently.

Table 45. HSG80 LU settings supported by the SAN Volume Controller

Option	HSG80 Default Setting	SAN Volume Controller Required Setting
TRANSFER_RATE_REQUESTED	20MHZ	n/a
TRANSPORTABLE/ NOTTRANSPORTABLE	NOTTRANSPORTABLE	NOTTRANSPORTABLE

Table 45. HSG80 LU settings supported by the SAN Volume Controller (continued)

Option	HSG80 Default Setting	SAN Volume Controller Required Setting
ENABLE_ACCESS_PATH	ENABLE_ACCESS_PATH=ALL	ENABLE_ACCESS_PATH=ALL
DISABLE_ACCESS_PATH (See Note.)	NO DEFAULT	NO DEFAULT
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/ NOPREFERRED_PATH	NOPREFERRED_PATH is set	n/a
READ_CACHE/ NOREAD_CACHE	READ_CACHE	n/a
READAHEAD_CACHE/ NOREADAHEAD_CACHE	READAHEAD_CACHE	n/a
RUN/ NORUN	RUN	RUN
WRITE_LOG/NOWRITE_LOG	NOWRITE_LOG	NOWRITE_LOG
WRITE_PROTECT/ NOWRITE_PROTECT	NOWRITE_PROTECT	NOWRITE_PROTECT
WRITEBACK_CACHE/ NOWRITEBACK_CACHE	WRITEBACK_CACHE	WRITEBACK_CACHE
Note: DISABLE_ACCESS_PATH can be used to disable access from specific hosts. It should always be over ridden by using ENABLE_ACCESS_PATH=ALL on all connections to the SAN Volume Controller nodes.		

Related concepts

“Storage subsystems” on page 23

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

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Connection settings for the HP StorageWorks

The HP StorageWorks HSG80 includes options that are configurable at the connection level.

Table 46. HSG80 connection options and their required settings for the

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 concepts

“Storage subsystems” on page 23

A *storage subsystem* is a device that coordinates and controls the operation of

one or more disk drives. A storage subsystem also synchronizes the operation of the drives with the operation of the system as a whole.

Related tasks

“Configuring a balanced storage subsystem” on page 222

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings are applied to the device.

Related reference

“Configuration interface for the HP StorageWorks” on page 305

The Command Console configuration and service utility is the configuration interface for the HSG80 subsystem.

Mapping and virtualization settings for the HP StorageWorks

Consider these constraints for LUN mapping or masking and virtualization from the point of view of the HSG80 controller and their use in a SAN Volume Controller environment.

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 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 to restrict either the initiator ports or the Target ports which the SAN Volume Controller 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 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 is not supported.

Related reference

“Switch zoning for the SAN Volume Controller” on page 65

Ensure that you are familiar with the constraints for zoning a switch.

Configuring the HPQ Enterprise Virtual Array subsystem

This section provides information about configuring the HPQ Enterprise Virtual Array (EVA) disk controller system so it can attach to a SAN Volume Controller.

Related reference

“Supported models of the HPQ EVA”

The following table lists the models of the HPQ EVA that are supported by the SAN Volume Controller. See the following Web site for the latest supported models. <http://www.ibm.com/servers/storage/support/virtual/2145.html>

“Supported firmware levels for HPQ EVA” on page 315

The SAN Volume Controller supports HPQ EVA.

“User interface on HPQ EVA” on page 315

Ensure that you are familiar with the user interface that supports the HPQ EVA subsystem.

Sharing the HPQ EVA controller between a host and the SAN Volume Controller
The HPQ EVA controller can be shared between a host and a SAN Volume Controller.

Switch zoning limitations for the HPQ EVA subsystem

Consider the following limitations when planning switch zoning and connection to the SAN.

Quorum disks on HPQ EVA

Managed disks presented by the HPQ EVA controllers will be chosen by the SAN Volume Controller as quorum disks.

Advanced functions for HPQ EVA

Advanced copy functions for HPQ EVA (for example, VSnap and SnapClone) are not supported for disks that are managed by the SAN Volume Controller because the copy function does not extend to the SAN Volume Controller cache.

Logical unit configuration on the HPQ EVA

An EVA logical unit is referred to as a virtual disk (VDisk). An EVA subsystem can support up to 512 VDIsks. VDIsks are created within a set of physical disk drives, referred to as a disk group. A VDisk is striped across all the drives in the group.

Logical unit presentation

A virtual disk (VDisk) must be explicitly presented to a host before it can be used for I/O operations.

Configuration interface for the HPQ EVA

The HPQ EVA is configured, managed and monitored through a Storage Management Appliance. The Storage Management Appliance is a server that runs a software agent called Command View EVA. The Command View EVA is accessed using a graphical user interface that is provided by a standard Web browser.

Configuring settings for the HPQ EVA

The HPQ EVA configuration interface provides configuration settings and options that are supported with the SAN Volume Controller.

Supported models of the HPQ EVA

The following table lists the models of the HPQ EVA that are supported by the SAN Volume Controller. See the following Web site for the latest supported models. <http://www.ibm.com/servers/storage/support/virtual/2145.html>

Table 47. Supported HPQ EVA models

Model
EVA 3000
EVA 5000

Related information

“Configuring the HPQ Enterprise Virtual Array subsystem” on page 314
This section provides information about configuring the HPQ Enterprise Virtual Array (EVA) disk controller system so it can attach to a SAN Volume Controller.

Supported firmware levels for HPQ EVA

The SAN Volume Controller supports HPQ EVA.

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

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

Related information

“Configuring the HPQ Enterprise Virtual Array subsystem” on page 314
This section provides information about configuring the HPQ Enterprise Virtual Array (EVA) disk controller system so it can attach to a SAN Volume Controller.

Concurrent maintenance on the HPQ EVA

Concurrent maintenance is the capability to perform I/O operations to an HPQ EVA while simultaneously performing maintenance operations on it.

Important: All maintenance operations must be performed by an HPQ Field Engineer.

The SAN Volume Controller and HPQ EVA support concurrent hardware maintenance and firmware upgrade.

User interface on HPQ EVA

Ensure that you are familiar with the user interface that supports the HPQ EVA subsystem.

Storage Management Appliance

HPQ EVA systems are configured, managed and monitored through a Storage Management Appliance. The Storage Management Appliance is a PC server that runs a software agent called Command View EVA. The software agent is accessed using a user interface that is provided by a standard Web browser.

Command View EVA communicates in-band with the HSV controllers.

Related information

“Configuring the HPQ Enterprise Virtual Array subsystem” on page 314
This section provides information about configuring the HPQ Enterprise Virtual Array (EVA) disk controller system so it can attach to a SAN Volume Controller.

Sharing the HPQ EVA controller between a host and the SAN Volume Controller

The HPQ EVA controller can be shared between a host and a SAN Volume Controller.

- A host must not be connected to both a SAN Volume Controller and an HPQ EVA subsystem at the same time.
- LU's and RAID arrays must not be shared between a host and a SAN Volume Controller.

Related information

“Configuring the HPQ Enterprise Virtual Array subsystem” on page 314

This section provides information about configuring the HPQ Enterprise Virtual Array (EVA) disk controller system so it can attach to a SAN Volume Controller.

Switch zoning limitations for the HPQ EVA subsystem

Consider the following limitations when planning switch zoning and connection to the SAN.

Fabric zoning

The SAN Volume Controller switch zone must include at least one target port from each HSV controller in order to have no single point of failure.

Related information

“Configuring the HPQ Enterprise Virtual Array subsystem” on page 314

This section provides information about configuring the HPQ Enterprise Virtual Array (EVA) disk controller system so it can attach to a SAN Volume Controller.

Quorum disks on HPQ EVA

Managed disks presented by the HPQ EVA controllers will be chosen by the SAN Volume Controller as quorum disks.

Related information

“Configuring the HPQ Enterprise Virtual Array subsystem” on page 314

This section provides information about configuring the HPQ Enterprise Virtual Array (EVA) disk controller system so it can attach to a SAN Volume Controller.

Advanced functions for HPQ EVA

Advanced copy functions for HPQ EVA (for example, VSnap and SnapClone) are not supported for disks that are managed by the SAN Volume Controller because the copy function does not extend to the SAN Volume Controller cache.

Related information

“Configuring the HPQ Enterprise Virtual Array subsystem” on page 314

This section provides information about configuring the HPQ Enterprise Virtual Array (EVA) disk controller system so it can attach to a SAN Volume Controller.

Logical unit configuration on the HPQ EVA

An EVA logical unit is referred to as a virtual disk (VDisk). An EVA subsystem can support up to 512 VDIs. VDIs are created within a set of physical disk drives, referred to as a disk group. A VDisk is striped across all the drives in the group.

The minimum size of a disk group is eight physical drives. The maximum size of a disk group is all available disk drives.

EVA VDisks are created and deleted using the Command View EVA utility.

Note: A VDisk is formatted during the creation process, therefore, the capacity of the VDisk will determine the length of time it takes to be created and formatted. Ensure that you wait until the VDisk is created before you present it to the SAN Volume Controller.

A single VDisk can consume the entire disk group capacity or the disk group may be used for multiple VDisks. The amount of disk group capacity consumed by a VDisk depends on the VDisk capacity and the selected redundancy level. There are three redundancy levels:

- Vraid 1 - High redundancy (mirroring)
- Vraid 5 - Moderate redundancy (parity striping)
- Vraid 0 - No redundancy (striping)

Related reference

Logical unit creation and deletion on the HPQ EVA
EVA VDisks are created and deleted using the Command View EVA utility.

Related information

“Configuring the HPQ Enterprise Virtual Array subsystem” on page 314
This section provides information about configuring the HPQ Enterprise Virtual Array (EVA) disk controller system so it can attach to a SAN Volume Controller.

Logical unit creation and deletion on the HPQ EVA

EVA VDisks are created and deleted using the Command View EVA utility.

VDisks are formatted during creation. The time it takes to format the VDisks depends on the capacity.

Note: Selecting a host for presentation at creation time is not recommended. Ensure that you wait until the VDisk has been created before presenting it to the SAN Volume Controller.

Logical unit presentation

A virtual disk (VDisk) must be explicitly presented to a host before it can be used for I/O operations.

The SAN Volume Controller supports LUN masking on an HPQ EVA controller. When presenting a VDisk, the LUN can be specified or allowed to default to the next available value.

The SAN Volume Controller supports LUN virtualization on an HPQ EVA controller. The LUN-host relationship is set on a per-host basis.

Note: All nodes and ports in the SAN Volume Controller cluster must be represented as one host to the HPQ EVA.

Special LUs

The Console LU is a special VDisk that represents the SCSI target device. It is presented to all hosts as LUN 0.

Related information

“Configuring the HPQ Enterprise Virtual Array subsystem” on page 314
 This section provides information about configuring the HPQ Enterprise Virtual Array (EVA) disk controller system so it can attach to a SAN Volume Controller.

Configuration interface for the HPQ EVA

The HPQ EVA is configured, managed and monitored through a Storage Management Appliance. The Storage Management Appliance is a server that runs a software agent called Command View EVA. The Command View EVA is accessed using a graphical user interface that is provided by a standard Web browser.

In band

The Command View EVA subsystem communicates in-band with the HSV controllers.

Related information

“Configuring the HPQ Enterprise Virtual Array subsystem” on page 314
 This section provides information about configuring the HPQ Enterprise Virtual Array (EVA) disk controller system so it can attach to a SAN Volume Controller.

Configuring settings for the HPQ EVA

The HPQ EVA configuration interface provides configuration settings and options that are supported with the SAN Volume Controller.

The settings and options can have a scope of the following:

- Subsystem (global)
- Logical unit (LU)
- Host

Related reference

Global settings for the HPQ EVA

Global settings apply across a HPQ EVA subsystem.

Logical unit settings for the HPQ EVA

Logical unit (LU) settings are configurable at the LU level.

Host settings for the HPQ EVA

Host settings are configurable for the HPQ EVA.

Related information

“Configuring the HPQ Enterprise Virtual Array subsystem” on page 314
 This section provides information about configuring the HPQ Enterprise Virtual Array (EVA) disk controller system so it can attach to a SAN Volume Controller.

Global settings for the HPQ EVA

Global settings apply across a HPQ EVA subsystem.

Table 48 lists the subsystem options that can be accessed using the Command View EVA.

Table 48. HPQ EVA global settings supported by the SAN Volume Controller

Option	HPQ EVA default setting	SAN Volume Controller required setting
Console LUN ID	0	Any
Disk replacement delay	1	Any

Logical unit settings for the HPQ EVA

Logical unit (LU) settings are configurable at the LU level.

Table 49 describes the options that must be set for each LU that is accessed by other hosts. LUs that are accessed by hosts can be configured differently.

Table 49. HPQ EVA LU settings supported by the SAN Volume Controller

Option	HPQ EVA Default Setting	SAN Volume Controller Required Setting
Capacity	None	Any
Write cache	Mirrored Write-back	Mirrored
Read cache	On	On
Redundancy	Vraid0	Any
Preferred path	No preference	No preference
Write protect	Off	Off

Host settings for the HPQ EVA

Host settings are configurable for the HPQ EVA.

Table 50 lists the host options that can be accessed using the Command View EVA.

Table 50. HPQ EVA host settings supported by the SAN Volume Controller

Option	HPQ EVA Default Setting	SAN Volume Controller Required Setting
OS type	-	Windows
Direct eventing	Disabled	Disabled

Chapter 8. Installing or upgrading the SAN Volume Controller Console overview

The SAN Volume Controller Console can be installed on a Windows 2000 Server or Windows 2003 Server operating system.

Before you install the SAN Volume Controller Console, you must know how to administer commands on a Windows 2000 Server or a Windows 2003 Server operating system.

Modes of installation

You can use one of two modes to install or upgrade the SAN Volume Controller Console: *graphical* or *unattended*. Graphical mode requires that you are present for the installation process, while the unattended mode does not require you to be present.

Review the following list of installation and configuration tasks *before* you install or upgrade the SAN Volume Controller Console:

1. Check the hardware and software requirements.
2. Download the SAN Volume Controller Console software package from the following IBM Web site:
<http://www.ibm.com/servers/storage/support/virtual/2145.html>
3. If the secure shell (SSH) client software called PuTTY is not installed on your system, you must install the SSH client software. You can get more information about PuTTY and download the program from the following PuTTY Web site:
<http://www.chiark.greenend.org.uk/~sgtatham/putty/>
4. Install or upgrade the SAN Volume Controller Console either in graphical mode with the help of an installation wizard or in unattended (silent) mode.
5. Verify that the following services that are associated with the SAN Volume Controller Console are installed and started:
 - Service Location Protocol
 - IBM CIM Object Manager - SVC
 - IBM Websphere Application Server V5 - SVC
6. Use a Web browser to access the SAN Volume Controller Console.
7. Identify the clusters that are to be managed by the SAN Volume Controller Console.
8. Remove the SAN Volume Controller Console. You do not have to remove the SAN Volume Controller Console unless errors are generated during the installation process.

Related tasks

Installing or upgrading the SAN Volume Controller Console in graphical mode
If you choose to install or upgrade the SAN Volume Controller Console in graphical mode, you must satisfy all installation requirements before you start the installation.

Installing or upgrading the SAN Volume Controller Console in unattended (silent) mode

The unattended (silent) mode install or upgrade option enables you to run the installation unattended.

Verifying the Windows services associated with the SAN Volume Controller Console

You must verify that the Windows services associated with your SAN Volume Controller Console are correctly installed and started.

Post installation tasks

Complete the steps in the following sections to start using the SAN Volume Controller Console.

Removing the SAN Volume Controller Console

You can remove the SAN Volume Controller Console from your Windows system.

Related information

SAN Volume Controller Console hardware installation requirements

Before starting the installation, ensure that your system satisfies the following hardware installation prerequisites for installing the SAN Volume Controller Console.

SAN Volume Controller Console workstation space requirements

Before you start the installation, ensure that your system satisfies the workstation space prerequisites for installing the SAN Volume Controller Console.

SAN Volume Controller Console software installation requirements

Before you start the installation, ensure that your system satisfies the software installation prerequisites for installing the SAN Volume Controller Console.

SAN Volume Controller Console hardware installation requirements

Before starting the installation, ensure that your system satisfies the following hardware installation prerequisites for installing the SAN Volume Controller Console.

The following hardware is required:

- Any Intel-based PC running Windows 2000 Server SP 3 or Windows Server 2003
- Intel® Pentium® processor at 1 GHz, or faster
- Support for a communications adapter
- Minimum 2 GB RAM

SAN Volume Controller Console workstation space requirements

Before you start the installation, ensure that your system satisfies the workstation space prerequisites for installing the SAN Volume Controller Console.

The following amount of space is required on your workstation:

- 350 MB of disk space

Note: You might need to increase the total available disk space on your hard drives if the SAN Volume Controller Console and other associated products are split between more than one logical drive. Also, the SAN Volume Controller Console 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

SAN Volume Controller Console software installation requirements

Before you start the installation, ensure that your system satisfies the software installation prerequisites for installing the SAN Volume Controller Console.

The following software is required:

- Operating system:
 - Windows 2000 Server SP3 or Windows Server 2003 (Standard or Enterprise editions)
- 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 and download the program from the following PuTTY Web site:
<http://www.chiark.greenend.org.uk/~sgtatham/putty/>
- TCP/IP
- Adobe Acrobat Reader version 4.0 or later

You need the Adobe Acrobat Reader to read the License Agreement and product information from the SAN Volume Controller Console LaunchPad. You can download the Adobe Acrobat Reader from the following Web site:

<http://www.ibm.com/servers/storage/software/virtualization/svc>

Installing or upgrading the SAN Volume Controller Console in graphical mode

If you choose to install or upgrade the SAN Volume Controller Console in graphical mode, you must satisfy all installation requirements before you start the installation.

You must download the SAN Volume Controller Console zip file from the following IBM Web site:

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

After you have downloaded the zip file, you can extract the contents and write it to a CD or you can extract the contents to a directory on your system and perform the installation tasks from that directory.

The IBM TotalStorage SAN Volume Controller Console Launchpad application allows you to select from the following options:

SVC Console overview

Provides information about the SAN Volume Controller Console and its components.

Readme file

Provides any last minute product information that is not provided in the topics for installing the SAN Volume Controller Console.

Configuration guide

Provides instructions for installing and configuring the SAN Volume Controller Console.

License agreement

Provides information about the license for the SAN Volume Controller Console.

SAN Volume Controller Web site

Opens the SAN Volume Controller product Web site.

Installation wizard

Starts the SAN Volume Controller Console installation program.

Post installation tasks

Details information about validating the installation, accessing the SAN Volume Controller Console URL and adding the SAN Volume Controller Console cluster to the SAN Volume Controller Console management facility.

Exit Exits the SAN Volume Controller Console LaunchPad program.

The installation wizard determines if this is a reinstallation or an upgrade of the SAN Volume Controller Console. If the installation wizard determines that the SAN Volume Controller Console 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 that is 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.

Perform the following steps to install the SAN Volume Controller Console:

1. Log on to the system as a local system administrator.
2. Perform one of the following steps:
 - If you wrote the contents of the zip file to a CD and you have autorun mode set on your system, insert the CD into the drive. The IBM TotalStorage SAN Volume Controller Console Launchpad application starts.
 - If you wrote the contents of the zip file to a CD and you do not have autorun mode set on your system, insert the CD into the drive. Open a command prompt window and change to the W2K directory on the CD.
Issue the following command:
Launchpad
The IBM TotalStorage SAN Volume Controller Console Launchpad application panel is displayed.
 - If you did not write the contents of the zip file to a CD, open a command prompt window and change to the following directory:
extract_directory\W2K
Where *extract_directory* is the directory where you extracted the zip file.
Issue the following command:
Launchpad
The IBM TotalStorage SAN Volume Controller Console Launchpad application panel is displayed.
3. Click **Readme file** in the LaunchPad window to read installation information that is specific to this SAN Volume Controller Console software level.
4. Click **Installation wizard** in the LaunchPad window 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. You can click **Exit** if you want to close LaunchPad.

There might be a slight delay while the software loads on your system. After the software loads, a command prompt window opens to display the following message:

```
Initializing InstallShield Wizard...
Preparing Java <tm> Virtual Machine .....
.....
.....
```

The Welcome panel for the installation wizard is displayed. The Welcome panel provides the names of the documentation that you should read before you continue with the installation.

5. Click **Next** to continue or **Cancel** to exit the installation. If you click Next, the license agreement panel is displayed.
6. Read the license agreement information and perform one of the following steps:
 - Select **I accept the terms of the license agreement** and click **Next** to accept the license agreement.
 - Select **I do not accept the terms of the license agreement** and click **Cancel** to exit the installation.
7. Wait while the installation wizard verifies that your system meets all of the requirements. You might have to perform additional steps before the installation process can start if any of the following apply:
 - If you do not have PuTTY installed on your system, you must install PuTTY before you can continue with the installation. You can use the **putty-<version>-installer.exe** file that is located in the SSHClient/PuTTY folder that is included as part of the SAN Volume Controller Console zip file to install PuTTY on your system.
 - If you have a Service Location Protocol (SLP) service that is different from the SLP that the SAN Volume Controller Console requires, the installation wizard displays an error and asks you to stop the installation and remove this SLP service from the system.
 - If the SLP, the IBM CIM Object Manager (CIMOM), or WebSphere Application Server V5 - SVC services are started, you are asked if you want to continue the installation. If you choose to continue the installation, you must stop all the applications that use these services.

When the panel with the option to Preserve Configuration is displayed, you can choose to preserve the current configuration. If you chose to preserve the current configuration, the installation program skips the next steps and goes directly to the Installation Confirmation panel. If you do not preserve the current configuration, the Destination Directory panel is displayed.

8. Select one of the following options from the Destination Directory panel:
 - Click **Next** to accept the default directory.
 - Click **Browse** to select a different directory for installation and then click **Next** to continue the installation process.
 - Click **Cancel** to exit the installation process.

Note:

- The directory name, including the drive letter, can be a maximum of 44 characters.
- If the program detects insufficient space for the SAN Volume Controller Console 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 click **Back**, and select a different destination.

After you click **Next**, the PuTTY Configuration panel is displayed.

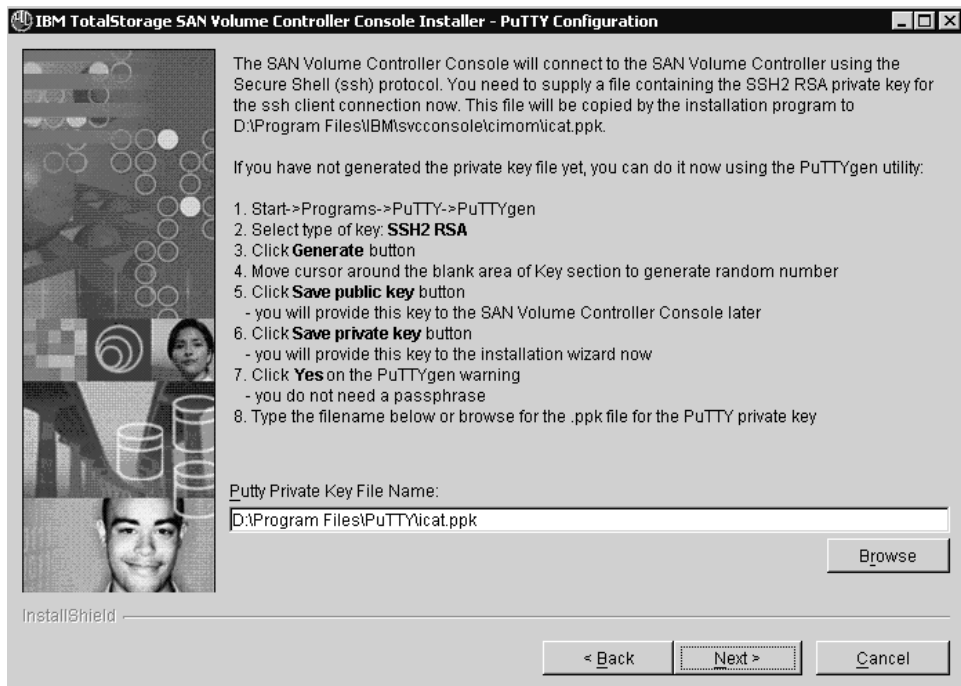


Figure 29. PuTTY Configuration panel

9. Enter the name and location of your PuTTY SSH2 RSA private key file or click **Browse** to select the private key file. If you do not have a PuTTY private key file, follow the steps that are displayed on the PuTTY configuration panel to generate a private and public key. Click **Next** to continue. The Updating Ports panel is displayed
10. Update the default port assignments and the default communication protocol by typing the unique port numbers and selecting the desired communication protocol for the products that have been registered on your system. To check the ports that are in use, issue the `netstat -a` command and view the `C:\WINNT\system32\drivers\etc\services` file. Click **Next** to continue. The Updating Embedded WAS Ports panel is displayed.

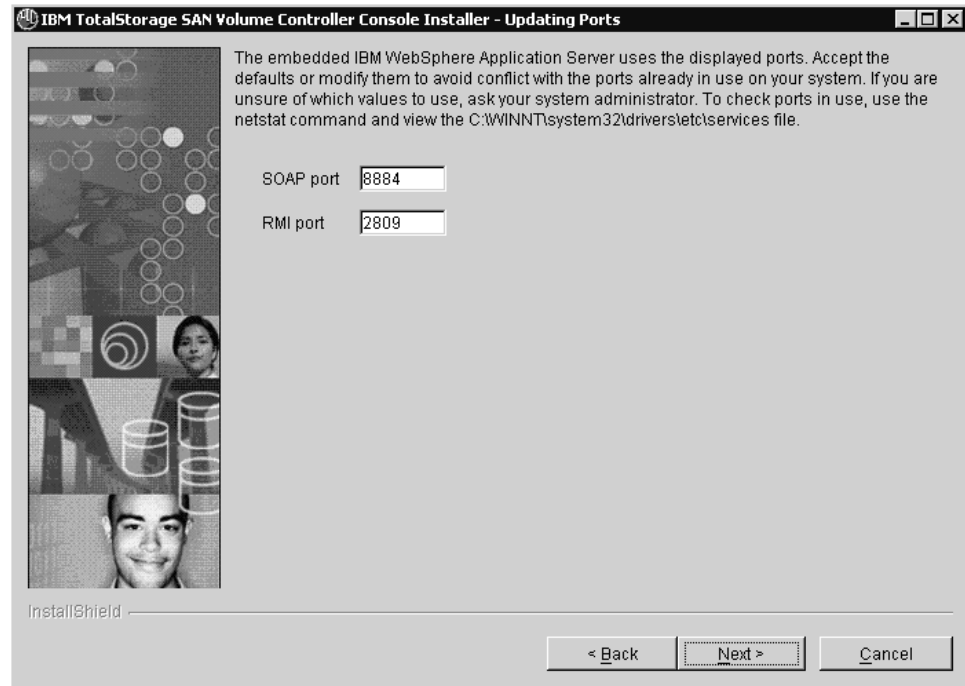


Figure 30. Updating Embedded WAS Ports panel

11. Update the default ports assignments by typing unique port numbers for the products that have been registered on your system. To check the ports that are in use, issue the `netstat -a` command and view the `C:\WINNT\system32\drivers\etc\services` file. Click **Next** to continue.
12. Click **Install** to confirm the installation location and file size and to start the installation. Click **Cancel** to exit the installation wizard or click **Back** to go to the previous panel. The Installation Progress panel indicates 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 and asks you to confirm the cancellation of the installation wizard. Click **Yes** to confirm the cancellation or click **No** to continue the installation. If you confirm the cancellation, the information that you entered or selected in the previous panel is not saved and you must restart the installation process.

After the completion of a successful installation of the SAN Volume Controller Console, the installer attempts to start the following services:

- Service Location Protocol
- IBM CIM Object Manager - SVC
- IBM WebSphere Application Server V5 - SVC

13. Review the log file for error messages when the Finish panel is displayed. The log file is located in `install_directory\logs\install.log`, where `install_directory` is the directory where the SAN Volume Controller Console was installed. The `install.log` file contains a trace of the installation process.

Note: At the bottom of the Finish panel, there is a check box labeled **View post installation tasks**. If you check this box and click **Finish**, the

wizard will exit and the post installation tasks text file is displayed. If you do not check this box, you can view the post installation tasks from the LaunchPad window.

14. Click **Finish** to exit the installation wizard.

Note: If the installation wizard determines that a system restart is necessary, you must restart your system. After you restart the system, the installation wizard continues with the installation.

15. If you did not review the post-installation tasks from the installation Finish panel, review the post installation tasks from the LaunchPad window.
 - a. Click **Post installation tasks** in the LaunchPad window.
 - b. Follow the instructions in this file to complete the post installation tasks for the SAN Volume Controller Console.
16. Click **Exit** to exit the LaunchPad window.
17. Use the Services component of the Windows Computer Management utility to verify that the following services Status is set to Started and Startup Type is set to Automatic:
 - Service Location Protocol
 - IBM CIM Object Manager - SVC
 - IBM WebSphere Application Server V5 - SVC

Installing or upgrading the SAN Volume Controller Console in unattended (silent) mode

The unattended (silent) mode install or upgrade option enables you to run the installation unattended.

You must satisfy all of the installation requirements before you start the SAN Volume Controller Console installation.

You must download the SAN Volume Controller Console zip file from the following IBM Web site:

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

After you have downloaded the zip file, you can extract the contents and write it to a CD or you can extract the contents to a directory on your system and perform the installation tasks from that directory.

Use the unattended method of installation to customize a response file. The response file is a template that is included in the zip file.

You can also create a standard response file to ensure that the product is installed consistently on multiple systems.

The installation wizard determines if this is a reinstallation or an upgrade of the SAN Volume Controller Console. If the installation wizard determines that the SAN Volume Controller Console 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 that is 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.

Perform the following steps to install the SAN Volume Controller Console:

1. Log on to the system as a local system administrator.
2. Perform one of the following steps:
 - If you wrote the contents of the zip file to a CD, insert the CD into the drive. If you have autorun mode set on your system, the SAN Volume Controller Console program starts within 30 seconds. When the IBM TotalStorage SAN Volume Controller Console Launchpad application begins, click **Exit**.
 - If you did not write the contents of the zip file to a CD, go the directory where you extracted the zip file.
3. Locate the response file that is in the W2K directory. If you wrote the contents of the zip file to a CD, you must copy and paste the response file to a location on your system.
4. Supply a file that contains the SSH2 RSA private key for the SSH client connection. The SAN Volume Controller Console connects to the SAN Volume Controller using the Secure Shell (SSH) protocol.

The installation program copies the SSH2 RSA private key to the following location:

inst_dir\cimom\icat.ppk

Where *inst_dir* is the location that you will install the SAN Volume Controller Console.

For example, if you use the default location, the file is copied in the following directory:

C:\ProgramFiles\IBM\svconsole\cimom\icat.ppk

- If you have not generated the private key file before, you can generate one now using the PuTTYgen utility. Perform the following steps to generate the private key using the PuTTYgen utility:
 - a. Click **Start** → **Programs** → **PuTTY** → **PuTTYgen**.
 - b. Select **SSH2 (RSA)** as the type of key to generate.
 - 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 must provide this key to the SAN Volume Controller Console in a later step.
 - f. Click **Save private key**. You must provide this key to the installation wizard using the option that is in the response file.
 - g. Click **Yes** on the PuTTYgen warning window. You have the option to set a passphrase for your key. You can enter a passphrase in the Key passphrase and Confirm passphrase fields. The passphrase is used to encrypt the key on the disk. If you choose to use a passphrase, you must enter the passphrase before you are allowed to use the key.

Attention: Do *not* enter anything in the Key passphrase and Confirm passphrase fields when you generate the key pair for the master console.
 - h. Use a text editor to set the value of the `<-W puttyConfiguration.puttyPrivateKeyFile>` option in the response file to the name of the file that contains the PuTTY private key.
5. Use a text editor to modify the default options in the response file with the values that you want to supply to the installation program:

- a. 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 ("").
- b. Use the following guidelines to activate the appropriate lines in the response file:

Note: If a response file line is active but inappropriate to the mode (new, reinstall or upgrade), it is 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 WebSphere Application Server - V5 SVC using the update ports variables options. If you want to change a specific port that is used for a particular WebSphere service, remove the # character from the beginning of the line that contains the value of the option and set it to the desired value. The following options are available for the embedded WebSphere ports:
 - `<-W ports.portSOAP="8884">`
 - `<-W ports.portRMI="2809">`
 - `<-W ports.portHTTP="9080">`
 - `<-W ports.portHTTPS="9443">`
- Change the default port values and default server communication type for the IBM CIM Object Manager (CIMOM) server using the following variables options:

Note: 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 desired value.

- `<-W cimObjectManagerPorts.port="5999">`
- `<-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. Do *not* save the response file with a .txt extension.

Reinstallation or Upgrade:

- The `<-G startUpgrade>` option must be enabled to permit the new SAN Volume Controller Console to be reinstalled (having the same version) or upgraded (installed at a higher version). To enable this option, remove the # character from the corresponding line and change the value of the option to yes.
- The `<-G stopProcessesResponse>` option tells the install program if it must automatically stop SLP, CIMOM, and WebSphere Application Server - V5 SAN Volume Controller 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 CIMOM, remove the # character from the corresponding line and change its value to yes.

- The `<-G saveConfiguration>` option specifies if the configuration files must be saved 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 must make the following lines active or accept the default values:
 - Change the default ports values for the embedded WebSphere Application Server - V5 SAN Volume Controller 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 that contains the value of the option and set it to the desired value. The following options are available for the embedded WebSphere ports:
 - `<-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 CIMOM server using the following variable options:

Note: If you want to change a specific port or the default server communication type, remove the # character from the beginning of the line that contains the option's value and set it to the desired value.

- `<-W cimObjectManagerPorts.port="5999">`
- `<-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. Do *not* save the response file with a .txt extension.

6. Issue the following command in a command prompt window to start the installation:

```
Directory\W2K\install -options response file path\responsefile
```

Where *Directory* is the directory of your CD drive or the directory where you extracted the zip file and *response file path* is the directory of the response file that you copied or extracted in step 3 on page 329 and customized in step 5 on page 329.

Note: The directory name, including the drive letter, can be a maximum of 44 characters.

Example 1:

You extracted the zip file to a folder called SVCCEXTRACT on the C: drive of your system and left the response file in the W2K directory. Issue the following command to start the installation:

```
C:\SVCCEXTRACT\W2K>install -options responsefile
```

Example 2:

You are using a CD to install the SAN Volume Controller Console. Your CD drive is E: and you copied the response file to your C: drive. Issue the following command to start the installation:

```
E:\W2K\>install -options C:\responsefile
```

The following output is displayed during the installation:

```
C:\SVCCEXtract\W2K>install -options responsefile

Initializing InstallShield Wizard...
Preparing Java(tm) Virtual Machine...
.....
.....
.....
.....
```

When the install is complete, the command prompt returns.

7. Check the install.log file for errors. This file is initially created in the system temporary file under the subdirectory named, cimagent. After all the prerequisites checks are performed, the log file is copied to the <dest-path>\logs directory.

The following is an example of an install.log file:

```
(May 15, 2003 9:36:06 AM), This summary log is an overview of the
sequence of the installation of the IBM TotalStorage SAN Volume
Controller Console 1.0.0.12
(May 15, 2003 9:38:22 AM), IBM TotalStorage SAN Volume Controller
Console installation process started with the following install
parameters:
Target Directory: C:\Program Files\IBM\svconconsole
SOAP port: 8884
RMI port: 2809
(May 15, 2003 9:38:28 AM), Copying Service Location Protocol Files ...
(May 15, 2003 9:38:29 AM), Service Location Protocol successfully installed
(May 15, 2003 9:38:29 AM), Copying CIM Object Manager Files ...
(May 15, 2003 9:39:26 AM), The PuTTY private key successfully copied
into file C:\Program Files\IBM\svconconsole\cimom\icat.ppk
(May 15, 2003 9:39:51 AM), The file setupCmdLine.bat successfully updated.
(May 15, 2003 9:39:51 AM), Compile MOF files started ...
(May 15, 2003 9:40:06 AM), MOF files successfully compiled.
(May 15, 2003 9:40:06 AM), Generate a certificate store started ...
(May 15, 2003 9:40:19 AM), Certificate store called truststore
successfully generated.
(May 15, 2003 9:40:20 AM), IBM CIM Object Manager successfully installed
(May 15, 2003 9:40:20 AM), Installing embedded version of IBM WebSphere
Application Server ...
(May 15, 2003 9:41:42 AM), Websphere Application Server - SVC
successfully installed.
(May 15, 2003 9:43:20 AM), Copying SAN Volume Controller Console Ear Files...
(May 15, 2003 9:46:11 AM), The ICACONSOLE application successfully installed.
(May 15, 2003 9:47:24 AM), The SVCConsole application successfully installed.
(May 15, 2003 9:48:06 AM), The help application successfully installed.
(May 15, 2003 9:48:27 AM), The "C:\Program Files\IBM\svconconsole\console\
embeddedWAS\bin\expressPorts\UpdateExpressMultiPorts.bat" -soap 8884
-boot 2809 -remove" command updated successfully embedded WAS ports
in configuration files.
(May 15, 2003 9:48:27 AM), Command to be executed : net start cimomsvr
(May 15, 2003 9:48:49 AM), Command to be executed : net start
"IBMWAS5Service - SVC"
(May 15, 2003 9:50:15 AM), The following services started successfully:
Service Location Protocol
IBM CIM Object Manager
IBM WebSphere Application Server V5 - SVC
(May 15, 2003 9:50:15 AM), INSTSUCC: The IBM TotalStorage SAN Volume
Controller Console has been successfully installed.
```

8. Issue the `exit` command to close the command prompt window. After the completion of the successful installation of the SAN Volume Controller Console, the installer attempts to start the following services:
 - Service Location Protocol
 - IBM CIM Object Manager - SVC
 - IBM WebSphere Application Server V5 - SVC
9. Perform the following steps to view the post installation tasks:
 - a. Open a command prompt window.
 - b. Go to the W2K directory on your CD drive or the W2K directory that was created during the extract.
 - c. Issue the following command to open LaunchPad:
LaunchPad
 - d. Click **Post installation tasks** in the LaunchPad window.
 - e. Follow the instructions in this file to complete the post installation tasks for the SAN Volume Controller Console.
10. Use the Services component of the Windows Computer Management utility to verify that the following services Status is set to Started and Startup Type is set to Automatic:
 - Service Location Protocol
 - IBM CIM Object Manager - SVC
 - IBM WebSphere Application Server V5 - SVC

Verifying the Windows services associated with the SAN Volume Controller Console

You must verify that the Windows services associated with your SAN Volume Controller Console are correctly installed and started.

Perform the following steps to verify your SLP, CIM Object Manager (CIMOM), and WebSphere Application Server V5 - SVC services were correctly installed:

1. Verify the installation of the SLP.
 - a. Verify that the SLP is started. Select **Start** → **Settings** → **Control Panel**.
 - b. Double-click the **Administrative Tools** icon.
 - c. Double-click the **Services** icon.
 - d. Find **Service Location Protocol** in the **Services** list. For this component, the **Status** column should be marked Started.
 - e. If the SLP 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.
 - f. Do not close the Services window because you will also use it to verify the CIM Object Manager (CIMOM) service.
2. Verify the installation of the SAN Volume Controller Console.
 - a. Find **IBM CIM Object Manager - SVC** in the **Services** list. For this component, the **Status** column should be marked Started.
 - b. If the CIMOM is not started, right click on **IBM CIM Object Manager - SVC** 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 WebSphere Application Server V5 - SVC service.

3. Verify the installation of the WebSphere Application Server V5 - SVC service.
 - a. Find the WebSphere Application Server V5 - SVC in the **Services** list. For this component, the **Status** column should be marked Started.
 - b. If the **WebSphere Application Server V5 - SVC** service is not started, right click on **WebSphere Application Server V5 - SVC** 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.

Post installation tasks

Complete the steps in the following sections to start using the SAN Volume Controller Console.

After you have installed the SAN Volume Controller and the services (IBM CIM Object Manager, IBM WebSphere Application Server V5 - SVC, Service Location Protocol) have started, you will use a browser to access the Web pages of the Console for purposes of administering the SAN Volume Controller as well as configuring SAN Volume Controller clusters.

Each time you want to add a SAN Volume Controller cluster to the collection of clusters managed by the SAN Volume Controller, you must store the PuTTY secure shell (SSH) client public key located on the SAN Volume Controller system, on the SAN Volume Controller cluster.

Attention: If you do not store the SSH public key on the SAN Volume Controller cluster, the SAN Volume Controller Console software cannot connect to the cluster.

When you installed the SAN Volume Controller Console, you provided the name and location of the PuTTY SSH client private key. At the time you used PuTTYGen to generate the PuTTY SSH private key, you also generated an SSH public key. Familiarize yourself with the name and location of the PuTTY SSH public key on the SAN Volume Controller Console system.

Note: This is a long-term administrative task and not just a post-installation task.

To enter config mode: `switch#config-t`

To enable ssh: `switch (config)#ssh server enable`

This document has an overview of the steps necessary to get to the web page where you identify the PuTTY public key to the clusters. These steps are documented in more detail in other sections of this manual and references are included to the relevant section titles.

1. Start your Web browser to access the SAN Volume Controller Console. It is recommended that you log on to the SAN Volume Controller Console system from a browser on which the SAN Volume Controller Console is installed to complete uploading the client public SSH key for each cluster that you want to manage. You can access the SAN Volume Controller Console by entering 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 on to the SAN Volume Controller Console 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 SAN Volume Controller Console 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 clusters to the SAN Volume Controller Console. The steps you might need to perform to add SAN Volume Controller clusters to the SAN Volume Controller Console collection of managed clusters, depends on the current status of the cluster in which you are interested.

Choose one of the following two steps, depending on whether the cluster has completed the cluster creation (initialization) process:

- a. Uninitialized SAN Volume Controller cluster.

If you have not yet created a SAN Volume Controller cluster using the front panel of the SAN Volume Controller cluster, you will need to perform that phase of the cluster creation first. You will be given a special password by the customer engineer (CE) to be used in later steps of initializing the SAN Volume Controller Console.

After you create the SAN Volume Controller cluster using the front panel of cluster, you will need to complete the creation of the cluster by using the SAN Volume Controller Console Web pages.

Enter the IP address of the cluster and check **Create (Initialize) Cluster**. When you click the **OK** button, the Create a Cluster wizard will take over and present you with the panels you need to complete initializing the cluster.

The browser will then prompt you to enter the network password. Enter the user name `admin` and the password provided to you by the customer engineer (CE) during the cluster front panel creation phase which is configured for the cluster.

During the initializing of the cluster, using the SAN Volume Controller Console, you will be taken to a Web page to provide the PuTTY SSH client public key to upload the key to the cluster. Step 5 on page 336 below continues with the SSH public key input description. This PuTTY SSH client public key is the other key of the key pair you provided to the SAN Volume Controller Console during the installation program.

- b. Previously initialized SAN Volume Controller cluster.

If the SAN Volume Controller cluster has completed the initialization (creation) process but is not yet registered with the SAN Volume Controller Console, click the **Add SAN Volume Controller Cluster** button and then add the cluster IP address but *do not* check **Create (Initialize) Cluster**, which is above the **OK** button. When you click the **OK** button, you will be

taken to the Web page to provide the PuTTY SSH client public key to upload to the cluster. Step 5 below continues with the SSH key input description.

The browser will then prompt you to enter the network password. Enter the user name `admin` and the password which is configured for the cluster. Then Click **OK**.

5. Store the SAN Volume Controller Console system SSH public key on the SAN Volume Controller Console. This PuTTY client SSH public key is the other key in the key pair you provided to the SAN Volume Controller Console during the installation program. Each key is associated with an ID string that you define that can consist of up to 30 characters. Up to 100 keys can be stored on a cluster. You can add keys to provide either *administrator* access or *service* access. Perform the following steps to store the SSH public key on the cluster:
 - a. Enter the SSH public key name and directory location on your local browser system in the field labeled **Public Key (file upload)** or click **Browse** to identify the key on the local system. Alternatively, you can paste the SSH key into the **Public Key (direct input)** field.
 - b. Enter an ID string in the field labeled **ID**. This is a unique ID to distinguish the key and is not related to a user name.
 - c. Select the *administrator* **Access Level** radio button.
 - d. Click **Add Key** to store this SSH public key on the cluster.
6. Launch the secondary Web browser window to manage your specific cluster. After you have identified the SAN Volume Controller clusters to the SAN Volume Controller Console you can see a summary of all of the clusters. 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. A new view will be 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 Web application. Now you can work with the specific SAN Volume Controller cluster which you selected.

Note: The `ClusterName` parameter in the browser location URL, identifies the cluster with which you are working.

For example:

```
http://9.43.147.38:9080/svc/Console?Console.login  
Token=79334064:f46d035f31:-7ff1&Console.  
ClusterName=9.43.225.208
```

Select **Manage Cluster** and click **View Cluster Properties** in the portfolio section.

This completes the verification of the connection to the SAN Volume Controller.

Removing the SAN Volume Controller Console

You can remove the SAN Volume Controller Console from your Windows system.

1. Log onto the system as a local system administrator.

2. Stop the CIM Object Manager (CIMOM), WebSphere Application Server V5 - SVC, and the SLP services if they are started.
 - a. Click **Start** → **Settings** → **Control Panel**.
 - b. In the Control Panel window, double-click on the **Administrative Tools** icon.
 - c. Double-click the **Services** icon. The Services window opens.
 - d. Stop the CIMOM service:
 - 1) In the Services window, find IBM CIM Object Manager (CIMOM). Click on the service to select it.
 - 2) If the **Status** column shows Started, right-click the service, then click **Stop** on the menu.
 - e. Stop the WebSphere Application Server V5 - SVC service:
 - 1) In the Services window, find IBM WebSphere Application Server V5 - SVC. 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.
 - f. Stop the SLP service:

Note: You must be careful if you have other applications that use the SLP service. In this case, you must stop these applications before stopping SLP service, because during the removal process the SLP service will be deleted. You must also stop the configuration utilities for the SAN Volume Controller Console, if they are running.

- 1) In the Services window, find 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.

Note: If you did not stop the CIMOM service, the system now asks if you want to stop it. Because the CIMOM service is dependent on the SLP service which you just stopped, you must click **Yes** to stop the CIMOM.

- 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 SAN Volume Controller Console and the SLP components.
 - a. From the Windows menu bar, click **Start** → **Settings** → **Control Panel**. Double-click **Add/Remove Programs**.
 - b. Click **IBM TotalStorage SAN Volume Controller Console** from the list of currently installed programs and click **Remove** to remove the product. The Welcome panel for the Uninstaller opens.
 4. Click **Next** to continue or click **Cancel** to stop the removal of the SAN Volume Controller Console. The program detects whether the SLP, CIMOM, and the WebSphere Application Server V5 - SVC 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 SAN Volume Controller Console are dependent on the services.

5. Click **Next** to have the program stop the services for you or 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 337. You must then restart the removal process from the Windows Add/Remove facility. The Confirmation panel opens.
6. Click **Remove** to continue or click **Cancel** to stop the removal of the SAN Volume Controller Console. Click **Back** to return to the previous panel. The Uninstallation Progress panel opens.
7. Wait for the program to remove the SAN Volume Controller Console product. The Finish panel for the Uninstaller opens.
8. 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.
10. If the system has not been restarted since SAN Volume Controller Console was removed, do so now.
11. Log on to the system as a local system administrator.

Note: The removal process saves files uniquely related to the configuration in a backup directory under the destination path where you installed the SAN Volume Controller Console. 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\svccconsole.

12. Empty your Windows Recycle Bin to reclaim the disk space that was made available during the removal process.

Chapter 9. IBM TotalStorage support for Microsoft Volume Shadow Copy Service

The SAN Volume Controller provides support for the Microsoft Volume Shadow Copy Service. The Microsoft Volume Shadow Copy Service can provide a point-in-time (shadow) copy of a Windows host volume while the volume is mounted and files are in use.

The following components are used to provide support for the service:

- SAN Volume Controller
- SAN Volume Controller master console
- IBM TotalStorage hardware provider, known as the IBM TotalStorage Support for Microsoft Volume Shadow Copy Service
- Microsoft Volume Shadow Copy Service

The IBM TotalStorage hardware provider is installed on the Windows host.

To provide the point-in-time shadow copy, the components complete the following process:

1. A backup application on the Windows host initiates a snapshot backup.
2. The Volume Shadow Copy Service notifies the IBM TotalStorage hardware provider that a copy is needed.
3. The SAN Volume Controller prepares the volumes for a snapshot.
4. The Volume Shadow Copy Service quiesces the software applications that are writing data on the host and flushes file system buffers to prepare for the copy.
5. The SAN Volume Controller creates the shadow copy using the FlashCopy Copy Service.
6. The Volume Shadow Copy Service notifies the writing applications that I/O operations can resume, and notifies the backup application that the backup was successful.

The Volume Shadow Copy Service maintains a free pool of virtual disks (VDisks) for use as a FlashCopy target and a reserved pool of VDisks. These pools are implemented as virtual host systems on the SAN Volume Controller.

Installation overview

The steps for implementing the IBM TotalStorage Support for Microsoft Volume Shadow Copy Service must be completed in the correct sequence.

Before you begin, you must have experience with or knowledge of administering a Windows operating system.

You must also have experience with or knowledge of administering a SAN Volume Controller.

Complete the following tasks:

1. Verify that the system requirements are met.
2. Install the SAN Volume Controller Console if it is not already installed.
3. Install the IBM TotalStorage hardware provider.
4. Verify the installation.

5. Create a free pool of volumes and a reserved pool of volumes on the SAN Volume Controller.

System requirements for the IBM TotalStorage hardware provider

Ensure that your system satisfies the following requirements before you install the IBM TotalStorage hardware provider on a Windows Server 2003 operating system.

The following software is required:

- SAN Volume Controller Console version 2.1.0 or later. You must install the SAN Volume Controller Console *before* you install the IBM TotalStorage hardware provider.
- SAN Volume Controller version 2.1.0 or later with the FlashCopy feature enabled.
- IBM TotalStorage Support for Microsoft Volume Shadow Copy Service version 2.3 or later.
- Windows Server 2003 operating system. The following editions of Windows Server 2003 are supported:
 - Standard Server Edition
 - Enterprise Edition, 32-bit version

Installing the IBM TotalStorage hardware provider

This section includes the steps to install the IBM TotalStorage hardware provider on a Windows server.

Important: You must install the SAN Volume Controller Console before you install the IBM TotalStorage hardware provider.

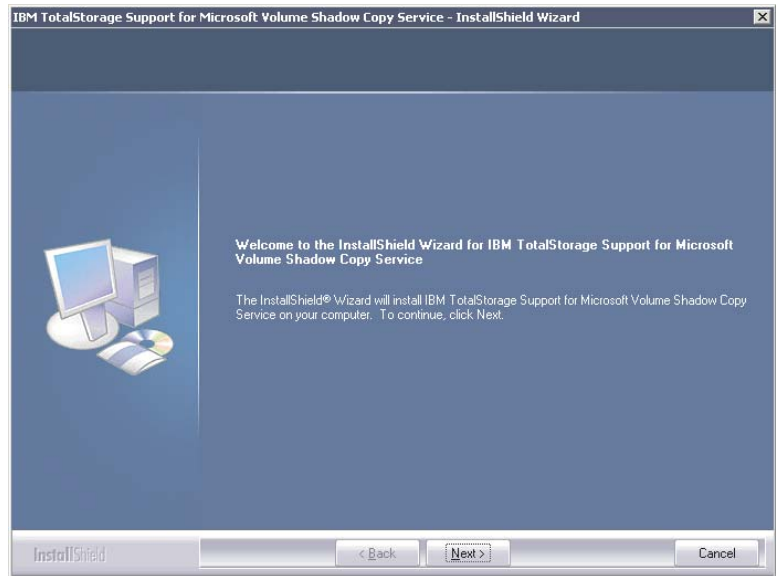
You must satisfy all of the prerequisites that are listed in the system requirements section before starting the installation.

During the installation, you are prompted to enter information about the SAN Volume Controller Console, including the location of the truststore file. You must copy this file to a location that is accessible to the IBM TotalStorage hardware provider on the Windows server.

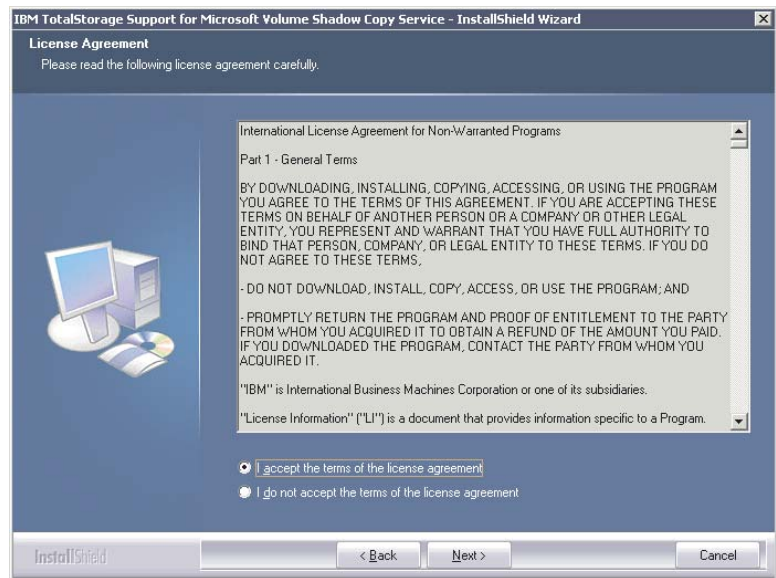
When the installation is complete, the installation program might prompt you to restart the system.

Perform the following steps to install the IBM TotalStorage hardware provider on the Windows server:


1. Log on to Windows as an administrator.
2. Download the IBM VSS Host Installation Package file from the following Web site:
<http://www.ibm.com/servers/storage/support/virtual/2145.html>
3. Double click on the name of the file that you downloaded in step 2 to start the installation process. The Welcome panel is displayed.



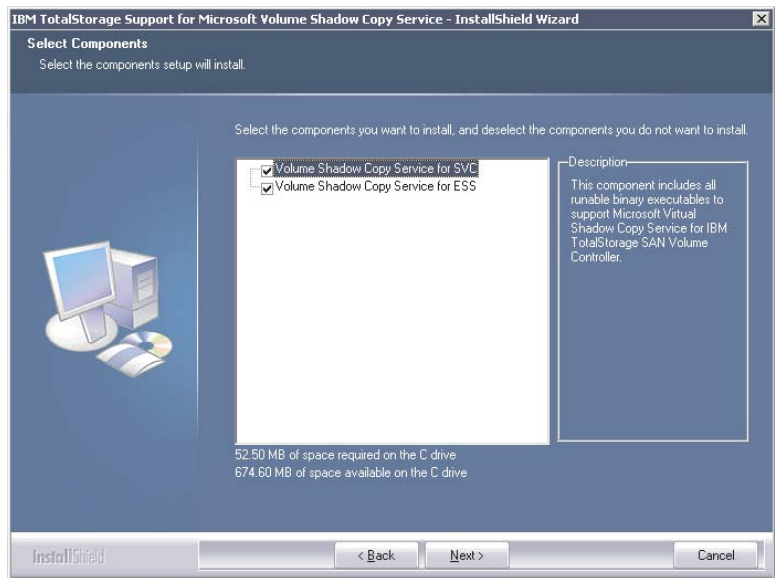
4. Click **Next** to continue with the InstallShield Wizard. The License Agreement panel is displayed. You can click **Cancel** at any time to exit the installation. To move back to previous screens while using the wizard, click **Back**.




5. Read the license agreement information. Select whether you accept the terms of the license agreement, and click **Next**. If you do not accept, you cannot continue with the installation. The Choose Destination Location panel is displayed.



6. Click **Next** to accept the default directory where the setup program will install the files, or click **Change** to select a different directory. Click **Next**. The Select Components panel is displayed.



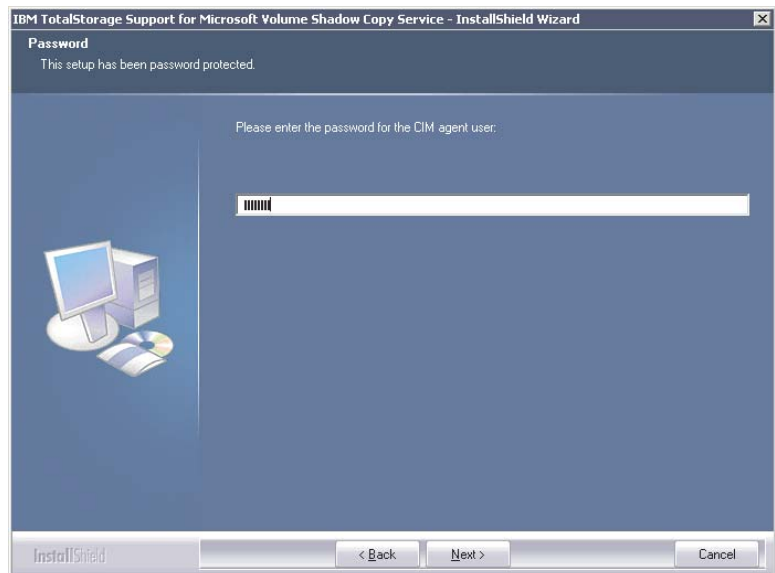
7. Select **Volume Shadow Copy Service for SVC**. Click **Next**. The Edit Data panel is displayed.




8. Perform the following steps from the Edit Data panel:

- In the **CIM Host** box, type the name of the server where the SAN Volume Controller Console is installed.
- In the **CIM Port** box, type the port number of the server where the SAN Volume Controller Console is installed. The default value is 5999.
- In the **CIM User** box, type the user name that the IBM TotalStorage hardware provider will use to gain access to the server where the SAN Volume Controller Console is installed.

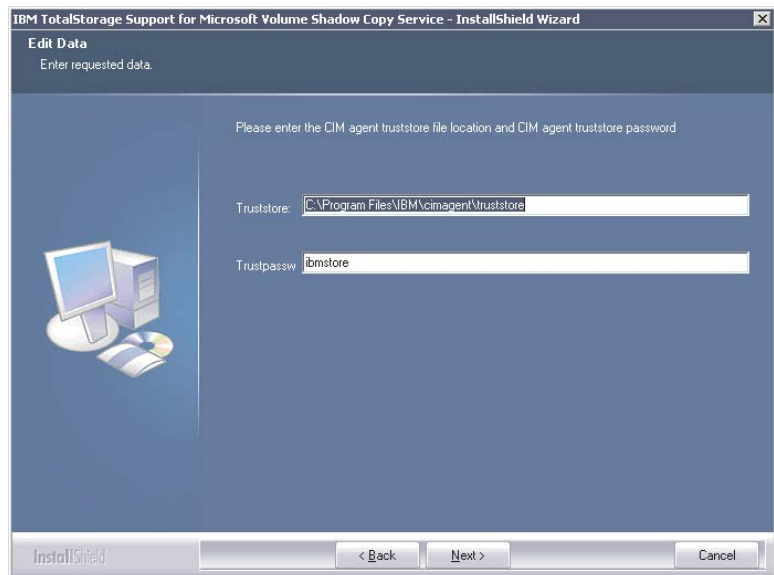
Click **Next**. The Password panel is displayed.



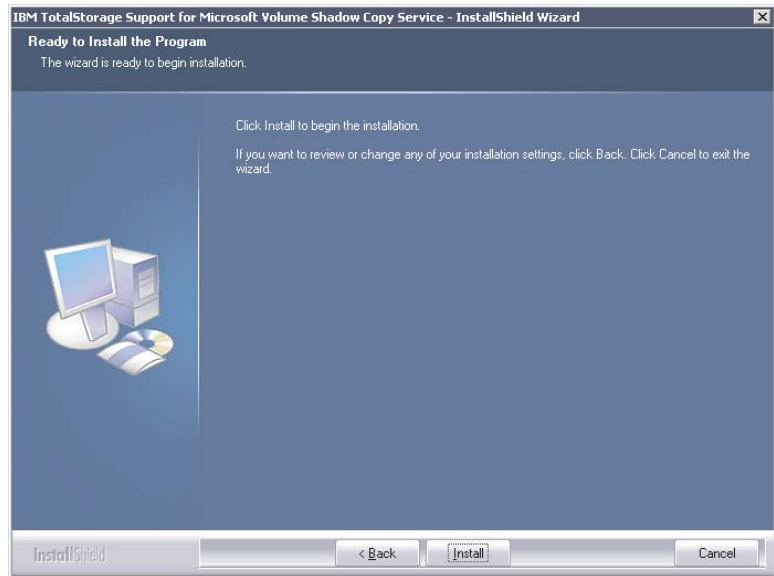
9. Enter the password for the user name that the IBM TotalStorage hardware provider will use to gain access to the SAN Volume Controller Console and click **Next**. The Setup Type panel is displayed.



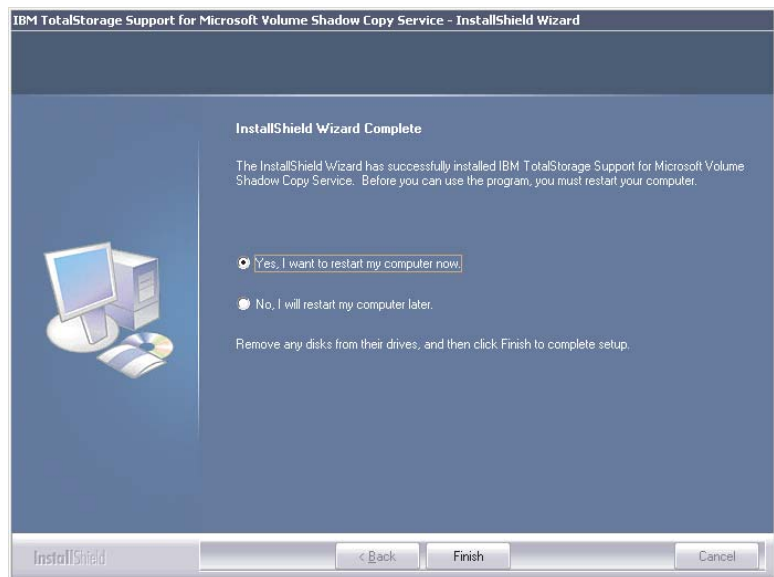
10. Select whether you want to use the Secure Sockets Layer (SSL) protocol to connect to the SAN Volume Controller Console. The default communication protocol for the SAN Volume Controller Console is SSL. Click **Next**. The Edit Data panel is displayed.



11. Perform the following steps from the Edit Data panel.
 - In the **Truststore** box, enter the path to the SAN Volume Controller Console truststore file.
 - In the **Trustpassword** box, enter the truststore password. The default truststore password is ibmstore.Click **Next**. The Ready to Install the Program panel is displayed.



12. Click **Install** to start the installation. The Setup Status panel is displayed.
13. The InstallShield Wizard Complete is displayed.



- If you want to restart the system now, click **Yes**, and then click **Finish** to restart the system.
- If you want to restart the system later, click **No**, and then click **Finish** to close the wizard.
- If you are not prompted to restart the system, click **Finish** to close the wizard.

Validating the truststore certificate expiration

In order to successfully log onto the master console, you must ensure that you have a valid truststore certificate.

When signing onto the master console, you might receive a message similar to the following:

CMMUI8304E The Administrative server is unable to find a valid certificate in the truststore file.

This message is displayed when a certificate in the truststore file expires. The Administrative server uses the certificates in the truststore file to create a secure connection with the CIM agent. Because the Administrative server cannot find a valid certificate for the CIM agent in the truststore file, no authentication can occur.

To resolve the problem, you must verify that the truststore file was created correctly. If you have any problems, contact your service representative.

Perform the following steps to regenerate a truststore certificate:

1. Go to the C:\Program Files\IBM\svconconsole\cimom directory.
2. Double-click on the **mkcertificate.bat** file. A "Generating Certificates" message is displayed. The new certificate is generated and stored in the C:\Program Files\IBM\svconconsole\cimom directory.
3. Copy the truststore file to the following sub directories:

Note: Each directory begins with C:\Program Files\IBM\svconconsole\console\embeddedWAS...

C:\...\config\cells\DefaultNode\applications\
ICAConsole.ear\deployments\ICAConsole\ICAConsole.war\
WEB-INF

C:\...\config\cells\DefaultNode\applications\
SVCConsole.ear\deployments\SVCConsole\SVCConsole.war\
WEB-INF

C:\...\config\installedApps\DefaultNode\
ICAConsole.ear\ICAConsole.war\WEB-INF

C:\...\config\installedApps\DefaultNode\
SVCConsole.ear\SVCConsole.war\WEB-INF

4. Stop and then restart the following applications. The following services are located in **Start ► Settings ► Control Panel ► Administrative Tools ► Component Services**.

- IBM CIM Object Manager
- IBM WebSphere Application Server V5 - SVC

To stop and then restart the services, right-click on the application and select **Stop**, then **Start**.

Note: If the stop command times-out in the IBM WebSphere application, you can restart the master console because this restarts the application, as well.

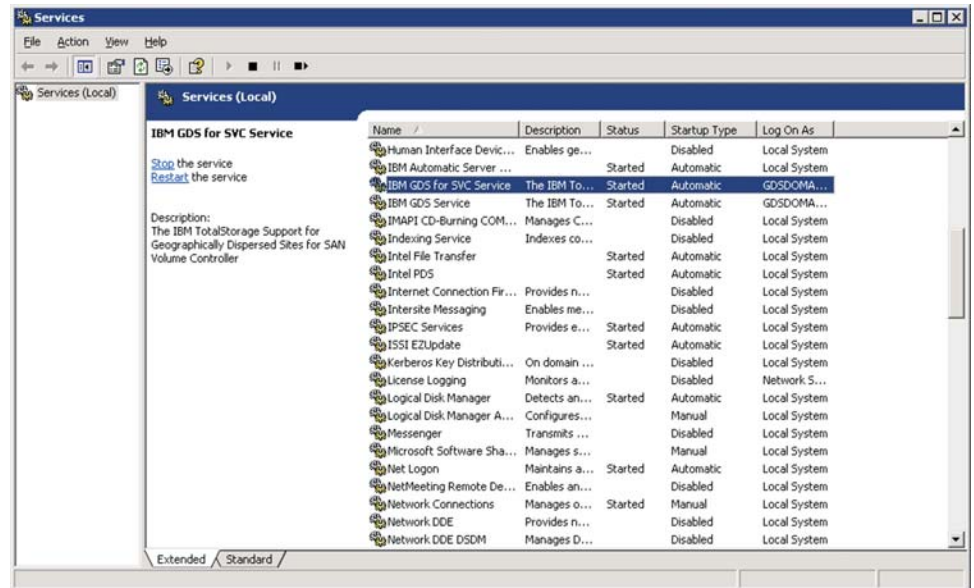
5. Ensure that both applications are running again. Launch the SAN Volume Controller Console and logon.

Verifying the installation

This task verifies that the IBM TotalStorage Support for Microsoft Volume Shadow Copy Service is correctly installed on the Windows server.

Perform the following steps to verify the installation:

1. Click **Start** → **All Programs** → **Administrative Tools** → **Services** from the Windows server task bar. The **Services** panel is displayed.



2. Ensure that the service named IBM TotalStorage Support for Microsoft Volume Shadow Copy Service appears and that **Status** is set to **Started** and **Startup Type** is set to **Automatic**.
3. Open a command prompt window and issue the following command:
vssadmin list providers
4. Ensure that the service named IBM TotalStorage Support for Microsoft Volume Shadow Copy Service is listed as a provider.

If you are able to successfully perform all of these verification tasks, the IBM TotalStorage hardware provider was successfully installed on the Windows server.

Creating the free and reserved pools of volumes

The IBM TotalStorage hardware provider maintains a free and a reserved pool of volumes. Because these objects do not exist on the SAN Volume Controller, the free and reserved pool of volumes are implemented as virtual host systems. You must define these two virtual host systems on the SAN Volume Controller.

When a shadow copy is created, the IBM TotalStorage hardware provider selects a volume in the free pool, assigns it to the reserved pool, and then removes it from the free pool. This protects the volume from being overwritten by other Volume Shadow Copy Service users.

To successfully perform a Volume Shadow Copy Service operation, there must be enough virtual disks (VDisks) mapped to the free pool. The VDisks must be the same size as the source VDisks.

Use the SAN Volume Controller Console or the SAN Volume Controller command-line interface (CLI) to perform the following steps:

1. Create a host for the free pool of VDisks.
 - You can use the default name VSS_FREE or specify a different name.

- Associate the host with the worldwide port name (WWPN) 5000000000000000 (15 zeroes).
2. Create a virtual host for the reserved pool of volumes.
 - You can use the default name VSS_RESERVED or specify a different name.
 - Associate the host with the WWPN 5000000000000001 (14 zeroes).
 3. Map the logical units (VDisks) to the free pool of volumes.

Restriction: The VDisks cannot be mapped to any other hosts.

- If you already have VDisks created for the free pool of volumes, you must assign the VDisks to the free pool.
4. Create VDisk-to-host mappings between the VDisks selected in step 3 and the VSS_FREE host to add the VDisks to the free pool. Alternatively, you can use the **ibmvcfg add** command to add VDisks to the free pool.
 5. Verify that the VDisks have been mapped.

If you do not use the default WWPNs 5000000000000000 and 5000000000000001, you must configure the IBM TotalStorage hardware provider with the WWPNs.

Changing the configuration parameters

You can change the parameters that you defined when you installed the IBM TotalStorage hardware provider. You must use the `ibmvcfg.exe` utility to change the parameters.

Table 51 describes the configuration commands that are provided by the `ibmvcfg.exe` utility.

Table 51. Configuration commands

Command	Description	Example
<code>ibmvcfg showcfg</code>	Lists the current settings.	<code>ibmvcfg showcfg</code>
<code>ibmvcfg set username <username></code>	Sets the user name to access the SAN Volume Controller master console.	<code>ibmvcfg set username johnny</code>
<code>ibmvcfg set password <password></code>	Sets the password of the user name that will access the master console.	<code>ibmvcfg set password mypassword</code>
<code>ibmvcfg set targetSVC <ipaddress></code>	Specifies the IP address of the SAN Volume Controller on which the VDisks are located when VDisks are moved to and from the free pool with the <code>ibmvcfg add</code> and <code>ibmvcfg rem</code> commands. The IP address is overridden if you use the <code>-s</code> flag with the <code>ibmvcfg add</code> and <code>ibmvcfg rem</code> commands.	<code>set targetSVC 64.157.185.191</code>
<code>set backgroundCopy</code>	Sets the background copy rate for FlashCopy.	<code>set backgroundCopy 80</code>
<code>ibmvcfg set trustpassword <trustpassword></code>	Sets the password for the truststore file. The default value is <code>ibmstore</code> .	<code>ibmvcfg set trustpassword ibmstore</code>

Table 51. Configuration commands (continued)

Command	Description	Example
ibmvcfg set truststore <path>	Specifies the truststore file location.	ibmvcfg set truststore c:\truststore
ibmvcfg set usingSSL	Specifies whether to use Secure Sockets Layer protocol to connect to the master console.	ibmvcfg set usingSSL yes
ibmvcfg set cimomPort <portnum>	Specifies the master console port number. The default value is 5999.	ibmvcfg set cimomPort 5999
ibmvcfg set cimomHost <server name>	Sets the name of the server where the master console is installed.	ibmvcfg set cimomHost cimomserver
ibmvcfg set namespace <namespace>	Specifies the namespace value that master console is using. The default value is \root\ibm.	ibmvcfg set namespace \root\ibm
ibmvcfg set vssFreeInitiator <WWPN>	Specifies the WWPN of the host. The default value is 5000000000000000. Modify this value only if there is a host already in your environment with a WWPN of 5000000000000000.	ibmvcfg set vssFreeInitiator 5000000000000000
ibmvcfg set vssReservedInitiator <WWPN>	Specifies the WWPN of the host. The default value is 5000000000000001. Modify this value only if there is a host already in your environment with a WWPN of 5000000000000001.	ibmvcfg set vssReservedInitiator 5000000000000001

Adding and removing volumes

You can use the `ibmvcfg.exe` utility to perform the pool management tasks of adding, removing, and listing volumes.

The Microsoft Volume Shadow Copy service maintains a free pool of volumes and a reserved pool of volumes. These pools are implemented as virtual host systems on the SAN Volume Controller.

Table 52 describes the `ibmvcfg.exe` commands for adding or removing volumes from the free pool of volumes.

Table 52. Pool management commands

Command	Description	Example
ibmvcfg listvols	Lists all virtual disks (VDisks), including information about size, location, and VDisk to host mappings.	ibmvcfg listvols

Table 52. Pool management commands (continued)

Command	Description	Example
ibmvcfg listvols all	Lists all VDisks, including information about size, location, and VDisk to host mappings.	ibmvcfg listvols all
ibmvcfg listvols free	Lists the volumes that are currently in the free pool.	ibmvcfg listvols free
ibmvcfg listvols unassigned	Lists the volumes that are currently not mapped to any hosts.	ibmvcfg listvols unassigned
ibmvcfg add -s <i>ipaddress</i>	Adds one or more volumes to the free pool of volumes. Use the -s parameter to specify the IP address of the SAN Volume Controller where the VDisks are located. The -s parameter overrides the default IP address that is set with the <code>ibmvcfg set targetSVC</code> command.	ibmvcfg add vdisk12 ibmvcfg add 600507 68018700035000000 0000000BA -s 66.150.210.141
ibmvcfg rem -s <i>ipaddress</i>	Removes one or more volumes from the free pool of volumes. Use the -s parameter to specify the IP address of the SAN Volume Controller where the VDisks are located. The -s parameter overrides the default IP address that is set with the <code>ibmvcfg set targetSVC</code> command.	ibmvcfg rem vdisk12 ibmvcfg rem 600507 68018700035000000 0000000BA -s 66.150.210.141

Error codes

The IBM TotalStorage hardware provider for the Microsoft Volume Shadow Copy Service logs error messages in the Windows Event Viewer and in private log files.

You can view error messages by going to the following locations on the Windows server where the IBM TotalStorage hardware provider is installed:

- The Windows Event Viewer in Application Events. Check this log first.
- The log file `ibmVSS.log`, which is located in the directory where the IBM TotalStorage hardware provider is installed.

Table 53 lists the errors messages that are reported by the IBM TotalStorage hardware provider.

Table 53. Error messages for the IBM TotalStorage hardware provider

Code	Message	Symbolic name
1000	JVM Creation failed.	ERR_JVM
1001	Class not found: %1.	ERR_CLASS_NOT_FOUND
1002	Some required parameters are missing.	ERR_MISSING_PARAMS

Table 53. Error messages for the IBM TotalStorage hardware provider (continued)

Code	Message	Symbolic name
1003	Method not found: %1.	ERR_METHOD_NOT_FOUND
1004	A missing parameter is required. Use the configuration utility to set this parameter: %1.	ERR_REQUIRED_PARAM
1600	The recovery file could not be created.	ERR_RECOVERY_FILE_ CREATION_FAILED
1700	ibmGetLunInfo failed in AreLunsSupported.	ERR_ARELUNSSUPPORTED_ IBMGETLUNINFO
1800	ibmGetLunInfo failed in FillLunInfo.	ERR_FILLLUNINFO_IBMGETLUNINFO
1900	Failed to delete the following temp files: %1	ERR_GET_TGT_CLEANUP
2500	Error initializing log.	ERR_LOG_SETUP
2501	Unable to search for incomplete Shadow Copies. Windows Error: %1.	ERR_CLEANUP_LOCATE
2502	Unable to read incomplete Shadow Copy Set information from file: %1.	ERR_CLEANUP_READ
2503	Unable to cleanup snapshot stored in file: %1.	ERR_CLEANUP_SNAPSHOT
2504	Cleanup call failed with error: %1.	ERR_CLEANUP_FAILED
2505	Unable to open file: %1.	ERR_CLEANUP_OPEN
2506	Unable to create file: %1.	ERR_CLEANUP_CREATE
2507	HBA: Error loading hba library: %1.	ERR_HBAAPI_LOAD
3000	An exception occurred. Check the ESSService log.	ERR_ESSSERVICE_EXCEPTION
3001	Unable to initialize logging.	ERR_ESSSERVICE_LOGGING
3002	Unable to connect to the CIM agent. Check your configuration.	ERR_ESSSERVICE_CONNECT
3003	Unable to get the Storage Configuration Service. Check your configuration.	ERR_ESSSERVICE_SCS
3004	An internal error occurred with the following information: %1.	ERR_ESSSERVICE_INTERNAL
3005	Unable to find the VSS_FREE controller.	ERR_ESSSERVICE_FREE_CONTROLLER
3006	Unable to find the VSS_RESERVED controller. Check your configuration.	ERR_ESSSERVICE_RESERVED_ CONTROLLER
3007	Unable to find suitable targets for all volumes.	ERR_ESSSERVICE_INSUFFICIENT_ TARGETS

Table 53. Error messages for the IBM TotalStorage hardware provider (continued)

Code	Message	Symbolic name
3008	The assign operation failed. Check the CIM agent log for details.	ERR_ESSSERVICE_ASSIGN_FAILED
3009	The withdraw FlashCopy operation failed. Check the CIM agent log for details.	ERR_ESSSERVICE_WITHDRAW_FAILED

Uninstalling the IBM TotalStorage hardware provider

You must use Windows to uninstall the IBM TotalStorage hardware provider from the Windows server.

Perform the following steps to uninstall the IBM TotalStorage hardware provider:

1. Log on to the Windows server as the local administrator.
2. Click **Start** → **Control Panel** from the task bar. The Control Panel window is displayed.
3. Double-click **Add or Remove Programs**. The Add or Remove Programs window is displayed.
4. Select **IBM TotalStorage Support for Microsoft Volume Shadow Copy Service** and click **Remove**.
5. Click **Yes** when you are prompted to verify that you want to completely remove the program and all of its components.
6. Click **Finish**.

The IBM TotalStorage hardware provider is no longer installed on the Windows server.

Appendix A. Configuring IBM Director overview

You can configure IBM Director for the SAN Volume Controller Call Home and e-mail notification services.

Perform the following steps to configure the IBM Director:

1. Set up an Event Action Plan.
2. Set up a correctly formatted e-mail.

Setting up an event action plan

You can configure the IBM Director to notify your system administrator when errors or events are logged by the SAN Volume Controller.

The IBM Director must receive a trap from the SAN Volume Controller before it can present the correct SAN Volume Controller information and set up an event action plan.

Perform the following steps to set up an event action plan:

1. Cause a temporary error on the SAN Volume Controller to force an SNMP trap to be sent to the master console IP address. For example, temporarily remove one of the SAN Volume Controller fibre-channel cables, which will cause error code 1060 to be displayed on the front panel of the SAN Volume Controller node. After this error is displayed, replace the fibre-channel cable, and delete the entry in the SAN Volume Controller error log.
2. Log on to the master console.
3. Double-click the **IBM Director console** icon to open the IBM Director console.
4. Log on to the IBM Director console.
5. Perform the following steps to ensure that IBM Director received the traps that were sent by SAN Volume Controller:
 - a. Double click **Event Log** in the right column of IBM Director console.
 - b. Verify that the SNMP traps were received. SAN Volume Controller traps can be identified by displaying the Event Type field of the log. For SAN Volume Controller traps, the field contains text that starts with `SNMP.iso.org.dod.internet.private.enterprises.ibm.ibmProd.190`.
If the traps were not received:
 - Contact your network administrator to ensure that there was not a networking problem.
 - Verify that the error notification setting on the SAN Volume Controller is not set to none.
 - Verify that the master console IP address has been configured.
 - c. Close the **Event Log**.
 - d. Click **Tasks** → **Event Action Plan Builder** in the IBM Director main panel. The Event Action Plan Builder panel is displayed.
 - e. Expand the **Send an Internet (SMTP) E-mail** hierarchy in the right column of the Event Action Plan Builder panel.
 - f. Double click **2145EventNot**.
 - g. Enter the following information in the displayed form:
 - **Internet E-mail Address**

- Enter an e-mail address (for example, the e-mail address of the system administrator).
 - **Reply to**
 - Enter the e-mail address to which you want replies to be directed.
 - **SMTP E-mail server**
 - Enter the name or IP address of the SMTP mail server.
 - **SMTP port**
 - Enter the port number through which e-mail is sent to your e-mail server. The default is 25.
 - **Subject of E-mail Message**
 - Enter the following text: 2145 Event Notification.
 - **Body of E-mail Message**
 - Enter any information that you want to be sent to the recipient of the e-mail (for example, machine location information). The body of the e-mail will also contain all the SNMP trap data and the details of the event.
- h. Click **File** → **Save**.
- i. Close the Event Action Plan Builder panel.
- j. Close the main IBM Director panel.

Setting up an e-mail notification to IBM

You can set up an e-mail call home to IBM if the IBM Director has been installed on a separate machine or is reinstalled on the master console.

1. From the IBM Director Console menu bar, select **Tasks** → **Event Action Plan Builder**.
2. In the **Actions** column, right-click on **Send an Internet (SMTP) E-mail** and select **Customize**.
3. In the resulting **Customize Action: Send an Internet (SMTP) E-mail** panel fill-in:

Internet E-mail Address

- Enter the IBM Retain E-mail address
 - CALLHOME1@de.ibm.com for customers in North America, Latin America, South America and Caribbean Islands
 - CALLHOME0@de.ibm.com for customer's outside of the USA.

Reply to

- Enter the E-mail address that you require any replies to be directed

SMPT E-mail Server

- Enter the address of your E-mail server

SMPT Port

- Change this, if required to your SMTP Server port number

Subject of E-mail Message

- Fill in 2145 Error Notification.

Body of the E-mail Message

- Fill in the following information:
 - Contact name.....not required in the E-mail to Admin

Note: There is a limitation of 72 characters per field.

- Contact phone number.....not required in the E-mail to Admin
- Offshift phone number.....not required in the E-mail to Admin
- Machine location
- Record Type = 1

&iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.1
 &iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.2
 &iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.3
 &iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.4
 &iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.5
 &iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.6
 &iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.7
 &iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.8
 &iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.9
 &iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.10
 &iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.11
 &iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.12

4. Click **Save** to save the information, using the name **2145CallHome**.
5. From the **Send an Internet (SMTP) E-mail** list select the newly created **2145CallHome** E-mail and Drag and Drop it on to the **2145 Error** action plan icon in the **Event Action Plan** column. This action causes the **2145CallHome** to be call when the **2145 Error** filter is satisfied.

Setting up an e-mail user notification

You can set up the e-mails if the IBM Director has been installed on a separate machine or is reinstalled on the master console.

1. From the IBM Director Console menu bar, select **Tasks -> Event Action Plan Builder**.
2. In the **Actions** column, right-click on **Send an Internet (SMTP) E-mail** and select **Customize**.
3. In the resulting **Customize Action: Send an Internet (SMTP) E-mail** panel fill-in :

Internet E-mail Address

- Enter the E-mail address you require for notification

Reply to

- Enter the E-mail address that you require any replies to be directed

SMPT E-mail Server

- Enter the address of your E-mail server

SMPT Port

- Change this, if required to your SMTP Server port number

Subject of E-mail Message

- Fill in 2145 Error Notification.

Body of the E-mail Message

- Fill in the following information:
 - # Machine location = xxxx

iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.1
 iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.2
 iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.3
 iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.4

iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.5
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.6
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.7
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.8
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.9
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.10
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.11
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.12

Where xxxx is information relevant to your organization.

4. Click **Save** to save the information, using the name **2145ErrorNot**.
5. From the **Send an Internet (SMTP) E-mail** list select the newly created **2145ErrorNot** E-mail and Drag and Drop it on to the **2145 Event** action plan icon in the **Event Action Plan** column. This action causes the **2145ErrorNot** to be call when the **2145 Event** filter is satisfied.

Setting up SNMP traps

You can set up SNMP traps if the master console has been installed on a separate machine.

Prerequisites

There are two steps required to enable the Call Home service:

1. Set up the SAN Volume Controller SNMP Trap destination, a specific machine (IP Address)
2. Set up IBM Director to send a correctly formatted e-mail

Overview

The SAN Volume Controller SNMP trap destination is normally set up as part of the SAN Volume Controller installation process, but can also be set up using the SAN Volume Controller command-line interface (CLI) or the SAN Volume Controller Console.

Appendix B. Event codes

The system generates information and configuration event codes.

There are two different types of event codes:

- Information event codes
- Configuration event codes

Information event codes provide information on the status of an operation. Information event codes are recorded in the error log and an SNMP trap and sometimes an e-mail is generated if the corresponding management flag is set in the Preference cache.

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.

Information event codes

The information event codes provide information on the status of an operation.

Information event codes are recorded in the error log and in an SNMP trap. An e-mail is generated if the corresponding management flag is set in the Preference cache.

Information event codes are reported in an SNMP trap as either information type (I) descriptions or warning type (W) descriptions. You can use the SNMP trap type to determine if the information event resulted from an expected or unexpected condition. An information event report for an SNMP trap of type (W) might require user attention.

Table 54. Information event codes

Event code	Type	Description
980221	I	Error log cleared.
980310	I	Degraded or offline Managed Disk group is now online.
980435	W	Failed to obtain directory listing from remote node.
980440	W	Failed to transfer file from remote node.
980446	I	Secure Delete complete.
980500	W	Featurization Violation.
981001	W	Cluster Fabric View has been updated by a multiphase discovery
981007	W	Preferred port is not being used for Managed Disk access.
981014	W	LUN Discovery failed. Cluster has a connection to a device through this node but this node cannot discovery the Managed Disks associated LUN correctly.
981015	W	LUN capacity equals or exceeds maximum, only first 2 TB of disk will be accessed.
981020	W	Managed Disk error count warning threshold met.
982003	W	Insufficient Virtual Extents.

Table 54. Information event codes (continued)

Event code	Type	Description
982004	W	Migration suspended due to insufficient Virtual Extents or too many media errors on the source MDisk.
982007	W	Migration Stopped.
982009	I	Migration Complete.
982010	W	Copied disk I/O medium error.
983001	I	FlashCopy prepared.
983002	I	FlashCopy complete.
983003	W	FlashCopy stopped.
984001	W	First customer data being pinned in a Virtual Disk working set.
984002	I	All customer data in a Virtual Disk working set now unpinned.
984003	W	Virtual Disk working set cache mode being changed to synchronous destage because too much pinned data has now been unpinned for that Virtual Disk working set.
984004	I	Virtual Disk working set cache mode now allows asynchronous destage because enough customer data has now been unpinned for that Virtual Disk working set.
985001	I	Metro Mirror, background copy complete.
985002	I	Metro Mirror ready to restart.
985003	W	Unable to find path to disk in remote cluster within timeout.
987102	W	Node power-off requested from power switch.
987103	W	Coldstart.
987301	W	Connection to a configured remote cluster has been lost.
987400	W	The node unexpectedly lost power but has now been restored to the cluster.
988100	W	An overnight maintenance procedure has failed to complete. Resolve any hardware and configuration problems that you are experiencing on the SAN Volume Controller cluster. If the problem persists, contact your IBM service representative for assistance.

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 55. Configuration event codes

Event code	Description
990101	Modify cluster (attributes in the svctask chcluster command)
990105	Delete node from cluster (attributes in the svctask rmnode command)

Table 55. Configuration event codes (continued)

Event code	Description
990106	Create host (attributes in the svctask mkhost command)
990117	Create cluster (attributes in the svctask mkcluster command)
990118	Modify node (attributes in the svctask chnode command)
990119	Configure set controller name
990120	Shut down node (attributes in the svctask stopcluster command)
990128	Modify host (attributes in the svctask chhost command)
990129	Delete node (attributes in the svctask rmnode command)
990138	Virtual Disk Modify (attributes in the svctask chvdisk command)
990140	Virtual Disk Delete (attributes in the svctask rmvdisk command)
990144	Modify Managed Disk group (attributes in the svctask chmdiskgrp command)
990145	Delete Managed Disk group (attributes in the svctask rmdiskgrp command)
990148	Create Managed Disk group (attributes in the svctask mkmdiskgrp command)
990149	Modify Managed Disk (attributes in the svctask chmdisk 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 svctask mkvdisk command)
990173	Add a Managed Disk to Managed Disk group (attributes in the svctask addmdisk command)
990174	Delete a Managed Disk from Managed Disk group (attributes in the svctask rmdisk command)
990178	Add a port to a Host (attributes in the svctask addhostport command)
990179	Delete a port from a Host (attributes in the svctask rmhostport command)
990182	Create a Virtual Disk to Host SCSI mapping (attributes in the svctask mkvdiskhostmap command)
990183	Delete an Virtual Disk to Host SCSI mapping (attributes in the svctask rmdiskhostmap command)
990184	Create a FlashCopy mapping (attributes in the svctask mkfcmap command)
990185	Modify a FlashCopy mapping (attributes in the svctask chfcmap command)
990186	Delete a FlashCopy mapping (attributes in the svctask rmfcmap command)
990187	Prepare a FlashCopy mapping (attributes in the svctask prestartfcmap command)
990188	Prepare a FlashCopy consistency group (attributes in the svctask prestartfcconsistgrp command)

Table 55. Configuration event codes (continued)

Event code	Description
990189	Trigger a FlashCopy mapping (attributes in the svctask startfcmap command)
990190	Trigger a FlashCopy consistency group (attributes in the svctask startfcconsistgrp command)
990191	Stop a FlashCopy mapping (attributes in the svctask stopfcmap command)
990192	Stop a FlashCopy consistency group (attributes in the svctask stopfcconsistgrp command)
990193	FlashCopy set name
990194	Delete a list of ports from a Host (attributes in the svctask rmhostport command)
990196	Shrink a Virtual Disk.
990197	Expand a Virtual Disk (attributes in the svctask expandvdisksize command)
990198	Expand single extent a Virtual Disk
990199	Modify govern a Virtual Disk
990203	Initiate manual Managed Disk discovery (attributes in the svctask detectmdisk command)
990204	Create FlashCopy consistency group (attributes in the svctask mkfconsistgrp command)
990205	Modify FlashCopy consistency group (attributes in the svctask chfconsistgrp command)
990206	Delete FlashCopy consistency group (attributes in the svctask rmfconsistgrp command)
990207	Delete a list of Hosts (attributes in the svctask rmhost command)
990213	Change the HWS a node belongs to (attributes in the svctask chiogrp command)
990216	Apply software upgrade (attributes in the svcservicetask applysoftware command)
990219	Analyze error log (attributes in the svctask finderr command)
990220	Dump error log (attributes in the svctask dumperrlog command)
990222	Fix error log entry (attributes in the svctask cherrstate command)
990223	Migrate a single extent (attributes in the svctask migrateexts command)
990224	Migrate a number of extents
990225	Create Metro Mirror relationship (attributes in the svctask mkrcrelationship command)
990226	Modify Metro Mirror relationship (attributes in the svctask chrcrelationship command)
990227	Delete Metro Mirror relationship (attributes in the svctask rmrcrelationship command)
990229	Start Metro Mirror relationship (attributes in the svctask startrcrelationship command)
990230	Stop Metro Mirror relationship (attributes in the svctask stoprcrelationship command)

Table 55. Configuration event codes (continued)

Event code	Description
990231	Switch a Metro Mirror relationship (attributes in the svctask switchrcrelationship command)
990232	Start Metro Mirror consistency group (attributes in the svctask startrcconsistgrp command)
990233	Stop Metro Mirror consistency group (attributes in the svctask stoprcconsistgrp command)
990234	Switch a Metro Mirror consistency group (attributes in the svctask switchrcconsistgrp 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 svctask mkpartnership command)
990238	Modify partnership with remote cluster (attributes in the svctask chpartnership command)
990239	Delete partnership with remote cluster (attributes in the svctask rmpartnership command)
990240	Create Metro Mirror consistency group (attributes in the svctask mkrcconsistgrp command)
990241	Modify Metro Mirror consistency group (attributes in svctask chrcconsistgrp)
990242	Delete Metro Mirror consistency group (attributes in the svctask rmrcconsistgrp command)
990245	Node pend
990246	Node remove
990247	Node unpend
990380	Time zone changed (attributes in the svctask settimezone command)
990383	Change cluster time (attributes in the svctask setclustertime command)
990385	System time changed
990386	SSH key added (attributes in the svctask addsshkey command)
990387	SSH key removed (attributes in the svctask rmsshkey command)
990388	All SSH keys removed (attributes in the svctask rmallsshkeys command)
990390	Add node to the cluster
990395	Shutdown or reset node
990410	Software Install started
990415	Software Install completed
990420	Software Install failed
990430	Planar Serial Number changed
990501	The featurization has changed. See feature log for details.
990510	Configuration limits have been changed.
991024	IO tracing has finished, trigger occurred for given Managed Disk.

Appendix C. Object types

You can use the object code to determine the object type.

Table 56 lists the object codes and corresponding object types.

Table 56. Object types

Object code	Object type
1	mdisk
2	mdiskgrp
3	vdisk
4	node
5	host
7	iogroup
8	fcgrp
9	rcgrp
10	fcmap
11	rcmap
12	wwpn
13	cluster
15	hba
16	device
17	SCSI lun
18	quorum
19	time seconds
20	ExtSInst
21	ExtInst
22	percentage
23	system board
24	processor
25	processor cache
26	memory module
27	fan
28	fc card
29	fc device
30	software
31	front panel
32	ups
33	port
34	adapter
35	migrate

Appendix D. Valid combinations of FlashCopy and Metro Mirror functions

The following table outlines the combinations of FlashCopy and Metro Mirror functions that are valid for a single virtual disk (VDisk).

Table 57. Valid combinations of FlashCopy and Metro Mirror interactions

FlashCopy	Metro Mirror Primary	Metro Mirror Secondary
FlashCopy source	Supported	Supported
FlashCopy target	Not supported	Not supported

Appendix E. Moving data between MDisk groups with Copy Services

You cannot use the SAN Volume Controller data migration function to move a virtual disk (VDisk) between managed disk (MDisk) groups that have different extent sizes. However, you can use Copy Services to move the data by copying a VDisk between MDisk groups that have different extent sizes.

To copy a VDisk between MDisk groups that have different extent sizes, you have the following options:

- Use FlashCopy to copy a VDisk between a source and a destination MDisk group.
- Use intracluster Metro Mirror to copy a VDisk between a source and a destination MDisk group.

Using FlashCopy

This option is available if you have licensed the FlashCopy feature.

Use the following guidelines for using FlashCopy to copy a VDisk between a source and a destination MDisk group that have different extent sizes:

- The VDisk must not be in another FlashCopy or Metro Mirror relationship.
- Stop all I/O operations from the hosts while the VDisk is being copied.
- Once the copy is complete, configure host mappings for the new VDisk and configure the hosts to access the destination VDisk rather than the source VDisk.

Using Metro Mirror

This option is available if you have licensed the Metro Mirror feature.

Use the following guidelines for using Metro Mirror to copy a VDisk between a source and a destination MDisk group that have different extent sizes:

- The VDisk must not be in another FlashCopy or Metro Mirror relationship.
- Create an intracluster Metro Mirror relationship for the VDIs.
- I/O operations from the host can continue while the copy is being performed, however there will be some performance degradation while write I/O operations are mirrored.
- Once the copy is complete, stop all I/O operations from the hosts. Configure host mappings for the new VDisk and configure the hosts to access the destination VDisk rather than the source VDisk.

Accessibility

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

Features

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

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

Navigating by keyboard

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

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

Accessing the publications

You can view the publications for the SAN Volume Controller in Adobe Portable Document Format (PDF) using the Adobe Acrobat Reader. The PDFs are provided at the following Web site:

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

Related reference

“SAN Volume Controller library and related publications” on page xx

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

Notices

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

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

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

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

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

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATIONS "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

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

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

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

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

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

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

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

Related reference

“Trademarks”

Trademarks

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

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

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

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

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

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

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

Glossary

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

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

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 run from within the SAN Volume Controller application and are documented in the service guide.

disconnected

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

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 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*.

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, 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, 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, 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, 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 through a fibre-channel interface.

host ID

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

host zone

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

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

input/output (I/O)

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

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

N

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

node rescue

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

non-RAID

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

O

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

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

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

paused

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

pend To cause to wait for an event.

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

primary virtual disk

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

PuTTY

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

Q**quorum disk**

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

R

RAID See *redundant array of independent disks*.

RAID 0

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

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

RAID 1

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

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

RAID 10

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

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 Global Mirror, 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*.

Global Mirror

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

S

SAN See *storage area network*.

SAN Volume Controller fibre-channel port fan in

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

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

Simple Network Management Protocol (SNMP)

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

SNMP See *Simple Network Management Protocol*.

stand-alone relationship

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

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

stopped

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

storage area network (SAN)

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

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 Global Mirror, 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 reinitiate 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)

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

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

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 xv
- Access Logix 239
- accessibility
 - keyboard 369
 - shortcut keys 369
- adding
 - managed disks 160
 - managed disks (MDisks) 110, 113
 - nodes 96
 - storage controllers
 - using the 230
 - using the CLI (command-line interface) 231
- analyzing error logs 140
- audience xv

B

- backup cluster configuration files
 - creating 196
- backup configuration files
 - creating 195
 - deleting 201
 - using the CLI 202
 - restoring 199
- BladeCenter fabric support 65
- book
 - about this xv

C

- certificate
 - truststore 345
- changing
 - cluster password 95
 - passwords 191
- CLI (command-line interface)
 - getting started 147
 - issuing commands from a PuTTY SSH client system 150
 - preparing SSH client systems 148, 149
 - upgrading software 203
 - using to set cluster features 152
- clusters
 - adding managed disks (MDisks) 110
 - adding nodes 96
 - backing up configuration file 14, 195
 - backing up configuration file using the CLI 196
 - changing password 95
 - creating
 - from the front panel 80
 - deleting nodes 105, 187
 - error logs 192
 - gateway address
 - changing 189
 - including managed disks (MDisks) 110

- clusters (*continued*)
 - IP address
 - changing 189
 - IP failover 15
 - logs 192
 - maintaining 138
 - Metro Mirror partnerships
 - deleting 137
 - overview 13
 - properties 96, 153
 - recovering nodes 177
 - removing nodes 105, 187
 - renaming 107
 - resetting the SSH fingerprint 145
 - restoring backup configuration files 199
 - setting
 - features 152
 - time 152
 - setting time 94
 - shutting down 108, 193
 - subnet mask
 - changing 189
 - viewing feature logs 192
 - viewing properties 96
- codes
 - configuration events 358
 - events 357
 - information events 357
- command-line interface (CLI)
 - getting started 147
 - issuing commands from a PuTTY SSH client system 150
 - preparing SSH clients 148, 149
 - upgrading software 203
 - using to set cluster features 152
 - using to set cluster time 152
- commands
 - ibmvcfg add 349
 - ibmvcfg listvols 349
 - ibmvcfg rem 349
 - ibmvcfg set cimomHost 348
 - ibmvcfg set cimomPort 348
 - ibmvcfg set FlashCopyVer 348
 - ibmvcfg set namespace 348
 - ibmvcfg set password 348
 - ibmvcfg set trustpassword 348
 - ibmvcfg set truststore 348
 - ibmvcfg set username 348
 - ibmvcfg set usingSSL 348
 - ibmvcfg set vssFreeInitiator 348
 - ibmvcfg set vssReservedInitiator 348
 - ibmvcfg showcfg 348
 - svcconfig backup 196
 - svcconfig restore 199
 - svcinfo caterrlog 204
 - svcinfo caterrlogbyseqnum 204
 - svcinfo lserrlogbyxxxx 204
 - svcinfo lsxxxx 204

- commands *(continued)*
 - svcinfolxxxxcandidate 204
 - svcinfolxxxxdumps 204
 - svcinfolxxxxextent 204
 - svcinfolxxxxmember 204
 - svcinfolxxxxprogress 204
 - svcservicetask rmnode 204
 - svctask detectmdisk 227
 - svctask rmnode 204
- communications
 - determining between hosts and virtual disks 168
- configuration
 - event codes 358
 - maximum sizes 78
 - node failover 15
 - rules 52
- configuration requirements 76
- configuring
 - clusters 86
 - disk controllers 215, 216, 217, 218, 219
 - DS4000 series Storage Manager 266
 - Enterprise Storage Server 222, 259
 - error notification settings 139
 - FAST Storage Manager 222
 - FAST Storage Server 222
 - IBM DS4000 Storage Server 264
 - nodes 60
 - settings
 - error notification 139
 - switches 61
 - Web browsers 85
- consistency group, Metro Mirror 51
- consistency groups, FlashCopy 42
 - creating 130
 - deleting 132
 - filtering 130
 - modifying 131
 - starting 131
 - stopping 131
- console
 - SAN Volume Controller
 - portfolio 83
 - starting 85
 - task bar 83
 - work area 84
- controllers
 - adding
 - using the 230
 - using the CLI (command-line interface) 231
 - advanced functions
 - DS4000 series 270
 - EMC CLARiiON 246
 - EMC Symmetrix 254
 - HDS Lightning 280
 - HDS Thunder 289
 - HP StorageWorks 307, 308
 - HPQ EVA 316
 - IBM Enterprise Storage Server 262
 - concurrent maintenance
 - DS4000 series 269
 - EMC CLARiiON 243
 - controllers *(continued)*
 - concurrent maintenance *(continued)*
 - EMC Symmetrix 251
 - Enterprise Storage Server 261
 - HDS Lightning 278
 - HDS Thunder 286
 - HP StorageWorks 304
 - configuration
 - EMC CLARiiON 238, 239, 241, 242, 246
 - EMC Symmetrix 250, 255
 - Enterprise Storage Server 259
 - FAST 263
 - HDS Lightning 276
 - HDS Thunder 284
 - HP StorageWorks 296, 300, 302, 309, 318
 - HPQ EVA 314
 - IBM DS4000 264
 - controller settings
 - EMC CLARiiON 248
 - firmware
 - DS4000 series 268
 - EMC CLARiiON 243
 - EMC Symmetrix 251
 - HDS Lightning 277
 - HDS Thunder 286
 - HP StorageWorks 304
 - HPQ EVA 315
 - IBM Enterprise Storage Server 260
 - global settings
 - EMC CLARiiON 247
 - EMC Symmetrix 256
 - HDS Thunder 292
 - HPQ EVA 318
 - IBM DS4000 series 275
 - Lightning 283
 - host settings
 - HPQ EVA 319
 - interface
 - DS4000 series 272
 - HP StorageWorks 305, 318
 - logical unit creation and deletion
 - EMC CLARiiON 246
 - EMC Symmetrix 254
 - HDS Thunder 290
 - HP StorageWorks 308
 - HPQ EVA 316, 317
 - IBM DS4000 series 272
 - IBM Enterprise Storage Server 263
 - logical unit presentation
 - HPQ EVA 317
 - LU configuration
 - HDS Lightning 281
 - LU settings
 - EMC CLARiiON 249
 - EMC Symmetrix 257
 - HDS Thunder 294
 - HP StorageWorks 311
 - HPQ EVA 319
 - IBM DS4000 series 275
 - Lightning 284

- controllers (*continued*)
 - mapping settings
 - EMC Symmetrix 258
 - models
 - EMC CLARiiON 242
 - EMC Symmetrix 250
 - HDS Lightning 277
 - HDS Thunder 285
 - HP StorageWorks 303
 - HPQ EVA 315
 - HPQ XP 277
 - IBM DS4000 267
 - IBM Enterprise Storage Server 260
 - StorageTek 267
 - Sun StorEdge 277
 - port selection 227
 - port settings
 - EMC CLARiiON 248
 - EMC Symmetrix 256
 - HDS Lightning 283
 - HDS Thunder 293
 - HP StorageWorks 310
 - quorum disks
 - DS4000 series 270
 - EMC CLARiiON 245
 - EMC Symmetrix 254
 - HDS Lightning 280
 - HDS Thunder 288
 - HP StorageWorks 306
 - HPQ EVA 316
 - IBM Enterprise Storage Server 262
 - registering
 - EMC CLARiiON 240
 - removing 233
 - using the CLI (command-line interface) 234
 - settings
 - DS4000 series 274
 - HDS Thunder 291, 293
 - HP StorageWorks 310, 313
 - IBM DS4000 series 273
 - Lightning 282, 283
 - sharing
 - DS4000 series 269
 - EMC CLARiiON 244
 - EMC Symmetrix 252
 - HDS Lightning 279
 - HDS Thunder 287
 - HP StorageWorks 305
 - HPQ EVA 316
 - IBM Enterprise Storage Server 261
 - switch zoning
 - EMC CLARiiON 245
 - EMC Symmetrix 253
 - HP StorageWorks 306
 - HPQ EVA 316
 - IBM Enterprise Storage Server 262
 - target port groups
 - Enterprise Storage Server 261
 - user interface
 - EMC CLARiiON 244
 - EMC Symmetrix 252

- controllers (*continued*)
 - user interface (*continued*)
 - HDS Lightning 278
 - HDS Thunder 286
 - HPQ EVA 315
 - IBM DS4000 series 269
 - IBM Enterprise Storage Server 261
 - zoning 68
 - conventions xix
 - numbering xix
 - copy services
 - overview 35
 - creating
 - clusters
 - from the front panel 80
 - FlashCopy
 - mappings 128, 165, 167
 - FlashCopy consistency groups 130
 - managed disk (MDisk) groups 112
 - Metro Mirror
 - consistency groups 135
 - partnerships 136
 - quorum disks 236
 - virtual disk-to-host mappings 164
 - virtual disks (VDisks) 115

D

- data
 - migrating 219, 220
- data migration
 - DS4000 series 271
- deleting
 - backup configuration files 201
 - using the CLI 202
 - FlashCopy
 - mappings 130
 - hosts 127
 - Metro Mirror
 - consistency groups 136
 - partnerships 137
 - relationships 134
 - nodes 105, 187
 - virtual disks 120
- determining
 - communications between hosts and virtual disks 168
- discovering
 - managed disks 109, 157, 237
 - MDisks 109
- disk controller systems
 - renaming 229
- disk controllers
 - configuring 215, 216
 - overview 23
- disks
 - migrating 184
 - migrating image mode 187
- disruptive software upgrade
 - using the CLI (command-line interface) 211

E

- e-mail
 - setting up 354, 355
- emphasis in text xix
- error messages, IBM TotalStorage Support for Microsoft Shadow Copy Service 350
- errors
 - notification settings 139
- Ethernet
 - link failures 15
- events
 - codes 357
 - configuration 358
 - information 357
 - setting up an action plan for 353
- expanding
 - logical units 225
 - virtual disks 179, 180
- expiration 345
- expired
 - certificate 345
- extents
 - migrating
 - using the CLI (command-line interface) 182

F

- fabrics
 - BladeCenter support 65
- features
 - disabling features 138
 - enabling features 138
 - setting
 - using the CLI (command-line interface) 152
 - viewing features 138
 - viewing logs 138
- fibre-channel switches 61
- filtering
 - FlashCopy
 - consistency groups 130
 - mappings 128
 - Metro Mirror
 - consistency groups 135
 - relationships 132
- FlashCopy 40
 - background copy rate 46
 - consistency groups 42, 44
 - creating consistency groups 130
 - creating mappings 165, 167
 - for Volume Shadow Copy service 339
 - mappings 38, 46, 164
 - overview 35
 - renaming consistency groups 131
- free pool of volumes 347
- front panel
 - password 153

G

- gateway address
 - changing 189
- generating new truststore certificate 345
- getting started
 - using the CLI (command-line interface) 147
 - using the command-line interface (CLI) 147
 - using the SAN Volume Controller Console 93
- governing 19
- guide
 - about this xv
 - who should read xv
- guidelines
 - zoning 68

H

- hardware provider 339
- HBAs (host bus adapters)
 - configuration 59
 - replacing 127
- HDS Thunder
 - support 285
- host bus adapters (HBAs)
 - configuration 59
 - replacing 127
- host objects
 - creating 162
- hosts
 - creating 123
 - deleting 127
 - determining VDisk names 168
 - filtering 124
 - flushing data 37
 - mapped virtual disks (VDisks) 125
 - mapping virtual disks (VDisks) 164
 - overview 32
 - replacing HBA 127
 - viewing
 - general details 124
 - mapped I/O groups 125
 - port details 124
 - zoning 68

I

- I/O governing 19
- I/O groups
 - overview 17
 - renaming 107
- IBM Director
 - configuring 353
 - overview 353
- IBM TotalStorage hardware provider
 - described 339
 - installation procedure 339
 - system requirements 340
- IBM TotalStorage Metro Mirror for SAN Volume Controller 48

- IBM TotalStorage Support for Microsoft Shadow Copy Service
 - error messages 350
 - ibmvfcg.exe 348, 349
- IBM TotalStorage Support for Microsoft Volume Shadow Copy Service
 - creating pools of volumes 347
 - described 339
 - installation procedure 339
 - installing 340
 - system requirements 340
- ibmvfcg.exe 348, 349
- image mode
 - VDisks 120
- image mode VDisks
 - converting to managed mode
 - using 121
 - using CLI (command-line interface) 186
- including
 - managed disks (MDisks) 110
- information
 - event codes 357
- information center xx
- installing
 - IBM TotalStorage Support for Microsoft Volume Shadow Copy Service 340
 - SAN Volume Controller Console 323, 328
 - software 204
 - verification 333
- inter-switch link (ISL)
 - support for long links 11
- IP addresses
 - changing 189
 - modifying 95
- issuing
 - CLI commands 150

K

- keyboard 369
- keyboard shortcuts 369

L

- language 191
- links, physical 10
- listing
 - dump files 139
 - log files 139
- logical unit mapping 226
- logical units
 - expanding 225

M

- maintaining
 - passwords 95, 153
 - SSH keys 190
- maintenance procedures
 - clusters 138
- managed disk (MDisk) 24

- managed disk (MDisk) groups
 - adding
 - managed disks 113
 - creating 112
 - renaming 114
- managed disk groups
 - creating using the CLI 159
- managed disks (MDisks)
 - adding 110, 160
 - discovering 109, 157, 237
 - displaying groups 112
 - expanding 225
 - including 110
 - rebalancing access 157, 237
 - removing from a managed disk group 113
 - removing from an MDisk group 113
 - renaming 110
 - virtual disks (VDisks) relationships 169
- managed mode virtual disks
 - converting from image mode
 - using the 121
 - using the CLI (command-line interface) 186
- mapping events 40
- mappings, FlashCopy
 - background copy rate 46
 - creating 128
 - deleting 130
 - events 40
 - filtering 128
 - modifying 129
 - starting 128
 - stopping 129
- master console
 - error 345
- maximum configuration 78
- MDisk (managed disk) 24
- MDisk (managed disk) groups
 - deleting 114
 - forced 114
 - renaming 114
- MDisks (managed disks)
 - adding 160
 - VDisk (virtual disks) relationships 169
- measurements xix
- mesh configuration 52
- metro mirror
 - upgrading cluster software 203
- Metro Mirror
 - consistency groups
 - creating 135
 - deleting 136
 - filtering 135
 - starting 133, 135
 - stopping 133, 136
 - deleting partnerships 137
 - overview 47, 48, 51
 - partnerships 48
 - creating 136
 - relationships
 - deleting 134
 - filtering 132

- Metro Mirror *(continued)*
 - relationships *(continued)*
 - starting 133, 135
 - stopping 133, 136
 - zoning considerations 71
- migrating
 - data 219, 220
 - extents
 - using the CLI (command-line interface) 182
 - VDisks (virtual disks) 171
 - virtual disks (VDisks) 101, 122
- migration 115, 271
- mkcertificate.bat 345
- modifying 226
 - FlashCopy
 - consistency groups 131
 - mappings 129
 - Metro Mirror
 - partnerships 137
 - relationships 134
- monitoring
 - software upgrades 206, 213

N

- node
 - failover 15
- node status 16
- nodes
 - adding 96, 153
 - configuration 16, 60
 - deleting 105, 187
 - overview 13
 - removing 105, 187
 - renaming 105
 - replacing 101, 172
 - rescuing 212
 - returning to cluster 177
 - shutting down 109
 - viewing
 - general details 98, 156
- notices
 - legal 371

O

- object classes and instances 363
- object codes 363
- object types 363
- operating over long distances 73
- ordering publications xxi
- overview
 - SSH (secure shell) 205
 - zoning 66

P

- partnerships, Metro Mirror
 - modifying 137
- passwords
 - changing 191

- passwords *(continued)*
 - front panel 153
- plink utility
 - running 150
- power requirements 61
- preinstalled software
 - recovering from installation failures 213
- preparing
 - SSH client system
 - overview 147
 - to issue CLI commands 148, 149
- public SSH keys
 - storing 143
- publications
 - ordering xxi
- PuTTY
 - issuing CLI commands from 150
 - running the plink utility 150
- PuTTY scp
 - overview 205

Q

- quorum disks
 - creating 236
 - DS4000 series 270
 - setting 110

R

- rebalancing
 - managed disks (MDisks) access 157, 237
- recovering
 - offline virtual disks (VDisks) 118
 - using CLI 176
 - software automatically 213
- related information xx
- relationships, Metro Mirror
 - creating 132
 - deleting 134
 - modifying 134, 135
 - overview 49
 - starting 133, 135
 - stopping 133, 136
- removing
 - nodes 105, 187
 - storage controllers 233
 - using the CLI (command-line interface) 234
- renaming
 - a Metro Mirror consistency group 135
 - disk controller systems 229
 - I/O groups 107
 - managed disks 110
 - MDisks 110
 - nodes 105
- replacing
 - nodes 101, 172
- requirements 322
 - installing SAN Volume Controller Console software 321
 - SAN Volume Controller Console 322, 323

- requirements *(continued)*
 - upgrading SAN Volume Controller Console software 321
- reserved pool of volumes 347
- resetting
 - SSH fingerprint for a cluster 145
- running
 - cluster maintenance procedure 138
 - PuTTY plink utility 150

S

- SAN Volume Controller
 - adding to cluster 153
 - configuring nodes 60
 - Console
 - banner 82
 - layout 82
 - portfolio 83
 - post installation tasks 334
 - starting 85
 - task bar 83
 - work area 84
 - front panel password 153
 - overview 1
 - properties 156
 - removing 336
 - renaming 105
 - replacing nodes 172
 - shutting down 109
 - software upgrade problems 213
 - upgrading software 206
 - upgrading software automatically 206
 - upgrading software using the CLI 209
- SAN Volume Controller Console
 - backing up configuration file 195
 - banner 82
 - installing 323, 328
 - launching the Web application 93
 - layout 82
 - requirements 322, 323
 - starting 93
 - upgrading 323, 328
- SAN Volume Controller software
 - copying using PuTTY scp 205
- scanning
 - fibre-channel network 157, 237
 - rebalancing MDisk access 157, 237
- SDD 5
- secure shell (SSH)
 - client system
 - issuing CLI commands from 150
 - preparing to issue CLI commands 149
 - overview 205
- Secure Shell (SSH)
 - adding keys 190
 - client system
 - overview 147
 - preparing to issue CLI commands 148
 - keys
 - adding 142

- Secure Shell (SSH) *(continued)*
 - keys *(continued)*
 - replacing key pair 144
 - replacing private key 143
 - storing 143
 - listing keys 190
 - managing keys 142
 - resetting fingerprint 145
- setting
 - action plan for events 353
 - cluster features
 - using the CLI (command-line interface) 152
 - cluster time 94
 - using the CLI (command-line interface) 152
 - copy direction 134
 - e-mail account 354, 355
 - features
 - using the CLI (command-line interface) 152
 - quorum disks 110
 - time
 - using the CLI (command-line interface) 152
 - traps 356
- settings
 - error notification 190
 - language 191
- shadow copy 339
- shortcut keys 369
- shrinking
 - VDisks 117
- shutting down
 - clusters 108
 - nodes 109
- SNMP traps 139, 190, 356
- software
 - automatic recovery 213
 - automatic upgrades 206
 - copying using PuTTY scp 205
 - description 204
 - installing 204
 - installing SAN Volume Controller Console 321
 - manual recovery 213
 - recovering automatically 213
 - recovering manually 213
 - upgrading 204, 206
 - upgrading automatically 206
 - upgrading SAN Volume Controller Console 321
 - upgrading using the command-line interface (CLI) 209
- software upgrades
 - recovering 213
- software, upgrading
 - disruptive
 - using the CLI (command-line interface) 211
 - using the CLI (command-line interface) 203
- SSH (secure shell)
 - adding keys 190
 - client system
 - issuing CLI commands from 150
 - preparing to issue CLI commands 149
 - listing keys 190
 - resetting fingerprint 145

- SSH (Secure Shell)
 - client system
 - overview 147
 - preparing to issue CLI commands 148
 - keys
 - adding 142
 - replacing key pair 144
 - replacing private key 143
 - storing 143
 - managing keys 142
- starting
 - FlashCopy
 - consistency groups 131
 - mappings 128
 - Metro Mirror
 - consistency groups 133, 135
 - relationships 133, 135
- status
 - of node 16
- stopping
 - FlashCopy
 - mappings 129
 - Metro Mirror
 - consistency groups 133, 136
 - relationships 133, 136
 - Remote Copy
 - consistency groups 131
- storage controllers
 - adding
 - using the 230
 - using the CLI (command-line interface) 231
 - removing 233
 - using the CLI (command-line interface) 234
- storage subsystems
 - servicing 237
- storing
 - public SSH keys 143
- strategy
 - software upgrade
 - using the CLI (command-line interface) 203
- subnet mask
 - changing 189
- subsystem device driver (SDD) 5
- support
 - inter-switch link (ISL) 11
 - Web sites xxi
- switches
 - configuring 61
 - operating over long distances 73
 - zoning 66
- system requirements, IBM TotalStorage Support for Microsoft Volume Shadow Copy Service 340

T

- text emphasis xix
- time
 - setting
 - using the CLI (command-line interface) 152
- trademarks 372

- truststore
 - certificate 345
- truststore certificate 345

U

- uninterruptible power supply
 - configuration 21
 - operation 22
 - overview 19
- upgrading
 - SAN Volume Controller Console 323, 328
 - software 206
 - software automatically 206
 - software using the command-line interface (CLI) 209
- upgrading software
 - disruptive
 - using the CLI (command-line interface) 211
 - strategy
 - using the CLI (command-line interface) 203

V

- VDisk (virtual disk)
 - expanding 180
- VDisk (virtual disks)
 - determining mappings 169
- VDisks (virtual disks)
 - image mode 120
- VDisks (virtual disks)
 - converting
 - from image mode to managed mode 121, 186
 - creating 115, 161
 - deleting 120
 - determining name of 168
 - expanding 179
 - MDisks (managed disks) relationships 169
 - migrating 101, 122, 171, 185
 - moving offline 178
 - offline 119
 - overview 30
 - recovering from offline 118
 - using CLI 176
 - shrinking 117
- viewing
 - clusters
 - feature logs 138
- virtual disk-to-host mapping
 - description 33
- virtual disks (VDisks) 180
 - converting
 - from image mode to managed mode 121, 186
 - creating 115
 - determining mappings 169
 - determining name of 168
 - image mode 120
 - managed disks (MDisks) relationships 169
 - migrating 101, 115, 122, 171
 - moving offline 178
 - offline 119

- virtual disks (VDisks) *(continued)*
 - overview 30
 - recovering from offline 118
 - using CLI 176
 - shrinking 117
- virtualization
 - asymmetric 8
 - overview 6
 - symmetric 9
- Volume Shadow Copy service 339

W

- Web browsers
 - configuring 85
- Web sites xxi
- who should read this guide xv

Z

- zoning
 - considerations for Metro Mirror 71
 - controllers 68
 - guidelines 68
 - hosts 68
 - overview 66

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

IBM TotalStorage SAN Volume Controller
Configuration Guide
Version 3.1.0

Publication No. SC26-7543-05

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: 64P8401

Printed in USA

SC26-7543-05



(1P) P/N: 64P8401



Spine information:



IBM TotalStorage SAN Volume
Controller

SAN Volume Controller Configuration Guide

Version 3.1.0