

IBM System Storage SAN Volume Controller



Software Installation and Configuration Guide

Version 4.2.1

IBM System Storage SAN Volume Controller



Software Installation and Configuration Guide

Version 4.2.1

Note:

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

This edition applies to the IBM System Storage SAN Volume Controller, release 4.2.1, and to all subsequent releases and modifications until otherwise indicated in new editions. This edition replaces SC23-6628-00.

© **Copyright International Business Machines Corporation 2003, 2007. All rights reserved.**

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

Contents

Figures	xi	Long distance links for Metro Mirror and Global Mirror partnerships.	53
Tables	xiii	Using the intercluster link for host traffic	55
About this guide	xv	Metro Mirror and Global Mirror consistency groups	56
Who should use this guide	xv	Background copy bandwidth impact on foreground I/O latency	57
Summary of changes	xv	Backend storage controller requirements.	58
Summary of changes for SC23-6628-01 SAN Volume Controller Software Installation and Configuration Guide	xv	Migrating a Metro Mirror relationship to a Global Mirror relationship	59
Summary of changes for SC23-6628-00 SAN Volume Controller Configuration Guide.	xv	Using FlashCopy to create a consistent image before restarting a Global Mirror relationship	60
Emphasis	xvi	Monitoring Global Mirror performance with the IBM TotalStorage Productivity Center.	60
Numbering conventions.	xvii	The gmlinktolerance feature	61
SAN Volume Controller library and related publications.	xvii		
Related Web sites	xxi		
How to order IBM publications	xxi		
How to send your comments	xxi		
Chapter 1. SAN Volume Controller overview	1	Chapter 3. Configuring the SAN fabric	65
Virtualization	2	SAN fabric overview	66
Asymmetric virtualization	4	Configuration rules.	68
Symmetric virtualization	4	Storage subsystems.	68
SAN Volume Controller operating environment.	5	Host HBAs	72
Object overview	6	Nodes	73
Nodes and clusters	7	Fibre-channel switches.	74
I/O groups and uninterruptible power supply.	11	Port switches	77
Split cluster configurations	14	Physical links between SAN Volume Controller nodes and a switch.	78
Storage subsystems and MDisks	16	Zoning guidelines	78
MDisk groups and VDIsks	20	Example SAN environments.	80
Node management and support tools.	28	Zoning considerations for Metro Mirror and Global Mirror.	82
Master console	29	Switch operations over long distances	83
Master console software components	29	Limiting queue depth in large SANs	83
Secure Shell protocol	29	Queue depth	84
Assist On-site and remote service	30	Calculating a queue depth limit	84
Sending notifications	30	Homogeneous queue depth calculation	84
		Nonhomogeneous queue depth calculation	85
		Limiting the queue depth.	85
		Configuration requirements	86
		Supported fibre-channel extenders.	87
		Performance of fibre-channel extenders	87
Chapter 2. Copy Services features	35	Chapter 4. Installing master console software	89
FlashCopy.	35	Using the master console installation wizard	89
FlashCopy applications	36	Mirroring the boot drive	92
Host considerations for FlashCopy integrity	36		
FlashCopy mappings	37	Chapter 5. Configuring the master console	95
FlashCopy consistency groups	43	Changing the master console host name.	95
Grains and the FlashCopy bitmap	45	Configuring the internal IP network connection	96
FlashCopy indirection layer	46	Configuring the Web browser	96
Background and stopping copy.	48	Generating an SSH key pair using PuTTY	98
Metro Mirror and Global Mirror	49	Storing the private SSH key in the SAN Volume Controller Console software	99
Metro Mirror and Global Mirror relationships	50	Configuring the PuTTY session for the CLI.	99
Metro Mirror and Global Mirror relationships between two clusters	52		
Metro Mirror and Global Mirror partnerships	52		
Configuration requirements	52		

Chapter 6. Creating a SAN Volume Controller cluster.	101
Creating a cluster from the front panel	101
SAN Volume Controller Console layout	103
SAN Volume Controller Console banner	104
SAN Volume Controller Console task bar	104
SAN Volume Controller Console portfolio	104
SAN Volume Controller Console work area	105
Accessing the SAN Volume Controller Console	105
Creating a cluster using the SAN Volume Controller Console	106

Chapter 7. Using the SAN Volume Controller Console	109
Launching the SAN Volume Controller Console to manage a cluster	109
Setting the cluster date and time	110
Modifying the cluster IP addresses	111
Maintaining cluster passwords.	111
Viewing cluster properties	112
Adding nodes to a cluster	112
Viewing the node status	115
Increasing the size of a cluster	115
Adding a node to increase the size of a cluster	116
Migrating a VDisk to a new I/O group.	117
Replacing a faulty node with a spare node	118
Renaming a node	121
Deleting a node from a cluster	122
Renaming an I/O group.	123
Modifying a cluster	124
Shutting down a cluster	124
Shutting down a node	125
Discovering MDisks	125
Viewing discovery status	126
Renaming MDisks	126
Adding excluded MDisks to a cluster	126
Setting quorum disks.	127
Determining the relationship between MDisks and VDIs	127
Determining the relationship between MDisks and RAID arrays or LUNs	127
Displaying MDisk groups	128
Creating MDisk groups	128
Adding MDisks to MDisk groups	129
Removing MDisks from an MDisk group	129
Viewing the progress of an MDisk removal	130
Renaming MDisk groups	130
Displaying VDIs.	130
Deleting MDisk groups	131
Creating VDIs	131
Viewing the progress of VDisk formatting	131
Migrating VDIs	132
Viewing the progress of VDisk migration	133
Shrinking VDIs	133
Viewing virtual disk-to-host mappings	134
Determining the relationship between VDIs and MDisks	134
Recovering from offline VDIs	135
Deleting VDIs	136
Using image mode VDIs	137

Creating an image mode VDisk	137
Migration methods	139
Viewing the progress of image mode migration	139
Viewing the progress of extent migration	139
Creating hosts	140
Filtering hosts	140
Viewing host details	140
Viewing port details	141
Viewing mapped I/O groups	141
Displaying VDIs that are mapped to a host	141
Modifying a host	142
Adding ports to a host	142
Deleting ports from a host	143
Replacing an HBA in a host	143
Deleting hosts	144
Viewing fabrics.	144
Creating FlashCopy mappings.	144
Filtering FlashCopy mappings.	145
Starting FlashCopy mappings	145
Viewing the progress of a FlashCopy	145
Stopping FlashCopy mappings	146
Modifying FlashCopy mappings	146
Deleting FlashCopy mappings.	146
Creating FlashCopy consistency groups	147
Filtering FlashCopy consistency groups.	147
Starting FlashCopy consistency groups	147
Stopping FlashCopy consistency groups	148
Renaming FlashCopy consistency groups	148
Deleting FlashCopy consistency groups.	148
Creating Metro Mirror and Global Mirror relationships.	149
Filtering Metro Mirror and Global Mirror relationships.	149
Starting a Metro Mirror or Global Mirror copy process	149
Viewing the progress of Metro Mirror and Global Mirror copy processes	150
Stopping a Metro Mirror or Global Mirror copy process	150
Modifying Metro Mirror and Global Mirror relationships.	150
Switching the copy direction of a Metro Mirror or Global Mirror relationship	151
Deleting Metro Mirror or Global Mirror relationships.	151
Creating Metro Mirror or Global Mirror consistency groups	151
Filtering Metro Mirror or Global Mirror consistency groups	152
Renaming a Metro Mirror or Global Mirror consistency group	152
Starting a Metro Mirror or Global Mirror consistency group copy	152
Stopping a Metro Mirror or Global Mirror consistency group copy process	153
Deleting Metro Mirror and Global Mirror consistency groups	153
Creating Metro Mirror and Global Mirror partnerships.	154
Modifying Metro Mirror or Global Mirror partnerships.	154

Deleting Metro Mirror or Global Mirror partnerships	154	Migrating a VDisk to a new I/O group using the CLI	192
Viewing the feature log	155	Replacing a faulty node in the cluster using the CLI	193
Viewing and updating license settings	155	Recovering from offline VDIsks using the CLI	196
Running the cluster maintenance procedure	155	Recovering a node and returning it to the cluster using the CLI	197
Modifying error notification settings.	155	Moving offline VDIsks to the recovery I/O group using the CLI	198
Displaying and saving log and dump files.	156	Moving offline VDIsks to their original I/O group using the CLI	198
Analyzing the error log	157	Informing the SAN Volume Controller of changes to host HBAs using the CLI	199
Recovering a node and returning it to the cluster	158	Expanding VDIsks.	200
Managing SSH keys	159	Expanding a VDisk that is mapped to an AIX host	200
Adding SSH keys for hosts other than the master console	160	Expanding a VDisk that is mapped to a Windows 2000 host using the CLI	200
Adding subsequent SSH public keys to the SAN Volume Controller.	160	Shrinking a virtual disk using the CLI	202
Replacing the SSH key pair.	160	Migrating extents using the CLI	202
Resetting a refused SSH key	162	Migrating VDIsks between MDisk groups using the CLI.	204
Resetting the SSH fingerprint	162	Migrating a VDisk between I/O groups using the CLI.	206
Chapter 8. Using the CLI	165	Creating an image mode VDisk using the CLI	206
Preparing the SSH client system for the CLI	165	Migrating to an image mode virtual disk using the CLI.	207
Preparing the SSH client system to issue CLI commands	166	Deleting a node from a cluster using the CLI.	208
Preparing the SSH client on an AIX host	167	Performing the cluster maintenance procedure using the CLI	209
Issuing CLI commands from a PuTTY SSH client system	168	Changing the cluster IP address using the CLI	209
Running the PuTTY and plink utilities	168	Changing the cluster gateway address using the CLI.	210
Starting a PuTTY session for the CLI	170	Changing the cluster subnet mask using the CLI	210
Setting the cluster time using the CLI	170	Maintaining SSH keys using the CLI	210
Reviewing and setting the cluster features using the CLI	171	Setting up SNMP error notifications using the CLI	211
Displaying cluster properties using the CLI	171	Setting up e-mail error notifications using the CLI	212
Maintaining passwords for the front panel using the CLI	172	Changing cluster passwords using the CLI	212
Adding nodes to a cluster using the CLI	172	Changing the language setting using the CLI.	213
Displaying node properties using the CLI	176	Viewing the feature log using the CLI	213
Discovering MDIsks using the CLI	176	Analyzing the error log using the CLI	213
Creating MDisk groups using the CLI	178	Shutting down a cluster using the CLI	214
Adding MDIsks to MDisk groups using the CLI	180		
Creating VDIsks	181	Chapter 9. Backing up and restoring the cluster configuration	217
Creating host objects using the CLI	182	Backing up the cluster configuration.	217
Creating VDisk-to-host mappings using the CLI	183	Backing up the cluster configuration using the CLI	219
Modifying the amount of memory available for Copy Service features using the CLI.	183	Downloading backup configuration data files	221
Creating FlashCopy mappings using the CLI.	184	Restoring the cluster configuration using the CLI	221
Creating a FlashCopy consistency group and adding mappings using the CLI	185	Deleting backup configuration files	223
Preparing and starting a FlashCopy mapping using the CLI	186	Deleting backup configuration files using the CLI	224
Preparing and starting a FlashCopy consistency group using the CLI	187		
Determining the WWPNs of a node using the CLI	188	Chapter 10. Upgrading the SAN Volume Controller software	225
Determining the VDisk name from the device identifier on the host	189	Installing or upgrading the SAN Volume Controller software	225
Determining the host that a VDisk is mapped to	189	Copying the SAN Volume Controller software upgrade files using PuTTY scp	226
Determining the relationship between VDIsks and MDIsks using the CLI	190	Upgrading the SAN Volume Controller software automatically	227
Determining the relationship between MDIsks and RAID arrays or LUNs using the CLI.	190		
Increasing the size of your cluster using the CLI	191		
Adding a node to increase the size of a cluster using the CLI	191		

Upgrading the SAN Volume Controller software using the SAN Volume Controller Console	228
Upgrading the SAN Volume Controller software using the CLI	230
Performing a disruptive software upgrade using the CLI	231
Performing the node rescue	232
Recovering from software upgrade problems automatically	233
Recovering from software upgrade problems manually	234

Chapter 11. Configuring and servicing storage subsystems 235

Identifying your storage subsystem	235
Configuration guidelines for storage subsystems	235
Logical disk configuration guidelines for storage subsystems	236
RAID array configuration guidelines for storage subsystems	236
Optimal MDisk group configuration guidelines for storage subsystems	237
FlashCopy mapping guidelines for storage subsystems	237
Image mode VDisks and data migration guidelines for storage subsystems	238
Configuring a balanced storage subsystem	241
Discovering logical units	244
Expanding a logical unit using the CLI	245
Modifying a logical unit mapping using the CLI	245
Accessing controller devices with multiple remote ports	246
Determining a storage subsystem name from its SAN Volume Controller name	247
Determining a storage subsystem name from its SAN Volume Controller name using the CLI	248
Renaming a storage subsystem	248
Renaming a storage subsystem using the CLI	248
Changing a configuration for an existing storage subsystem using the CLI	249
Adding a new storage controller to a running configuration	249
Adding a new storage controller to a running configuration using the CLI	250
Removing a storage subsystem	251
Removing a storage subsystem using the CLI	253
Removing MDisks that represent unconfigured LUs using the CLI	254
Creating a quorum disk	254
Manual discovery	255
Servicing storage subsystems	255
Configuring Bull FDA subsystems	256
Supported firmware levels for the Bull FDA	256
Logical unit creation and deletion for Bull FDA	256
Platform type for Bull FDA	256
Access control methods for Bull FDA	256
Setting cache allocations for Bull FDA	257
Snapshot Volume and Link Volume for Bull FDA	257
Configuring the EMC CLARiiON subsystem	257
Access Logix	257

Configuring the EMC CLARiiON controller with Access Logix installed	258
Configuring the EMC CLARiiON controller without Access Logix installed	260
Supported models of the EMC CLARiiON	261
Supported firmware levels for the EMC CLARiiON	261
Concurrent maintenance on the EMC CLARiiON	261
User interface on EMC CLARiiON	262
Sharing the EMC CLARiiON between a host and the SAN Volume Controller	262
Switch zoning limitations for the EMC CLARiiON	262
Quorum disks on the EMC CLARiiON	263
Advanced functions for the EMC CLARiiON	263
Logical unit creation and deletion on the EMC CLARiiON	264
Configuring settings for the EMC CLARiiON	264
Configuring the EMC Symmetrix and Symmetrix DMX subsystems	266
Supported models of the EMC Symmetrix and Symmetrix DMX controllers	266
Supported firmware levels for the EMC Symmetrix and Symmetrix DMX	266
Concurrent maintenance on the EMC Symmetrix and Symmetrix DMX	267
User interfaces on EMC Symmetrix and Symmetrix DMX	267
Sharing the EMC Symmetrix or Symmetrix DMX between a host and a SAN Volume Controller	268
Switch zoning limitations for the EMC Symmetrix and Symmetrix DMX	268
Quorum disks on EMC Symmetrix and Symmetrix DMX	269
Advanced functions for EMC Symmetrix and Symmetrix DMX	269
LU creation and deletion on EMC Symmetrix and Symmetrix DMX	269
Configuring settings for the EMC Symmetrix and Symmetrix DMX	270
Configuring the Fujitsu ETERNUS subsystems	272
Supported models of the Fujitsu ETERNUS	272
Supported firmware levels for the Fujitsu ETERNUS	272
User interfaces on the Fujitsu ETERNUS	272
Configuring the Fujitsu ETERNUS to use with the SAN Volume Controller	272
Zoning configuration for the Fujitsu ETERNUS	275
Migrating logical units from the Fujitsu ETERNUS to the SAN Volume Controller	275
Concurrent maintenance on the Fujitsu ETERNUS	275
Advanced functions for the Fujitsu ETERNUS	276
Configuring the IBM TotalStorage ESS subsystem	276
Configuring the IBM ESS	276
Supported models of the IBM ESS	277
Supported firmware levels for the IBM ESS	277
Concurrent maintenance on the IBM ESS	277
User interface on the IBM ESS	277

Sharing the IBM ESS between a host and the SAN Volume Controller	278	Supported firmware levels for HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS	296
Switch zoning limitations for the IBM ESS.	278	Concurrent maintenance on the HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS	296
Quorum disks on the IBM ESS	278	User interface on the HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS	296
Advanced functions for the IBM ESS	278	Sharing the HDS Thunder, HDS TagmaStore AMS, or HDS TagmaStore WMS between a host and the SAN Volume Controller	297
Logical unit creation and deletion on the IBM ESS.	278	Setting up an HDS Thunder with more than four ports	298
Configuring the IBM System Storage DS4000 (formerly FAStT) series subsystem	279	Quorum disks on HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS	298
Configuring IBM DS4000 series disk controllers for the storage server.	279	Advanced functions for HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS	298
Supported options of the IBM DS4000 series controller.	280	Logical unit creation and deletion on HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems	299
Supported models of the IBM DS4000 series of controllers	281	Configuring settings for HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems	300
Supported firmware levels for the IBM DS4000 series	282	Configuring the HDS TagmaStore USP and NSC subsystems	305
Concurrent maintenance on the IBM DS4000 series	282	Supported models of the HDS USP and NSC	305
User interface on the IBM DS4000 series	282	Supported firmware levels for HDS USP and NSC	305
Sharing the IBM DS4000 series controller between a host and the SAN Volume Controller.	282	User interface on the HDS USP and NSC	306
Quorum disks on the IBM DS4000 series	283	Logical units and target ports on the HDS USP and NSC	306
Advanced functions for the IBM DS4000 series	283	Switch zoning limitations for the HDS USP and NSC	306
Logical unit creation and deletion on the IBM DS4000 series	284	Concurrent maintenance on the HDS USP and NSC	307
Configuration interface for the IBM DS4000 series	284	Quorum disks on HDS USP and NSC	307
Controller settings for the IBM DS4000 series	284	Host type for HDS USP and NSC	307
Configuring the IBM System Storage DS6000 subsystem	286	Advanced functions for HDS USP and NSC	307
Configuring the IBM DS6000	287	Configuring HP StorageWorks MA and EMA subsystems	308
Supported firmware levels for the IBM DS6000	288	HP MA and EMA definitions	308
Supported models of the IBM DS6000 series	288	Configuring HP MA and EMA subsystems	311
User interfaces on the IBM DS6000	288	Supported models of HP MA and EMA subsystems	314
Concurrent maintenance on the IBM DS6000	288	Supported firmware levels for HP MA and EMA subsystems	314
Target port groups on the IBM DS6000	288	Concurrent maintenance on the HP MA and EMA	314
Configuring the IBM System Storage DS8000 subsystem	289	Configuration interface for the HP MA and EMA	315
Configuring the IBM DS8000	289	Sharing the HP MA or EMA between a host and a SAN Volume Controller	315
Supported firmware levels for the IBM DS8000	290	Switch zoning limitations for HP MA and EMA	316
Supported models of the IBM DS8000	290	Quorum disks on HP MA and EMA subsystems	316
User interfaces on the IBM DS8000	290	Advanced functions for HP MA and EMA.	316
Concurrent maintenance for the IBM DS8000	290	SAN Volume Controller advanced functions	317
Configuring the HDS Lightning series subsystem	290	LU creation and deletion on the HP MA and EMA	317
Supported models of the HDS Lightning	291	Configuring settings for the HP MA and EMA	318
Supported firmware levels for HDS Lightning	291	Configuring the HP StorageWorks EVA subsystem	321
Concurrent maintenance on the HDS Lightning	291	Supported models of the HP EVA	322
User interface on HDS Lightning	291		
Sharing the HDS Lightning 99xxV between a host and the SAN Volume Controller	291		
Quorum disks on HDS Lightning 99xxV	292		
Advanced functions for HDS Lightning	292		
Logical unit configuration for HDS Lightning	293		
Configuring settings for HDS Lightning	294		
Configuring the HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS	296		
Supported models of the HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS	296		

Supported firmware levels for HP EVA	322
Concurrent maintenance on the HP EVA	322
User interface on HP EVA	322
Sharing the HP EVA controller between a host and the SAN Volume Controller	322
Switch zoning limitations for the HP EVA subsystem	322
Quorum disks on HP EVA	323
Advanced functions for HP EVA	323
Logical unit configuration on the HP EVA	323
Logical unit presentation	323
Configuration interface for the HP EVA	324
Configuring settings for the HP EVA	324
Configuring HP MSA subsystems	325
Supported models of the HP MSA subsystem	325
Supported firmware levels for the HP MSA	325
User interfaces on the HP MSA	325
Logical unit creation, deletion and migration for HP MSA	326
Sharing the HP MSA between a host and the SAN Volume Controller	327
Concurrent maintenance on the HP MSA	327
Quorum disks on the HP MSA	328
Advanced functions for the HP MSA	328
Global settings for the HP MSA	328
Configuring NEC iStorage subsystems	328
Supported firmware levels for the NEC iStorage	328
Logical unit creation and deletion for NEC iStorage	328
Platform type for NEC iStorage	328
Access control methods for NEC iStorage	329
Setting cache allocations for NEC iStorage	329
Snapshot Volume and Link Volume for NEC iStorage	329
Configuring NetApp FAS subsystems	329
Supported models of the NetApp FAS subsystem	329
Supported firmware levels for the NetApp FAS	330
User interfaces on the NetApp FAS	330
Logical units and target ports on the NetApp FAS	330
Creating logical units on the NetApp FAS	330
Deleting logical units on the NetApp FAS	331
Creating host objects for the NetApp FAS	331
Presenting LUNs to hosts for NetApp FAS	332
Switch zoning limitations for the NetApp FAS	332
Concurrent maintenance on the NetApp FAS	333
Quorum disks on the NetApp FAS	333
Advanced functions for the NetApp FAS	333

Chapter 12. Maintaining the master console software 335

Upgrading the master console software.	335
Prerequisites for upgrading the master console	335
Upgrading using the master console installation wizard	336
Upgrading using the downloaded SAN Volume Controller Console installation program	338
Verifying the Windows services associated with the SAN Volume Controller Console.	349

Starting your initial SAN Volume Controller Console session.	349
Uninstalling master console software	352
Uninstalling IBM Director	353
Uninstalling Tivoli SAN Manager Agent	353
Uninstalling Tivoli SAN Manager	353
Uninstalling the DS4000 Storage Manager Client (FAStT Storage Manager Client)	354
Uninstalling the SAN Volume Controller Console	354
Uninstalling PuTTY	356
Uninstalling the master console	356
Changing the master console host name	356

Chapter 13. Troubleshooting the master console 359

Clearing the Microsoft Windows event logs	359
Troubleshooting unexpected shutdowns of the SAN Volume Controller Console	359
Troubleshooting Microsoft Windows boot problems	360
Starting the master console hardware from the mirrored disk	360
Replace a disk on the master console hardware	360

Chapter 14. IBM System Storage support for Microsoft Volume Shadow Copy Service 363

Installation overview	363
System requirements for the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software	364
Installing the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software	364
Creating the free and reserved pools of volumes	369
Verifying the installation	370
Changing the configuration parameters.	370
Adding and removing volumes	371
Error codes	372
Uninstalling the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software	374

Appendix A. Replacing nodes nondisruptively 375

Appendix B. Replacing nodes disruptively (rezoning the SAN) 379

Appendix C. Replacing nodes disruptively (moving VDisks to new I/O group) 381

Appendix D. Error Codes 383

Appendix E. Event codes	391
Information event codes.	391
Configuration event codes	392

Appendix F. SCSI error reporting . . .	397	Trademarks	407
Appendix G. Object types	401	Glossary	409
Accessibility	403	Index	421
Notices	405		

Figures

1.	Levels of virtualization	3	13.	Disk controller system shared between SAN Volume Controller and a host	70
2.	Asymmetrical virtualization	4	14.	IBM ESS LUs accessed directly with a SAN Volume Controller	71
3.	Symmetrical virtualization	5	15.	IBM DS4000 direct connection with a SAN Volume Controller on one host	72
4.	Configuration node	11	16.	Fabric with ISL between nodes in a cluster	76
5.	I/O group and UPS.	12	17.	Fabric with ISL in a redundant configuration	77
6.	Controllers and MDisks	18	18.	Basic frame layout	103
7.	MDisk group	20	19.	Task bar	104
8.	MDisk groups and VDIs	23	20.	Node-rescue-request display	233
9.	Hosts, WWPNs, and VDIs	27			
10.	Hosts, WWPNs, VDIs and SCSI mappings	28			
11.	Redundant fabrics	54			
12.	Example of a SAN Volume Controller cluster in a fabric	67			

Tables

1. Node state	10	29. HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems global settings supported by the SAN Volume Controller.	300
2. MDisk status	19	30. HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystem port settings supported by the SAN Volume Controller.	302
3. MDisk group status.	21	31. HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems LU settings for the SAN Volume Controller.	303
4. Capacities of the cluster given extent size	22	32. Determining partition usage	312
5. VDisk states	24	33. HSG80 controller container types for LU configuration	318
6. VDisk cache modes	25	34. HP MA and EMA global settings supported by the SAN Volume Controller.	318
7. FlashCopy mapping events	42	35. HSG80 controller settings that are supported by the SAN Volume Controller.	319
8. Intercluster heartbeat traffic in Mbps	54	36. HSG80 controller port settings supported by the SAN Volume Controller.	319
9. Configuration terms and definitions	65	37. HSG80 controller LU settings supported by the SAN Volume Controller.	320
10. Four hosts and their ports.	80	38. HSG80 connection default and required settings	321
11. Six hosts and their ports	81	39. HP EVA global settings supported by the SAN Volume Controller	324
12. Extent size	179	40. HP EVA LU settings supported by the SAN Volume Controller	325
13. Calculate the I/O rate.	241	41. HP EVA host settings supported by the SAN Volume Controller	325
14. Calculate the impact of FlashCopy mappings	242	42. Unsupported components and actions to take prior to upgrading.	335
15. Determine if the storage subsystem is overloaded	243	43. Configuration commands	370
16. Controller device port selection algorithm	247	44. Pool management commands	372
17. EMC CLARiiON global settings supported by the SAN Volume Controller.	264	45. Error messages for the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software	373
18. EMC CLARiiON controller settings supported by the SAN Volume Controller.	265	46. Error codes	383
19. EMC CLARiiON port settings supported by the SAN Volume Controller.	265	47. Information event codes	391
20. EMC CLARiiON LU settings supported by the SAN Volume Controller.	265	48. Configuration event codes	392
21. EMC Symmetrix and Symmetrix DMX global settings supported by the SAN Volume Controller.	270	49. SCSI status	397
22. EMC Symmetrix and Symmetrix DMX port settings supported by the SAN Volume Controller.	270	50. SCSI sense keys codes and qualifiers	398
23. EMC Symmetrix and Symmetrix DMX LU settings supported by the SAN Volume Controller.	271	51. Reason codes	399
24. IBM DS4000 series controller global settings supported by the SAN Volume Controller	286	52. Object types	401
25. HDS Lightning global settings supported by the SAN Volume Controller.	294		
26. HDS Lightning controller settings that are supported by the SAN Volume Controller	295		
27. HDS Lightning port settings supported by the SAN Volume Controller	295		
28. HDS Lightning LU settings for the SAN Volume Controller	295		

About this guide

The IBM System Storage SAN Volume Controller Configuration Guide provides information that helps you configure and use the IBM System Storage SAN Volume Controller.

The IBM System Storage 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 SAN Volume Controller.

Who should use this guide

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

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

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 SC23-6628-01 SAN Volume Controller Software Installation and 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, SC23-6628-00. The following sections summarize the changes that have since been implemented from the previous version.

This version includes the following new information:

- Internet Explorer 7.0 is now supported
- Incremental FlashCopy mappings are now supported
- Cascaded FlashCopy mappings are now supported
- The maximum extent size for managed disk groups is now 2048 MB
- Cluster addressability is now 8 PB

Summary of changes for SC23-6628-00 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, SC26-7902-00. The following sections summarize the changes that have since been implemented from the previous version.

This version includes the following new information:

- Added the following new chapters:
 - Configuring the master console
 - Maintaining the master console software
- Added the following new sections:
 - Configuring NEC iStorage subsystems
 - Configuring Fujitsu ETERNUS subsystems
 - Configuring Bull FDA subsystems
 - Configuring HP MSA subsystems
- Multiple target FlashCopy mappings are now supported
- Mesh configurations are now supported

Changed information

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

- There is a new SAN Volume Controller supported model. The SAN Volume Controller is now documented by model number. For example, this publication states four SAN Volume Controller model types: the SAN Volume Controller 2145-4F2, the SAN Volume Controller 2145-8F2, SAN Volume Controller 2145-8F4 and the new SAN Volume Controller 2145-8G4.

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 any of the SAN Volume Controller models. When the SAN Volume Controller is referred to as the SAN Volume Controller 2145-4F2, SAN Volume Controller 2145-8F2, SAN Volume Controller 2145-8F4, or the SAN Volume Controller 2145-8G4, the specific SAN Volume Controller is designated.

- The *IBM System Storage SAN Volume Controller Configuration Guide* is now titled *IBM System Storage SAN Volume Controller: Software Installation and Configuration Guide*.
- The *IBM System Storage SAN Volume Controller Installation Guide* is now titled *IBM System Storage SAN Volume Controller: Hardware Installation Guide*.
- The *IBM System Storage Master Console for SAN Volume Controller: Installation and User's Guide* and the *IBM System Storage Master Console for SAN Volume Controller Information Center* are no longer updated and distributed. Instead, all pertinent information from those information units has been incorporated into other SAN Volume Controller publications.
- Support for IBM N7000 series
- Support for TagmaStore Adaptable Modular Storage
- Support for TagmaStore Workgroup Modular Storage

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 System Storage SAN Volume Controller
- Other IBM publications that relate to the SAN Volume Controller

SAN Volume Controller library

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

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

Title	Description	Order number
<i>IBM System Storage SAN Volume Controller: CIM Agent Developer's Reference</i>	This reference guide describes the objects and classes in a Common Information Model (CIM) environment.	SC26-7904

Title	Description	Order number
<i>IBM System Storage SAN Volume Controller: Command-Line Interface User's Guide</i>	This guide describes the commands that you can use from the SAN Volume Controller command-line interface (CLI).	SC26-7903
<i>IBM System Storage SAN Volume Controller: Software Installation and Configuration Guide</i>	This guide provides guidelines for configuring your SAN Volume Controller.	SC23-6628
<i>IBM System Storage SAN Volume Controller: Host Attachment Guide</i>	This guide provides guidelines for attaching the SAN Volume Controller to your host system.	SC26-7905
<i>IBM System Storage SAN Volume Controller: Hardware Installation Guide</i>	This guide includes the instructions that the IBM service representative uses to install the SAN Volume Controller.	GC27-2132
<i>IBM System Storage SAN Volume Controller: Planning Guide</i>	This guide introduces the SAN Volume Controller and lists the features you can order. It also provides guidelines for planning the installation and configuration of the SAN Volume Controller.	GA32-0551
<i>IBM System Storage SAN Volume Controller: Service Guide</i>	This guide includes the instructions that the IBM service representative uses to service the SAN Volume Controller.	GC26-7901
<i>IBM Systems Safety Notices</i>	This guide contains translated caution and danger statements. Each caution and danger statement in the SAN Volume Controller documentation has a number that you can use to locate the corresponding statement in your language in the <i>IBM Systems Safety Notices</i> document.	G229-9054

Other IBM publications

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

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

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

Title	Description	Order number
<i>IBM System Storage Multipath Subsystem Device Driver: User's Guide</i>	This guide describes the IBM System Storage Multipath Subsystem Device Driver Version 1.6 for TotalStorage Products and how to use it with the SAN Volume Controller. This publication is referred to as the <i>IBM System Storage Multipath Subsystem Device Driver: User's Guide</i> .	GC27-2122
<i>IBM TotalStorage DS4300 Fibre Channel Storage Subsystem Installation, User's, and Maintenance Guide</i>	This guide describes how to install and configure the IBM TotalStorage DS4300 Fibre-Channel Storage Subsystem.	GC26-7722
<i>IBM eServer xSeries 306m (Types 8849 and 8491) Installation Guide</i>	This guide describes how to install the IBM eServer xSeries 306m, which is the hardware delivered for some versions of the hardware master console.	MIGR-61615
<i>IBM xSeries 306m (Types 8849 and 8491) User's Guide</i>	This guide describes how to use the IBM eServer xSeries 306m, which is the hardware delivered for some versions of the hardware master console.	MIGR-61901
<i>IBM xSeries 306m (Types 8849 and 8491) Problem Determination and Service Guide</i>	This guide can help you troubleshoot and resolve problems with the IBM eServer xSeries 306m, which is the hardware delivered for some versions of the hardware master console.	MIGR-62594
<i>IBM eServer xSeries 306 (Type 8836) Installation Guide</i>	This guide describes how to install the IBM eServer xSeries 306, which is the hardware delivered for some versions of the hardware master console.	MIGR-55080
<i>IBM eServer xSeries 306 (Type 8836) User's Guide</i>	This guide describes how to use the IBM eServer xSeries 306, which is the hardware delivered for some versions of the hardware master console.	MIGR-55079
<i>IBM eServer xSeries 306 (Types 1878, 8489 and 8836) Hardware Maintenance Manual and Troubleshooting Guide</i>	This guide can help you troubleshoot problems and maintain the IBM eServer xSeries 306, which is the hardware delivered for some versions of the hardware master console.	MIGR-54820

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

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

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

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

Related Web sites

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

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

How to order IBM publications

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

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

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

How to send your comments

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

- e-mail

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

starpubs@us.ibm.com

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

- Mail

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

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

Chapter 1. SAN Volume Controller overview

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

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

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

The SAN Volume Controller is analogous to a logical volume manager 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)
 - Global Mirror (asynchronous copy)
 - Data migration
 - Space management
 - Mapping that is based on desired performance characteristics
 - Metering of service quality

Each SAN Volume Controller node is a rack-mounted unit that you can install in a standard Electrical Industries Alliance (EIA) 19-inch rack. The nodes are always installed in pairs, with one-to-four pairs of nodes constituting a *cluster*. Each pair of nodes is known as an *I/O group*.

All I/O operations that are managed by the nodes in an I/O group are cached on both nodes. Each virtual volume is defined to an I/O group. I/O groups take the storage that is presented to the SAN by the storage subsystems as MDisks and translates the storage into logical disks, known as VDIs, that are used by applications on the hosts. Each node must reside in only one I/O group and provide access to the VDIs in that I/O group.

There are four models of SAN Volume Controller nodes:

- SAN Volume Controller 2145-8G4

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

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 RAIDs 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. Virtualization also provides the capability to move storage between storage subsystems transparently to host systems.

Types of virtualization

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

Asymmetric

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

Symmetric

A virtualization engine sits in the data path 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.

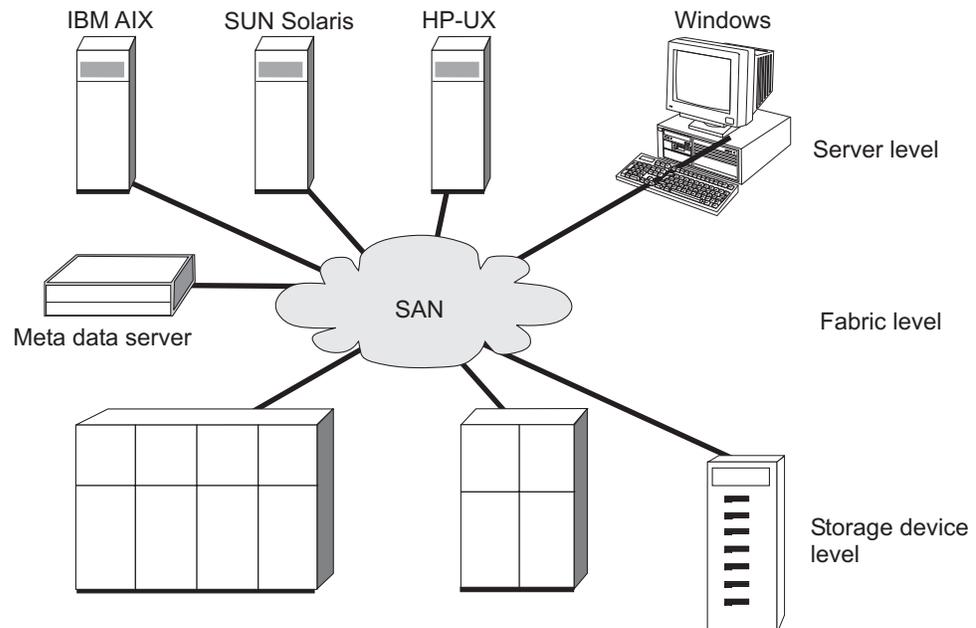


Figure 1. Levels of virtualization

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 2, 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.

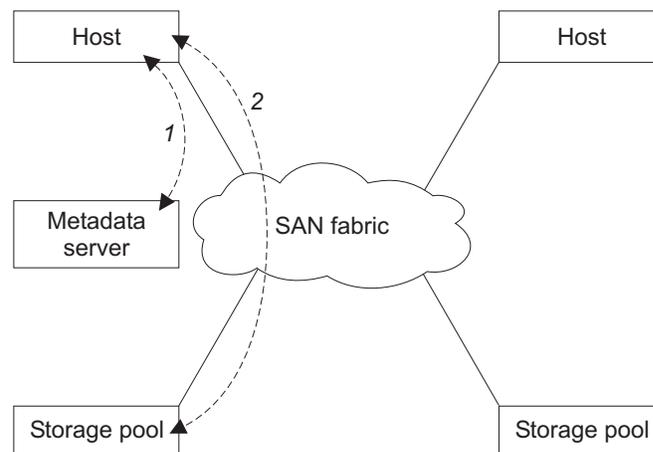


Figure 2. Asymmetrical virtualization

Symmetric virtualization

The SAN Volume Controller provides symmetric virtualization.

Virtualization splits the storage that is presented by the storage subsystems into smaller chunks 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 3 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.

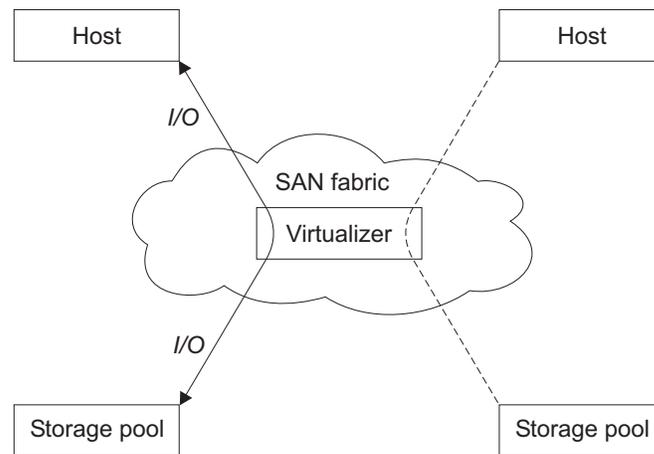


Figure 3. Symmetrical virtualization

The virtualization engine directly controls access to the storage and to the data that is written to the storage. As a result, locking functions that provide data integrity and advanced functions, such as cache and Copy Services, can be run in the virtualization engine itself. Therefore, the virtualization engine is a central point of control for device and advanced function 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.

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 requirements:

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

Note: You can order the master console for the SAN Volume Controller as a master console hardware option that comes preloaded with the master console software or a master console software option that you install on your own hardware.

Features of a SAN Volume Controller 2145-8G4 node

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

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

Supported hosts

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 through the master console:

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

Application programming interfaces

The SAN Volume Controller 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.

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 backend 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.

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.

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
- Global Mirror relationships
- Metro Mirror consistency groups
- Global Mirror consistency groups

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
- Processing

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 processing phase, the command sequence is run.

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 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.	

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 4 on page 11 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.

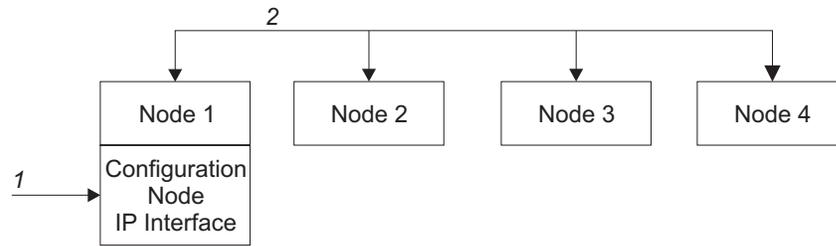


Figure 4. Configuration node

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*. The *I/O groups* are connected to the SAN so that all backend storage and all application servers are visible to all of the *I/O groups*. Each pair 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*. Nodes do 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 nodes 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 5 on page 12 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 5 also shows two UPS units correctly configured so that each node is in a different power domain.

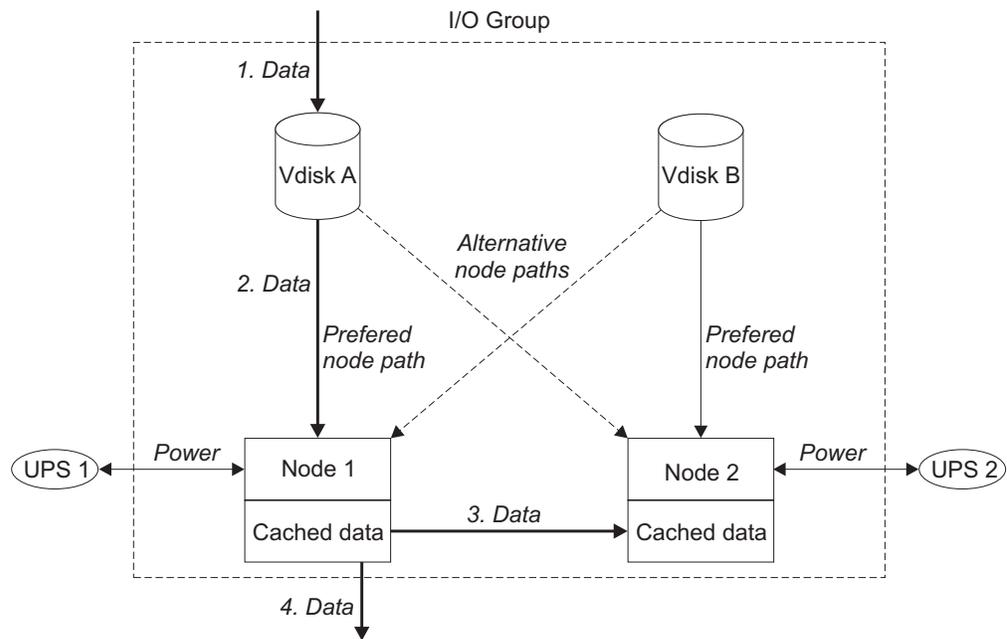


Figure 5. I/O group and UPS

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 and 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.

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 and Global Mirror primary and secondary VDisks as follows:

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

Using the 2145 UPS-1U

The 2145 uninterruptible power supply-1U (2145 UPS-1U) provides a SAN Volume Controller node 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.

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 internal disk of the SAN Volume Controller node. The UPS units are required to power the SAN Volume Controller nodes even if the input power source is considered uninterruptible.

The SAN Volume Controller 2145-8G4, SAN Volume Controller 2145-8F4, and SAN Volume Controller 2145-8F2 nodes can operate only with the 2145 UPS-1U. The SAN Volume Controller 2145-4F2 node can operate with either the 2145 UPS or the 2145 UPS-1U.

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

2145 UPS-1U configuration:

A 2145 UPS-1U powers one SAN Volume Controller node. All SAN Volume Controller model types are supported by the 2145 UPS-1U.

To make the SAN Volume Controller cluster more resilient to power failure, the 2145 UPS-1Us can be connected to the redundant ac power switch. If a redundant ac power switch is not used, connecting the two UPSs that are powering an I/O group to different, independent electrical power sources allows the SAN Volume Controller cluster to continue to operate with reduced capacity if a single power source fails.

Each UPS must be in the same rack as the node it powers.

Attention: Do not connect the UPSs to an input power source that does not conform to standards.

Each UPS includes power (line) cords that connect the UPS to either a redundant ac power switch, if one exists, a rack power distribution unit (PDU), if one exists, or to an external power source.

Each UPS is connected to a SAN Volume Controller node 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 cable enables the SAN Volume Controller node to read status and identification information from the UPS.

2145 UPS-1U operation:

Each SAN Volume Controller node 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 node 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 Controller node restarts and restores the original contents of the DRAM from the data saved on the disk drive.

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

Important: Do not shut down a UPS without first shutting down the SAN Volume Controller node that it supports. Data integrity can be compromised by pushing the 2145 UPS-1U on/off button when the node is still operating. However, in the case of an emergency, you can manually shut down the UPS by pushing the 2145 UPS-1U on/off button when the node is still operating. Service actions must then be performed before the node can resume normal operations. If multiple UPSs are shut down before the nodes they support, data can be corrupted.

Split cluster configurations

When possible, avoid using a configuration that splits a single cluster between two different physical locations. Split cluster configurations can only provide asymmetric disaster recovery facilities with substantially reduced performance. If you want to use a split cluster configuration, contact your IBM Regional Advanced Technical Specialist for more information on how to set up this type of configuration.

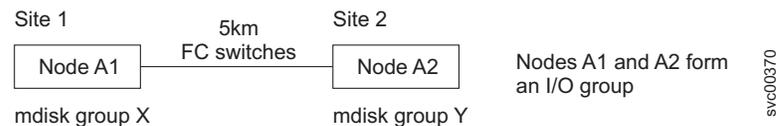
SAN Volume Controller cluster guidelines

- A cluster can be connected to application hosts, storage controllers, or other clusters through short wave or long wave optical fibre-channel connections with a distance of up to 300 m (short wave) or 10 km (long wave) between the cluster and the host, and other clusters and the storage controller. Longer distances are supported between clusters that use the intercluster Metro Mirror or Global Mirror features.
- For disaster recovery, a cluster is considered a single entity. This includes the backend storage that provides the quorum disks for the cluster. Collocate the cluster and the quorum disks. The components of a single cluster cannot be placed in different physical locations because this can cause maintenance, service, and quorum disk management issues.

- All nodes in a cluster should be located in close proximity to each other. Place all of the nodes for a cluster in the same set of racks and in the same room. You can have a large optical distance between the nodes in the same cluster. However, they must be physically collocated for service and maintenance purposes.
- All nodes in a cluster must be on the same IP subnet to allow the nodes in the cluster to assume the same cluster or service IP address.
- A node must be in the same rack as the UPS from which it is supplied.

Example configurations

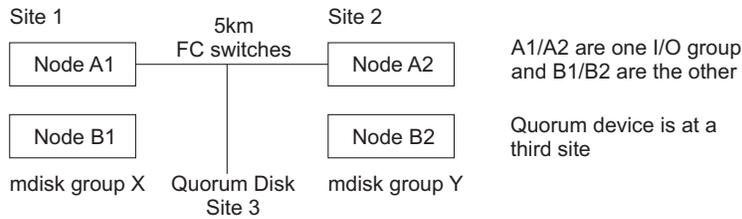
Example 1: One node from each I/O group is located at a primary site and a secondary site. The Metro Mirror or Global Mirror relationships are configured so that the primary VDisks at the primary site come from MDisks in group X and the secondary VDisks at the secondary site come from MDisks in group Y. This configuration appears to allow for disaster recovery. If the primary site fails, there is still a live I/O group in degraded mode at the secondary site that can perform the I/O workload.



This configuration introduces the following issues:

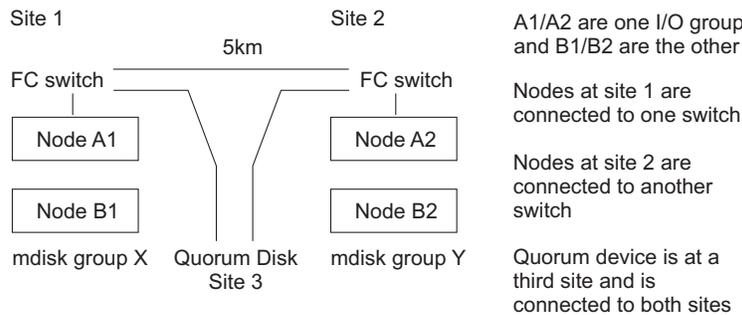
- If either site fails, only a degraded I/O group at the other site is available to continue I/O operations. As a result, performance is significantly impacted because the throughput of the cluster is reduced and the cluster caching is disabled.
- The disaster recovery solution is asymmetric. It is not possible to run applications on both the primary and secondary sites and allow either site to have a failure. One site must be regarded as the primary site and the other site must be regarded as a recovery site.
- If the quorum disk is at the secondary site and the primary site fails, the secondary site retains the quorum disk and can proceed to act as a disaster recovery site. If the secondary site fails, the primary site cannot act as a disaster recovery site because the primary site can only see half the nodes in the cluster and cannot see the quorum disk. The cluster components at the primary site cannot form an active cluster (error code 550). It is not possible to communicate with the nodes at the primary site in this state and all I/O operations immediately cease. An active cluster can only start operating at the primary site if the quorum disk reappears or if a node from the secondary site becomes visible.

Example 2: The quorum disks are located at a third site. A three site configuration has significant limitations because the cluster cannot change the path it uses to communicate with a quorum disk under all circumstances. To be tolerant of a single site failure, you must ensure that the path to the quorum disk from a node in one site does not go through a switch in the second site before it reaches the quorum disk in the third site.



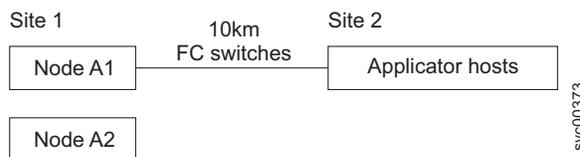
svc00371

Example 4: The nodes at the primary site are connected to one fibre-channel switch and the nodes at the secondary site are connected to a different fibre-channel switch. The quorum disks are located at a third site, but are connected to both the primary site and the secondary site.



svc00372

Example 5: Long distance connections that are 10 km are used for remote hosts or remote controllers. Both nodes in the I/O group are located at the primary site and the hosts are located at the secondary site. This configuration eliminates the issues that are discussed in examples 1 through 4.



svc00373

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.

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). To avoid this failure, do not configure your physical storage as RAID-0 arrays or JBODs.

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 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.

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 6 shows physical disks and MDisks.

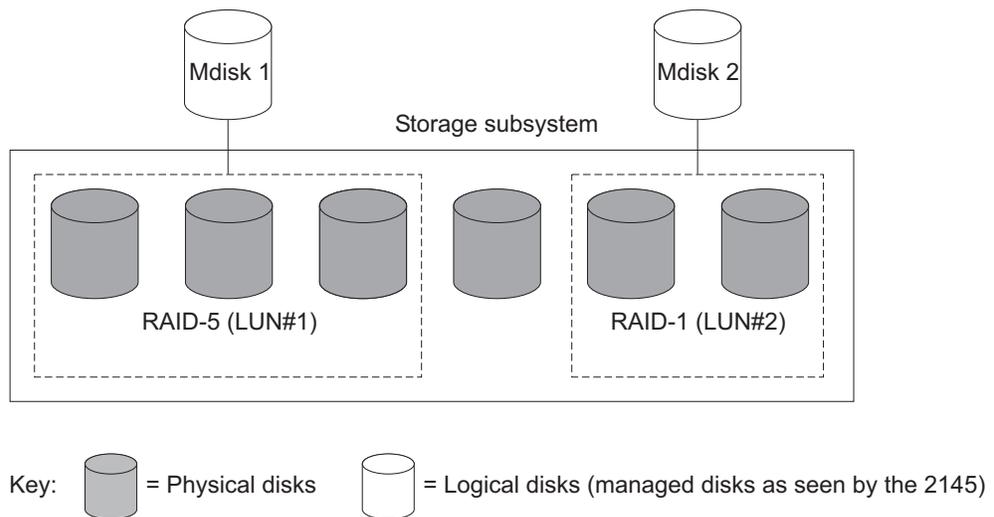


Figure 6. Controllers and MDisks

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.
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.

MDisk groups and VDisks

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

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 7 shows an MDisk group containing four MDisks.

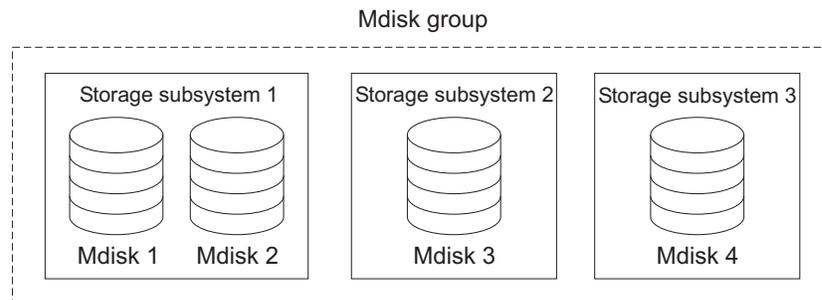


Figure 7. MDisk group

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 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:

- 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, 512, or 2048 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.
- Intracluster Metro Mirror or Global 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
2048 MB	8 PB

A cluster can manage 4 million extents ($4 \times 1024 \times 1024$). For example, with a 16 MB extent size, the cluster can manage up to $16 \text{ MB} \times 4 \text{ MB} = 64 \text{ TB}$ 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 64 MB, 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.

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 8 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.

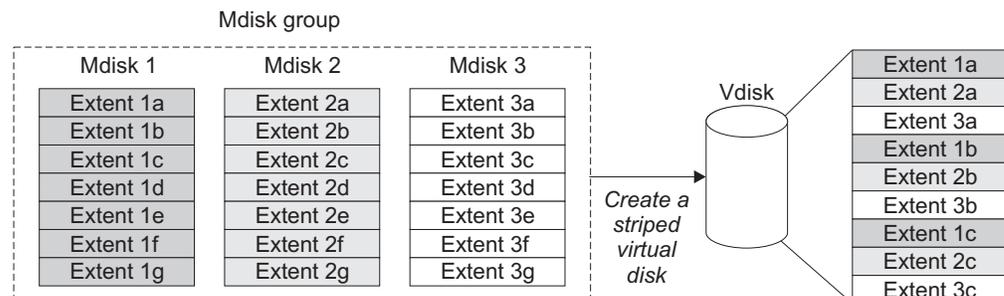


Figure 8. MDisk groups and VDisks

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 VDisks are special VDisks 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.

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, Metro Mirror or Global 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.

States

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 states

State	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. The VDisk can also be offline if the VDisk is the secondary of a Metro Mirror or Global Mirror relationship that is not synchronized.

Table 5. VDisk states (continued)

State	Description
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.

Cache modes

You can select to have read and write operations 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.

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 accesses the same storage, you can add HBA ports from several hosts to one host object to make a simpler configuration.

The cluster does not automatically present virtual disks (VDisks) on the fibre-channel ports. 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 VDIs are mapped to the host. The configuration interface provides a method that allows you to manually type the port names.

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.

Port masks

You can use a port mask to control the node target ports that a host can access. The port mask applies to logins from the host initiator port that are associated with the host object.

For each login between a host HBA port and node port, the node examines the port mask that is associated with the host object for which the host HBA is a member and determines if access is allowed or denied. If access is denied, the node responds to SCSI commands as if the HBA port is unknown.

The port mask is four binary bits. Valid mask values range from 0000 (no ports enabled) to 1111 (all ports enabled). For example, a mask of 0011 enables port 1 and port 2. The default value is 1111.

Multiple target ports

When you create a VDisk-to-host mapping, the host ports that are associated with the host object can see the LUN that represents the VDisk on up to eight fibre-channel ports. Nodes follow the ANSI FC standards for SCSI LUs that are accessed through multiple node ports. However, you must coordinate the nodes in an I/O group to present a consistent SCSI LU across all ports that can access it. The ANSI FC standards do not require that the same LUN is used on all ports; however, nodes always present the LU that represents a specific VDisk with the same LUN on all ports in an I/O group.

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 that expected for the current SAN zoning rules, then you may have a fabric problem.

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 cluster.

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 Controller nodes 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 VDisks are presented to the host. For example, if you present three VDisks to the host, and those VDisks 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 9 and Figure 10 on page 28 show two VDisks, and the mappings that exist between the host objects and these VDisks.

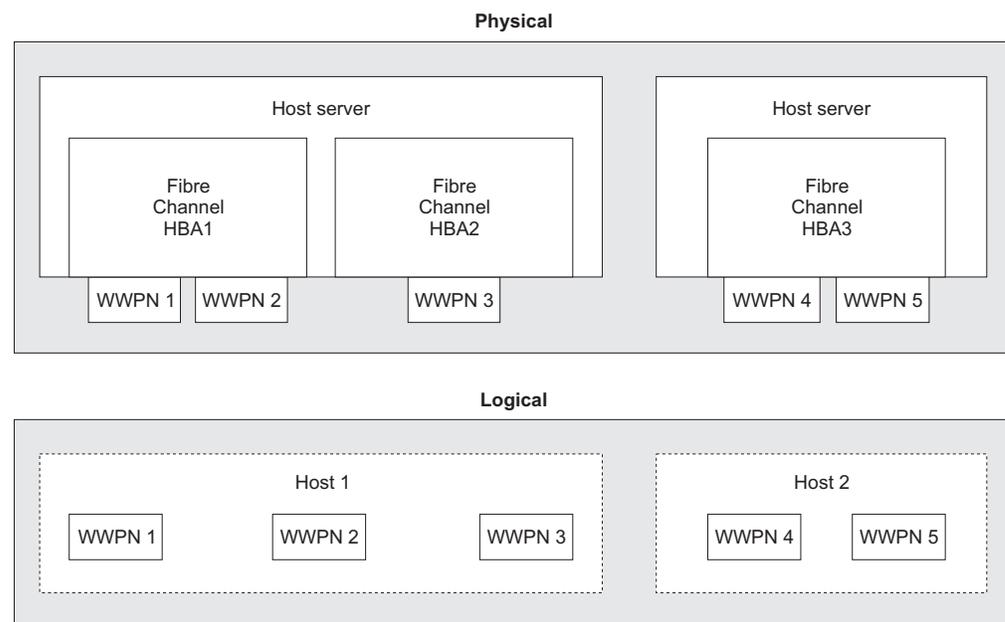


Figure 9. Hosts, WWPNs, and VDisks

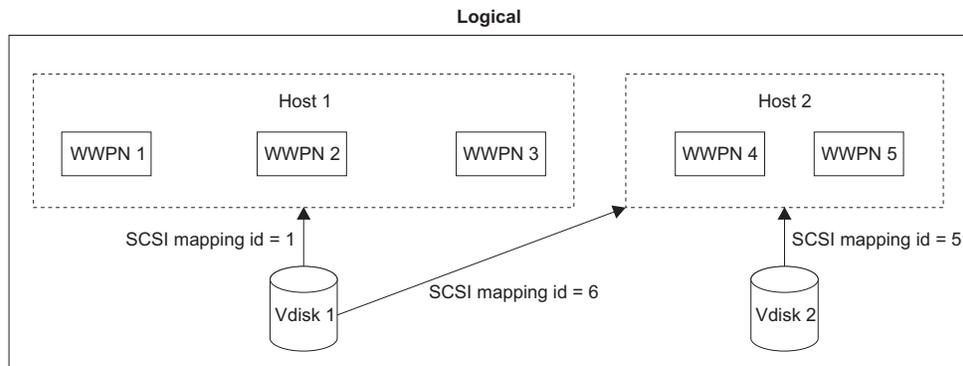


Figure 10. Hosts, WWPNs, VDisks and SCSI mappings

Standard and persistent reserves

The SCSI Reserve command and the SCSI Persistent Reserve command are specified by the SCSI standards. Servers can use these commands to prevent HBA ports in other servers from accessing the LUN.

This prevents accidental data corruption that is caused when a server overwrites data on another server. The Reserve and Persistent Reserve commands are often used by clustering software to control access to SAN Volume Controller virtual disks (VDisks).

If a server is not shut down or removed from the server cluster in a controlled way, the server reserves and persistent reserves are maintained. This prevents other servers from accessing data that is no longer in use by the server that holds the reservation. In this situation, you might want to release the reservation and allow a new server to access the VDisk.

When possible, you should have the server that holds the reservation explicitly release the reservation to ensure that the server cache is flushed and the server software is aware that access to the VDisk has been lost. In circumstances where this is not possible, you can use operating system specific tools to remove reservations. Consult the operating system documentation for details.

When you use the `svctask rmvdiskhostmap` CLI command or the SAN Volume Controller Console to remove VDisk-to-host mappings, SAN Volume Controller nodes with a software level of 4.1.0 or later can remove the server reservations and persistent reservations that the host has on the VDisk.

Maximum configuration

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

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

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

Node management and support tools

The SAN Volume Controller offers several management and support tools for you to maintain and manage your nodes.

The following node management tools are available with the SAN Volume Controller:

- Master console
- Secure Shell
- Assist On-site

Master console

The master console provides a single point from which to manage the SAN Volume Controller nodes. You can purchase the master console as a hardware product option (which includes the master console preinstalled software) or as a software-only option.

The two master console options are the same in function and software. However, the planning, installation, and configuration processes are slightly different:

Master console hardware option

The manufacturing plant installs the software on the hardware using the default settings. After the IBM service representative installs the hardware option, you must configure and customize the default factory settings.

Master console software-only option

You must provide your own hardware and perform both the installation and configuration processes.

The master console provides you with the following functions:

- A platform on which the subsystem configuration tools can be run
- A platform for remote service, which allows the desktop to be shared with remote IBM service personnel if assistance is required to resolve complex problems
- Access to the following components:
 - SAN Volume Controller Console, which is a graphical user interface application, through a Web browser
 - SAN Volume Controller command-line interface, through a Secure Shell (SSH) session

The master console can support up to four SAN Volume Controller clusters.

Master console software components

The master console feature is a collection of different software components.

The master console includes the following software components:

- SAN Volume Controller Console and CIM agent
- PuTTY (SSH client software)

Secure Shell protocol

Secure Shell (SSH) software is a client-server protocol that can be used from the master console or a host to enable you to control the SAN Volume Controller via a command-line interface (CLI).

The master console for the SAN Volume Controller distributes PuTTY, which provides an SSH client.

SSH provides a secure communications channel between systems. You can configure SSH to use a key pair (a private key and a public key) to establish the secure connection.

If you purchase the master console hardware option, you generate the SSH key pair when you configure the master console. If you purchase the master console software-only option, you generate the SSH key pair as part of the master console installation process. If you want to create an SSH connection (such as the SAN Volume Controller nodes), you must place the public key on every system.

Assist On-site and remote service

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

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

To use AOS, the master console must be able to access the Internet. The following Web site provides further information about AOS:

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

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

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

Sending notifications

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

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

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

Simple Network Management Protocol traps

Simple network management protocol (SNMP) is the standard protocol for managing networks and exchanging messages. SNMP enables the SAN Volume Controller to send external messages that notify personnel about an event. An SNMP manager allows you to view the messages that the SNMP agent sends. You can use the SAN Volume Controller Console or the SAN Volume Controller

command-line interface to configure and modify your SNMP settings. SNMP traps and Call Home e-mail can be sent simultaneously.

Call Home e-mail

The Call Home feature allows the transmission of operational and error-related data to you and IBM through a Simple Mail Transfer Protocol (SMTP) server connection in the form of an event notification e-mail. This function alerts IBM service personnel about machine conditions or sends data for error analysis and resolution.

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

Call Home support is initiated for the following reasons or types of data:

- Problem or event notification: Data is sent when there is a problem or event that might require the attention of IBM service personnel.
- Communication tests: You can test for the successful installation and communication infrastructure.
- Inventory information: A notification is sent to provide the necessary status and hardware information to IBM service personnel.

Call Home e-mails can contain any combination of the following types of information:

- Contact name
- Contact phone number
- Offshift phone number
- Machine location
- Record type
- Machine type
- Machine serial number
- Error ID
- Error code
- Software version
- FRU part number
- Cluster name
- Node ID
- Error sequence number
- Timestamp
- Object type
- Object ID
- Problem data

Inventory information e-mail

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

Inventory information that is sent to IBM can include the following information about the cluster on which the Call Home function is enabled:

- Timestamp
- Contact information, including name and phone number. This is initially set to the contact information that was set for the Call Home e-mail function. However, you can change the contact information specifically for inventory e-mail using the SAN Volume Controller Console or the **mkemailuser** or **chemailuser** CLI commands.
- Machine location. This is the machine location that is set for the Call Home e-mail function.
- Software level
- License information. This is the same information that it output from the **svcinfolicense** command.
- Cluster vital product data (VPD). The cluster VPD is the same information that is output from the **svcinfolcluster** command, including the following items:
 - Cluster name and IDs
 - Cluster location
 - Bandwidth
 - IP addresses
 - Memory capacities
 - SNMP settings
 - Time zone setting
 - E-mail settings
 - Microcode level
 - Fibre-channel port speed
- Node VPD for each node in the cluster. The node VPD is the same information that is output from the **svcinfolnodevpd** command, including the following items:
 - System part number
 - Number of various hardware parts, such as fans, processors, memory slots, fibre-channel cards, and SCSI/IDE devices
 - Part numbers of the various hardware parts
 - BIOS information
 - System manufacturing information, such as system product and manufacturer
 - Firmware level for the service processor
- Software VPD, including the following items:
 - Code level

- Node name
- Ethernet status
- Worldwide node name (WWNN)
- MAC address
- Processor information, including the following items for each processor:
 - Location of processor
 - Type of cache
 - Size of cache
 - Manufacturer
 - Version
 - Speed
 - Status (enabled or disabled)
- Memory information, including the following items:
 - Part number
 - Device location
 - Bank location
 - Size
- Fibre-channel card information, including the following items:
 - Part number
 - Port number
 - Device serial number
 - Manufacturer
- SCSI/IDE device information, including the following items:
 - Part number
 - Bus ID
 - Device ID
 - Model
 - Revision level
 - Serial number
 - Approximate capacity
- Front panel assembly information, including the following items:
 - Part number
 - ID
 - Location
- Universal power supply (UPS) information, including the following items:
 - Electronics part number
 - Battery part number
 - UPS assembly part number
 - Input power cable part number
 - UPS serial number
 - UPS type
 - UPS internal part number
 - ID
 - Firmware levels

Chapter 2. Copy Services features

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

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

FlashCopy

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

Metro Mirror

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

Global Mirror

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

FlashCopy

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

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

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

FlashCopy operations can occur on multiple source and target VDisks. FlashCopy management operations are coordinated to allow a common single point in time for copying target VDisks from their respective source VDisks. This allows a consistent copy of data that spans multiple VDisks. For SAN Volume Controller software level 4.2.0 or higher, the FlashCopy feature also allows multiple target VDisks to be copied from each source VDisk. This can be used to create images from different points in time for each source VDisk.

For SAN Volume Controller software level 4.2.1 or higher, the FlashCopy feature also allows a FlashCopy target VDisk to be the source VDisk of another FlashCopy mapping. SAN Volume Controller version 4.2.1 also introduces incremental FlashCopy, which potentially reduces the amount of time to complete the copy operation after the initial copy has completed. Only the differences are copied when the FlashCopy mapping restarts.

Source VDIsks and target VDIsks must meet the following requirements:

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

FlashCopy applications

You can use the FlashCopy feature to create consistent backups of dynamic data, test applications, and create copies for auditing purposes and for data mining.

To create consistent backups of dynamic data, use the FlashCopy feature 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.

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.

You can also use the FlashCopy feature to create restart points for long running batch jobs. This means that if a batch job fails several days into its run, it might be possible to restart the job from a saved copy of its data rather than rerunning the entire multiday job.

Host considerations for FlashCopy integrity

The SAN Volume Controller FlashCopy feature 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 the target VDisk has enough space available 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 metadata 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 be addressable 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 operation. You can flush the host cache by unmounting the source VDIsks from the source host before you start the FlashCopy operation.

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 operation.

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 operation. See 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 metadata on the target VDisk to allow the operating system to use the disk. See your host documentation on how to detect and mount the copied VDisks.

Flushing data from the host volumes

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

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

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

Note: If you can ensure that no reads and writes are issued to the source volumes after you unmount the drives, you can immediately remount and then start the FlashCopy operation.

FlashCopy mappings

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

The FlashCopy feature makes an instant copy of a VDisk at the time that it is started. To create an instant copy 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 VDisks must be of equal size.

A mapping can be created between any two VDIs in a cluster. The VDIs do not have to be in the same I/O group or managed disk (MDisk) group. When a FlashCopy operation starts, a checkpoint is made of the source VDI. 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 VDI has been copied. Each bit in the bitmap represents one region of the source VDI. Each region is called a *grain*.

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

During a read operation to the target VDI, 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 VDI. If the grain has not been copied, the data is read from the source VDI.

When you create a mapping, you specify a copy rate. When the mapping is in the copying state, the copy rate determines the priority that is given to the background copy process. If you want a copy of the whole source VDI so that a mapping can be deleted and still be accessed from the target VDI, you must copy all the data that is on the source VDI to the target VDI.

When you create a mapping, you also specify a clean rate. The clean rate is used to control the rate that data is copied from the target VDI of the mapping to the target VDI of a mapping that is either the latest copy of the target VDI, or is the next oldest copy of the source VDI. The clean rate is used in the following situations:

- The mapping is in the stopping state
- The mapping is in the copying state and has a copy rate of zero
- The mapping is in the copying state and the background copy has completed

You can use the clean rate to minimize the amount of time that a mapping is in the stopping state. If the mapping has not completed, the target VDI is offline while the mapping is stopping. The target VDI remains offline until the mapping is restarted.

While the mapping is in the copying state, you can set the copy rate to zero and the clean rate to a value other than zero to minimize the amount of time a mapping is in the stopping state.

The default values for both the clean rate and the copy rate is 50.

| When a mapping is started and the copy rate is greater than zero (or a value other
| than NOCOPY), the unchanged data is copied to the target VDI, 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 VDI, the whole VDI is copied
| to the target. The mapping returns to the idle_or_copied state and you can now
| restart the mapping at any time to create a new copy at the target.

If you use multiple target mappings, the mapping can stay in the copying state after all of the source data is copied to the target (the progress is 100%). This situation can occur if mappings that were started earlier and use the same source disk are not yet 100% copied.

If the 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. Unless the target VDisk already contains a complete copy of the source VDisk, this action makes the target inconsistent and the target VDisk is taken offline. The target VDisk remains offline until the mapping is restarted.

Multiple target FlashCopy mappings

SAN Volume Controller nodes with a software level of 4.2.0 or higher allow up to 16 target VDIs to be copied from a single source VDisk. Each relationship between a source and target VDisk is managed by a unique mapping such that a single VDisk can be the source VDisk in up to 16 mappings.

Each of the mappings from a single source can be started and stopped independently. If multiple mappings from the same source are active (in the copying or stopping states), a dependency exists between these mappings.

Example 1

Mapping A depends on mapping B if the following is true:

- Mapping A and mapping B both have the same source VDisk
- Mapping A and mapping B are both in the copying or stopping state
- Mapping B was started more recently than mapping A

Note: If both mappings were in the same consistency group and therefore started at the same time, the order of dependency is decided internally when the consistency group is started.

- Mapping A does not have a complete copy of the source because the copying progress for the mapping is less than 100.
- A mapping does not exist from the same source started more recently than A and later than B which has a complete copy of the source because the copying progress of the mapping is less than 100.

Example 2

Target VDisk A depends on target VDisk B if the mapping that VDisk A belongs to depends on the mapping that target VDisk B belongs to. The target VDisk of the most recently started mapping from the source VDisk depends on the source VDisk until there is a complete copy of the source (progress is 100%).

Incremental mappings

SAN Volume Controller nodes with a software level of 4.2.1 or higher allow the background copy process to only copy the parts of the source or target VDisk that have changed since the last FlashCopy process. This reduces the amount of time that it takes to recreate an independent FlashCopy image.

Cascaded mappings

SAN Volume Controller nodes with a software level of 4.2.1 or higher allow target VDIs to be the source of other mappings.

Up to 16 mappings can exist in a cascade. If cascaded mappings and multiple target mappings are used, a tree of up to 16 mappings can be created.

FlashCopy mapping states

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

Idle or copied

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

If the mapping is incremental and the background copy is complete, the mapping only records the differences between the source and target VDisks. If the connection to both nodes in the I/O group that the mapping is assigned to is lost, the source and target VDisks will be offline.

Copying

The copy is in progress. Read and write caching is enabled on the source and the target VDisks.

Prepared

The mapping is ready to start. The target VDisk is online, but is 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. If the mapping is incremental and a previous mapping has completed, the mapping only records the differences between the source and target VDisks. If the connection to both nodes in the I/O group that the mapping is assigned to is lost, the source and target VDisks will be offline.

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. If the mapping is incremental and a previous mapping has completed, the mapping only records the differences between the source and target VDisks. If the connection to both nodes in the I/O group that the mapping is assigned to is lost, the source and target VDisks will be offline.

Stopped

The mapping is stopped because either you issued a stop command or an I/O error occurred. The target VDisk is offline and its data is lost. To access the target VDisk, you must restart or delete the mapping. The source VDisk is accessible and the read and write cache is enabled. If the mapping is incremental, the mapping is recording write operations to the source VDisk. If the connection to both nodes in the I/O group that the mapping is assigned to is lost, the source and target VDisks will be offline.

Stopping

The mapping is in the process of copying data to another mapping.

- If the background copy process is complete, the target VDisk is online while the stopping copy process completes.
- If the background copy process is not complete, data is discarded from the target VDisk cache. The target VDisk is offline while the stopping copy process runs.

The source VDisk is accessible for I/O operations.

Suspended

The mapping started, but it did not complete. Access to the metadata is lost, which causes both the source and target VDisk to go offline. When access to the metadata is restored, the mapping returns to the copying or stopping state and the source and target VDIsks return online. The background copy process resumes. Any data that has not been flushed and has been written to the source or target VDisk before the suspension, is in cache until the mapping leaves the suspended state.

Notes:

1. If a FlashCopy source VDisk goes offline, any FlashCopy target VDIsks that depend on that VDisk also go offline.
2. If a FlashCopy target VDisk goes offline, any FlashCopy target VDIsks that depend on that VDisk also go offline. The source VDisk remains online.

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 Controller nodes; it passes straight through to the MDIsks. The prepare operation for the mapping might take some time to complete; 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 and start the mapping.

Note: The `svctask startfcmap` and `svctask startfcconsistgrp` commands can take some time to process.

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

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.

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	<p>A new FlashCopy mapping is created between the specified source virtual disk (VDisk) and the specified target VDisk. The operation fails if any of the following is true:</p> <ul style="list-style-type: none"> • For SAN Volume Controller software version 4.1.0 or earlier, the source or target VDisk is already a member of a FlashCopy mapping. • For SAN Volume Controller software version 4.2.0 or later, the source VDisk is already a member of 16 FlashCopy mappings. • For SAN Volume Controller software version 4.2.0 or later, the source or target VDisk is already a target VDisk of a FlashCopy mapping. • The node has insufficient bitmap memory. • The source and target VDIs are different sizes.
Prepare	<p>The prepare command is directed to either a consistency group for FlashCopy mappings that are members of a normal consistency group or to the mapping name for FlashCopy mappings that are stand-alone mappings. The prepare command places the FlashCopy mapping into the preparing state.</p> <p>Attention: The prepare command can corrupt any data that 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.</p>
Flush done	<p>The FlashCopy mapping 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 no longer valid.</p>
Start	<p>When all the FlashCopy mappings in a consistency group are in the prepared state, the FlashCopy mappings can be started.</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>

Table 7. FlashCopy mapping events (continued)

Modify	The following FlashCopy mapping properties can be modified: <ul style="list-style-type: none"> • FlashCopy mapping name • Clean rate • Consistency group • Copy rate (for background copy or stopping copy priority) • Automatic deletion of the mapping when the background copy is complete
Stop	There are two separate mechanisms by which a FlashCopy mapping can be stopped: <ul style="list-style-type: none"> • You have issued a command • An I/O error has occurred
Delete	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.
Flush failed	If the flush of data from the cache cannot be completed, the FlashCopy mapping enters the stopped state.
Copy complete	After all of the source data has been copied to the target and there are no dependent mappings, the state is set to copied. If the option to automatically delete the mapping after the background copy completes is specified, the FlashCopy mapping is automatically deleted. If this option is not specified, the FlashCopy mapping is not automatically deleted and can be reactivated by preparing and starting again.
Bitmap Online/Offline	The node has failed.

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 start 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 because they become controlled as part of that 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 operations 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.

Multiple target FlashCopy mappings

Consistency groups aggregate FlashCopy mappings, not the VDisks themselves. Therefore, a source VDisk with multiple FlashCopy mappings can be in different consistency groups. If a VDisk is the source VDisk for several FlashCopy mappings that are in the same consistency group, multiple identical copies of the source VDisk are created when the consistency group is started.

Cascaded FlashCopy mappings

To create a FlashCopy mapping in a consistency group, the source VDisk cannot be the target of a mapping in the same consistency group. In addition, the target VDisk cannot be the source of another FlashCopy mapping in the same consistency group. You cannot move a FlashCopy mapping into a consistency group that contains neighboring FlashCopy mappings in the cascade.

FlashCopy consistency group states

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

Idle_or_Copied

All FlashCopy Mappings in this consistency group are in the Idle or Copied state.

Preparing

At least one FlashCopy mapping in this consistency group is in the Preparing state.

Prepared

The consistency group is ready to start. While in this state, the target VDisks of all FlashCopy mappings in this consistency group are not accessible.

Copying

At least one FlashCopy mapping in the consistency group is in the Copying state and no FlashCopy mappings are in the Suspended state.

Stopping

At least one FlashCopy mapping in the consistency group is in the Stopping state and no FlashCopy mappings are in the Copying or Suspended state.

Stopped

The consistency group is stopped because either you issued a command or an I/O error occurred.

Suspended

At least one FlashCopy mapping in the consistency group is in the Suspended state.

Empty The consistency group does not have any FlashCopy mappings.

Dependent writes

To preserve the integrity of data that is being written, 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. However, if 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, the possibility that the database itself is copied slightly before the database log resulting in the target VDisks seeing writes (1) and (3) but not (2) must be excluded. In this case, if the database is restarted from a backup made from the FlashCopy target disks, the database log indicates that the transaction has completed successfully when, in fact, that is not the case. The transaction is lost and the integrity of the database is compromised.

You can perform a FlashCopy operation on multiple VDisks 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 can contain an arbitrary number of FlashCopy mappings, up to the maximum number of FlashCopy mappings that are supported by a SAN Volume Controller cluster. You can use the command-line interface (CLI) **svctask startfcconsistgrp** command to start the point-in-time copy for the entire consistency group. All of the FlashCopy mappings in the consistency group are started at the same time, resulting in a point-in-time copy that is consistent across all of the FlashCopy mappings that are contained in the consistency group.

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

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

Grains and the FlashCopy bitmap

When data is copied between virtual disks (VDisks), it is copied in units of address space known as grains.

The grain size is 64 KB or 256 KB. The FlashCopy bitmap contains one bit for each grain. The bit records if the associated grain has been split by copying the grain from the source to the target.

Write to target VDisk

A write to the newest target VDisk must consider the state of the grain for its own mapping and the grain of the next oldest mapping.

- If the grain of the intermediate mapping or the next oldest mapping has not been copied, it must be copied before the write is allowed to proceed. This is done to preserve the contents of the next oldest mapping. The data written to the next oldest mapping can come from a target or source.
- If the grain of the target that is being written has not been copied, the grain is copied from the oldest already copied grain in the mappings that are newer than the target (or the source if no targets are already copied). After the copy is complete, the write can be applied to the target.

Read to target VDisk

If the grain that is being read has been split, the read returns data from the target that is being read. If the read is to an uncopied grain on an intermediate target VDisk, each of the newer mappings are examined to determine if the grain has been split. The read is surfaced from the first split grain found or from the source VDisk if none of the newer mappings have a split grain.

FlashCopy indirection layer

The FlashCopy feature provides the semantics of a point-in-time copy by using an indirection layer which intercepts I/Os that are targeted at both the source and target virtual disks (VDisks).

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 that are in the consistency group.

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

- The VDisk 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 VDisk to the source VDisk or stalls the I/O while it arranges for data to be copied from the source VDisk to the target VDisk.

The following table provides an overview of the FlashCopy I/O path actions:

VDisk	Grain already copied?	Host I/O operation	
		Read	Write
Source	No	Read from source	Copies the grain to the most recently started target VDisk for this source VDisk and then writes to the source VDisk.
	Yes	Read from source	Write to source

VDisk	Grain already copied?	Host I/O operation	
Target	No	<p>When the grain has been copied, you can use the following algorithm to determine the VDisk that is being read:</p> <ol style="list-style-type: none"> 1. If newer target VDIsks exist for this source VDisk and the grain has already been copied, the read comes from the oldest target VDisk. 2. If there are no newer target VDIsks, the read comes from the source VDisk. <p>When the grain has not been copied, you can use the following algorithm to determine the VDisk that is being read:</p> <ol style="list-style-type: none"> 1. If newer target VDIsks exist for the source VDisk of the FlashCopy mapping that is the target VDisk being written to and the data has already been copied, the read is from the target VDisk. 2. If the source VDisk is not a target of another FlashCopy mapping, the read is from the source VDisk. 3. If newer target VDIsks exist for the source VDisk of the FlashCopy mapping that is the source VDisk being written to and the data has already been copied, the read is from the target VDisk. 	<ol style="list-style-type: none"> 1. If newer target VDIsks exist for this source VDisk and the grain has already been copied, the read comes from the oldest target VDisk. If there are no newer target VDIsks, the read comes from the source VDisk. 2. If the grain has not already been copied to the next oldest target VDisk for this source VDisk, the same data is also copied to the next oldest target VDisk. 3. Writes to target <p>When the grain has not been copied or overwritten, you can use the following algorithm:</p> <ol style="list-style-type: none"> 1. Use the algorithm for the corresponding read to determine the VDisk to read. 2. If there is an older target VDisk and the data has not been copied to this VDisk, the data is written to this VDisk. 3. If there is a target VDisk for this VDisk and the data has not been copied to this VDisk, the data is written to this VDisk. 4. Writes to target.
	Yes	Read from target	Write to target

Note: For cascaded FlashCopy operations, a VDisk can be both the source and the target. When the VDisk is both the source and target, the I/O path actions are handled as described for a target VDisk.

Source reads

Source reads are always passed through to the underlying source VDisk.

Target reads

To process a read from the target VDisk, the FlashCopy mapping must consult the FlashCopy bitmap. If the data has already been copied to the target VDisk, the read is sent to the target VDisk. If the data has not already been copied, the target read is either sent to the source VDisk, or to another target VDisk if multiple target FlashCopy mappings exist for the source VDisk. While the target read is outstanding, no writes that change the data that is being read are allowed to run.

Background and stopping copy

A FlashCopy mapping has a property called the copy rate. The 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. Because 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.

Note: For the command-line interface (CLI), the value NOCOPY is the same as setting the copy rate to 0 (zero).

The following table provides the relationship of the 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.

User-specified value	Data copied/sec	256 KB grains/sec	64 KB grains/sec
1 - 10	128 KB	0.5	2
11 - 20	256 KB	1	4
21 - 30	512 KB	2	32
41 - 50	2 MB	8	8
51 - 60	4 MB	16	64
61 - 70	8 MB	32	128
71 - 80	16 MB	64	256
81 - 90	32 MB	128	512
91 - 100	64 MB	256	1024

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. Background copy, stopping 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 moved to the other node in the I/O group in the event of the failure of the node that performs the background and stopping 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.

The stopping copy operation copies every grain that is split on the stopping map to the next map (if one exists) which is dependent on that grain. The operation starts searching with the grain that contains the highest LBAs and works in reverse towards the grain that contains LBA 0. Only those grains that other maps are dependent upon are copied.

Cleaning mode

When you create or modify a FlashCopy mapping, you can specify a cleaning rate for the FlashCopy mapping that is independent of the background copy rate. The cleaning rate controls the rate at which the cleaning process operates. The cleaning process copies data from the target VDisk of a mapping to the target VDIsks of other mappings that are dependent on this data. The cleaning process must complete before the FlashCopy mapping can go to the stopping state.

Cleaning mode allows you to activate the cleaning process when the FlashCopy mapping is in the copying state. This keeps your target VDisk accessible while the cleaning process is running. When operating in this mode, it is possible that host I/O operations can prevent the cleaning process from reaching 100% if the I/O operations keep copying new data to the target VDIsks. However, it is possible to minimize the amount of data that requires cleaning when the mapping is stopping.

Cleaning mode is active if the background copy progress has reached 100% and the mapping is in the copying state, or if the background copy rate is set to 0.

Metro Mirror and Global Mirror

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

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

For both Metro Mirror and Global Mirror copy types, one VDisk is designated the primary and the other VDisk is designated the secondary. Host applications write data to the primary VDisk, and updates to the primary VDisk are copied to the secondary VDisk. Normally, host applications do not perform I/O operations to the secondary VDisk.

The Metro Mirror feature provides a synchronous-copy process. When a host writes to the primary VDisk, it does not receive confirmation of I/O completion

until the write operation has completed for the copy on both the primary VDisk and the secondary VDisk. This ensures that the secondary VDisk is always up-to-date with the primary VDisk in the event that a failover operation must be performed. However, the host is limited to the latency and bandwidth limitations of the communication link to the secondary VDisk.

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

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

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

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

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

Metro Mirror and Global Mirror relationships

Metro Mirror and Global Mirror relationships define the relationship between two virtual disks (VDIsks): a master VDisk and an auxiliary VDisk.

Typically, 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 for disaster recovery.

The master and auxiliary VDIsks are defined when the relationship is created, and these attributes never change. However, either VDisk can operate in the primary or secondary role as necessary. The primary VDisk contains a valid copy of the application data and receives updates from the host application, analogous to a source VDisk. The secondary VDisk receives a copy of any updates to the primary VDisk, because these updates are all transmitted across the Mirror link. Therefore, the secondary VDisk is analogous to a continuously updated target VDisk. 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 from the command-line interface (CLI) or the SAN Volume Controller Console.

The two VDIsks in a relationship must be the same size. When the two VDIsks are in the same cluster, they must be in the same I/O group.

A relationship can be added to a 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 chrrelationship** command to add or remove a relationship to or from a consistency group. See the *IBM System Storage SAN Volume Controller: Command-Line Interface User's Guide*.

Copy types

A Metro Mirror copy ensures that updates are committed to both the primary and secondary VDIs before sending confirmation of I/O completion to the host application. This ensures that the secondary VDisk is synchronized with the primary VDisk in the event that a failover operation is performed.

A Global Mirror copy allows the host application to receive confirmation of I/O completion before the updates are committed to the secondary VDisk. If a failover operation is performed, the host application must recover and apply any updates that were not committed to the secondary VDisk.

States

When a Metro Mirror or Global Mirror relationship is created with two VDIs in different clusters, the distinction between the connected and disconnected states is important. These states apply to both clusters, the relationships, and the consistency groups. The following Metro Mirror and Global Mirror relationship states are possible:

Inconsistent (Stopped)

The primary VDisk is accessible for read and write 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 an **svctask startrelationship** command is issued to a consistency group in the InconsistentStopped state. This state is also entered when an **svctask startrelationship** 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 occur when a relationship was in the ConsistentSynchronized state and experiences an error that forces a freeze of the consistency group. This state can also occur 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 an auxiliary VDisk operates in the primary role. Consequently the VDisk is accessible for write I/O operations.

Idling (Disconnected)

The VDIs 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.

Metro Mirror and Global Mirror relationships between two clusters

Metro Mirror and Global Mirror relationships can exist simultaneously between two clusters. In this type of configuration, there can be impacts to performance because write data from both Metro Mirror and Global Mirror relationships is transported over the same intercluster links.

Metro Mirror and Global Mirror relationships handle heavy workload differently. Metro Mirror typically maintains the relationships that are in the copying or synchronized states, which causes the primary host applications to see degraded performance. Global Mirror requires a higher level of write performance to primary host applications. If the link performance is severely degraded, the link tolerance feature automatically stops Global Mirror relationships when the link tolerance threshold is exceeded. As a result, Global Mirror writes can suffer degraded performance if Metro Mirror relationships utilize most of the capability of the intercluster link.

Metro Mirror and Global Mirror partnerships

Metro Mirror and Global Mirror partnerships define the relationship between a local cluster and a remote cluster.

The SAN Volume Controller nodes must know not only about the relationship between the two VDisks but also about 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 specify 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 the *IBM System Storage SAN Volume Controller: Command-Line Interface User's Guide*.

You can also use the SAN Volume Controller Console to create Metro Mirror and Global Mirror partnerships.

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.

Configuration requirements

To use the Global Mirror feature, all components in the SAN must be capable of sustaining the workload that is generated by application hosts and the Global

Mirror background copy process. If all of the components in the SAN cannot sustain the workload, the Global Mirror relationships are automatically stopped to protect your application hosts from increased response times.

To use the Global Mirror feature, the following requirements must be met:

- Use IBM TotalStorage Productivity Center or an equivalent SAN performance analysis tool to monitor your SAN environment. The IBM TotalStorage Productivity Center provides an easy way to analyze the SAN Volume Controller performance statistics.
- Analyze the SAN Volume Controller performance statistics to determine the peak application write workload that the link must support. Gather statistics over a typical application I/O workload cycle.
- Set the background copy rate to a value that supports the intercluster link and the backend storage controllers at the remote cluster.
- Do not use cache-disabled VDisks in Global Mirror relationships.
- Set the gmlinktolerance parameter to an appropriate value. The default value is 300 seconds (5 minutes).
- When you perform SAN maintenance tasks, take one of the following actions:
 - Reduce the application I/O workload for the duration of the maintenance task.
 - Disable the gmlinktolerance feature or increase the gmlinktolerance value.

Note: If the gmlinktolerance value is increased during the maintenance task, do not set it to the normal value until the maintenance task is complete. If the gmlinktolerance feature is disabled for the duration of the maintenance task, enable it after the maintenance task is complete.

- Stop the Global Mirror relationships.
- Evenly distribute the preferred nodes for the Global Mirror VDisks between the nodes in the clusters. Each VDisk in an I/O group has a preferred node property that can be used to balance the I/O load between nodes in the I/O group. The preferred node property is also used by the Global Mirror feature to route I/O operations between clusters. A node that receives a write for a VDisk is normally the preferred node for that VDisk. If the VDisk is in a Global Mirror relationship, the node is responsible for sending the write to the preferred node of the secondary VDisk. The preferred node property alternates between the nodes of an I/O group as VDisks are created within the I/O group. Each node in the remote cluster has a set pool of Global Mirror system resources for each node in the local cluster. To maximize Global Mirror performance, set the preferred nodes for the VDisks of the remote cluster to use every combination of primary nodes and secondary nodes.

Long distance links for Metro Mirror and Global Mirror partnerships

For intracluster partnerships, all clusters can be considered as candidates for Metro Mirror or Global Mirror operations. For intercluster partnerships, cluster pairs must be separated by a number of moderately high bandwidth links.

Figure 11 on page 54 shows an example of a configuration that uses dual redundant fabrics. Part of each fabric is located at the local cluster and the remote cluster. There is no direct connection between the two fabrics.

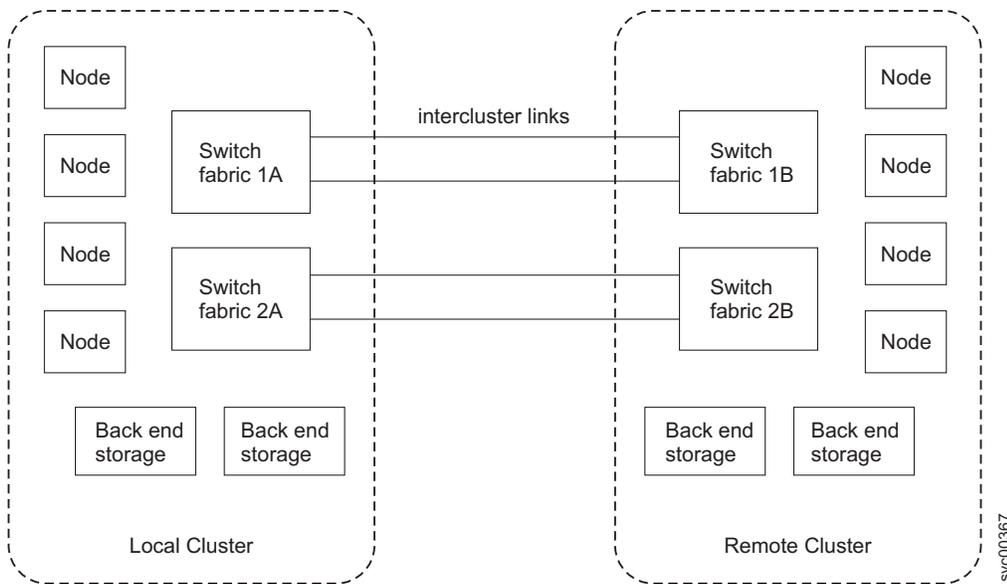


Figure 11. Redundant fabrics

You can use fibre-channel extenders or SAN routers to increase the distance between two clusters. Fibre-channel extenders transmit fibre-channel packets across long links without changing the contents of the packets. SAN routers provide virtual nPorts on two or more SANs to extend the scope of the SAN. The SAN router distributes the traffic from one virtual nPort to the other virtual nPort. The two fibre-channel fabrics are independent of each other. Therefore, nPorts on each of the fabrics cannot directly log into each other. See the following Web site for specific firmware levels and the latest supported hardware:

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

If you use fibre-channel extenders or SAN routers, you must meet the following requirements:

- For SAN Volume Controller software level 4.1.0, the round-trip latency between sites cannot exceed 68 ms for fibre-channel extenders or 20 ms for SAN routers.
- For SAN Volume Controller software level 4.1.1 or higher, the round-trip latency between sites cannot exceed 80 ms for either fibre-channel extenders or SAN routers.
- The configuration must be tested with the expected peak workloads.
- Metro Mirror and Global Mirror require a specific amount of bandwidth for intercluster heartbeat traffic. The amount of traffic depends on the number of nodes that are in both the local cluster and the remote cluster. Table 8 lists the intercluster heartbeat traffic for the primary cluster and the secondary cluster. These numbers represent the total traffic between two clusters when there are no I/O operations running on the copied VDisks. Half of the data is sent by the primary cluster and half of the data is sent by the secondary cluster so that traffic is evenly divided between all of the available intercluster links. If you have two redundant links, half of the traffic is sent over each link.

Table 8. Intercluster heartbeat traffic in Mbps

Cluster 1	Cluster 2			
	2 nodes	4 nodes	6 nodes	8 nodes
2 nodes	2.6	4.0	5.4	6.7

Table 8. Intercluster heartbeat traffic in Mbps (continued)

Cluster 1	Cluster 2			
4 nodes	4.0	5.5	7.1	8.6
6 nodes	5.4	7.1	8.8	10.5
8 nodes	6.7	8.6	10.5	12.4

- The bandwidth between two sites must meet the peak workload requirements and maintain the maximum round-trip latency between the sites. When you evaluate the workload requirement, you must consider the average write workload over a period of one minute or less and the required synchronization copy bandwidth. If there are no active synchronization copies and no write I/O operations for VDIs that are in the Metro Mirror or Global Mirror relationship, the SAN Volume Controller protocols operate with the bandwidth that is indicated in Table 8 on page 54. However, you can only determine the actual amount of bandwidth that is required for the link by considering the peak write bandwidth to VDIs that are participating in Metro Mirror or Global Mirror relationships and then adding the peak write bandwidth to the peak synchronization bandwidth.
- If the link between two sites is configured with redundancy so that it can tolerate single failures, the link must be sized so that the bandwidth and latency statements are correct during single failure conditions.
- The channel must not be used for links between nodes in a single cluster. Configurations that use long distance links in a single cluster are not supported and can cause I/O errors and loss of access.
- The configuration is tested to confirm that any failover mechanisms in the intercluster links interoperate satisfactorily with the SAN Volume Controller.
- All other SAN Volume Controller configuration requirements are met.

Limitations on host to cluster distances

There is no limit on the fibre-channel optical distance between SAN Volume Controller nodes and host servers. You can attach a server to an edge switch in a core-edge configuration with the SAN Volume Controller cluster at the core. SAN Volume Controller clusters support up to three ISL hops in the fabric. This means that the host server and the SAN Volume Controller cluster can be separated by up to five fibre-channel links. If you use longwave SFPs, four of the fibre-channel links can be 10 km long.

Using the intercluster link for host traffic

If you use the intercluster link for host traffic, ensure that you have sufficient bandwidth to support all sources of load.

Scenario: The hosts in a local cluster can read and write to the VDIs in a remote cluster

In this scenario, the hosts in the local cluster also exchange heartbeats with the hosts that are in the remote cluster. Because the intercluster link is being used for multiple purposes, you must have sufficient bandwidth to support the following sources of load:

- Global Mirror or Metro Mirror data transfers and the SAN Volume Controller cluster heartbeat traffic.
- Local host to remote VDI I/O traffic or remote host to local VDI I/O traffic.

- Local host to remote host heartbeat traffic. If the local host to remote VDisk I/O traffic is allowed to consume a high percentage of intercluster link bandwidth, the latency seen by the hosts that access SAN Volume Controller VDIsks that are participating in Metro Mirror or Global Mirror operations can be impacted. The bandwidth congestion can cause the Global Mirror link tolerance threshold to be exceeded. When the Global Mirror link tolerance threshold is exceeded, Global Mirror relationships are stopped.

Metro Mirror and Global Mirror consistency groups

You can group Metro Mirror or Global Mirror relationships into a consistency group so that they can be updated at the same time. A command that is issued to the consistency group is simultaneously applied to all of the relationships in the group.

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 can be updated. This ensures that the set of secondary VDIsks contain usable data. The key property is that these relationships are consistent.

Relationships can only belong to one consistency group; however, they do not have to belong to a consistency group. 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 and Global Mirror relationships cannot belong to the same consistency group. A copy type is automatically assigned to a consistency group when the first relationship is added to the consistency group. After the consistency group is assigned a copy type, only relationships of that copy type can be added to the consistency group. Each cluster can have a maximum of six different types of consistency groups. The following types of consistency groups are possible:

- Intracluster Metro Mirror
- Intercluster Metro Mirror from the local cluster to remote cluster
- Intercluster Metro Mirror from the remote cluster to local cluster
- Intracluster Global Mirror
- Intercluster Global Mirror from the local cluster to remote cluster
- Intercluster Global Mirror from the remote cluster to local cluster

States

A consistency group can be in one of the following states:

Inconsistent (stopped)

The primary VDIsks are accessible for read and write 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 VDIs 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 VDIs contain a consistent image, but it might be out-of-date with respect to the primary VDIs. This state can occur when a relationship was in the ConsistentSynchronized state and experiences an error that forces a freeze of the consistency group. This state can also occur when a relationship is created with the CreateConsistentFlag set to TRUE.

Consistent (synchronized)

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

Idling Both the primary VDIs and the secondary VDIs are operating in the primary role. Consequently the VDIs are accessible for write I/O operations.

Idling (disconnected)

The VDIs 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 VDIs 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 VDIs 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.

Empty The consistency group does not contain any relationships.

Background copy bandwidth impact on foreground I/O latency

The background copy bandwidth determines the rate at which the background copy for Metro Mirror or Global Mirror Copy Services are 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 intercluster link capacity, the following results can occur:
 - The background copy I/Os can back up on the intercluster link
 - For Metro Mirror, there is a delay in the synchronous secondary writes of foreground I/Os
 - For Global Mirror, the work is backlogged, which delays the processing of writes and causes the relationship to stop
 - The foreground I/O latency increases 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.
 - For Global Mirror, the work is backlogged and again the relationship is stopped

In order to set the background copy bandwidth optimally, you must 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. You must also consider concurrent host I/O because if other writes arrive at the primary cluster for copy to the remote site, these writes can be delayed by a high level of background copy and the hosts at the primary site receive poor write response times.

This provisioning can be done by the calculation above or by determining 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.

Example

If the bandwidth setting at the primary site for the secondary cluster is set to 200 Mbps (megabytes per second) and the relationships are not synchronized, the SAN Volume Controller attempts to resynchronize the relationships at a maximum rate of 200 Mbps with a 25 Mbps restriction for each individual relationship. The SAN Volume Controller cannot resynchronize the relationship if the throughput is restricted. The following can restrict throughput:

- The read response time of backend storage at the primary cluster
- The write response time of the backend storage at the secondary site
- Intercluster link latency

Backend storage controller requirements

The performance of applications at the local cluster can be limited by the performance of the backend storage controllers at the remote cluster.

Your set up must meet the following requirements to maximize the amount of I/O operations that applications can run on Global Mirror VDisks:

- The Global Mirror VDisks at the remote cluster must be in dedicated MDisk groups that only contain other Global Mirror VDisks.
- Configure storage controllers to support the Global Mirror workload that is required of them. The following guidelines can be used to fulfill this requirement:
 - Dedicate storage controllers to only Global Mirror VDisks
 - Configure the storage controller to guarantee sufficient quality of service for the disks that are being used by Global Mirror operations
 - Ensure that physical disks are not shared between Global Mirror VDisks and other I/O operations. For example, do not split an individual RAID array.
- For Global Mirror MDisk groups, use MDisks with the same characteristics. For example, use MDisks that have the same RAID level, physical disk count, and disk speed. This requirement is important to maintain performance when you use the Global Mirror feature.

You must provision the storage controllers that are attached to the remote cluster to accommodate the following:

- The peak application workload to the Global Mirror VDisks
- The specified background copy level
- All I/O operations that run on the remote cluster

Migrating a Metro Mirror relationship to a Global Mirror relationship

You can migrate a Metro Mirror relationship to a Global Mirror relationship.

Scenario: I/O operations to the secondary VDisk can be stopped during the migration

In this scenario, you have the ability to stop I/O operations to the secondary VDisk during the migration process.

To stop I/O operations to the secondary VDisk while migrating a Metro Mirror relationship to a Global Mirror relationship, you must specify the synchronized option when you create the Global Mirror relationship.

1. Stop all host I/O operations to the primary VDisk.
2. Verify that the Metro Mirror relationship is consistent.

Important: If the Metro Mirror relationship is not consistent when it is stopped, or if any host I/O operations run between the Metro Mirror relationship being stopped and the Global Mirror relationship being created, the updates are not copied to the secondary VDisk.

3. Delete the Metro Mirror relationship.
4. Create the Global Mirror relationship between the same two VDisks.

After the Global Mirror relationship is created, you can start the relationship and resume host I/O operations.

Scenario: I/O operations to the secondary VDisk cannot be stopped during the migration

In this scenario, you do not have the ability to stop I/O operations to the secondary VDisk during the migration process.

If I/O operations to the secondary VDisk cannot be stopped, the data on the secondary VDisk becomes out-of-date. When the Global Mirror relationship is started, the secondary VDisk is inconsistent until all of the recent updates are copied to the remote site.

If you do not require a consistent copy of the VDisk at the secondary site, perform the following steps to migrate from a Metro Mirror relationship to a Global Mirror relationship:

Important: The data on the secondary VDisk is not usable until the synchronization process is complete. Depending on your link capabilities and the amount of data that is being copied, this process

can take an extended period of time. You must set the background copy bandwidth for the intercluster partnerships to a value that does not overload the intercluster link.

1. Delete the Metro Mirror relationship.
2. Create and start the Global Mirror relationship between the same two VDisks.

If you require a consistent copy of the VDisk at the secondary site, perform the following steps to migrate from a Metro Mirror relationship to a Global Mirror relationship:

1. Delete the Metro Mirror relationship.
2. Create a Global Mirror relationship between VDisks that were not used for the Metro Mirror relationship. This preserves the VDisk so that you can use it if you require a consistent copy at a later time.

Alternatively, you can use the FlashCopy feature to maintain a consistent copy. Perform the following steps to use the FlashCopy feature to maintain a consistent copy:

1. Start a FlashCopy operation for the Metro Mirror VDisk.
2. Wait for the FlashCopy operation to complete.
3. Create and start the Global Mirror relationship between the same two VDisks. The FlashCopy VDisk is now your consistent copy.

Using FlashCopy to create a consistent image before restarting a Global Mirror relationship

For disaster recovery purposes, you can use the FlashCopy feature to create a consistent copy of an image before you restart a Global Mirror relationship.

When a consistent relationship is stopped, the relationship enters the `consistent_stopped` state. While in this state, I/O operations at the primary site continue to run. However, updates are not copied to the secondary site. When the relationship is restarted, the synchronization process for new data is started. During this process, the relationship is in the `inconsistent_copying` state. The secondary VDisk for the relationship cannot be used until the copy process completes and the relationship returns to the consistent state. When this occurs, start a FlashCopy operation for the secondary VDisk before you restart the relationship. While the relationship is in the copying state, the FlashCopy feature can provide a consistent copy of the data. If the relationship does not reach the synchronized state, you can use the FlashCopy target VDisk at the secondary site.

The SVCTools package that is available on the IBM Alphaworks Web site provides an example script that demonstrates how to manage the FlashCopy process. See the `copymanager` script that is available in the SVCTools package. You can download the SVCTools package from the following Web site:

<http://www.alphaworks.ibm.com/tech/svctools/download>

Monitoring Global Mirror performance with the IBM TotalStorage Productivity Center

You can use the IBM TotalStorage Productivity Center to monitor key Global Mirror performance measurements.

It is important to use a Storage Area Network (SAN) performance monitoring tool to ensure that all SAN components are performing correctly. This is particularly

important when you use an asynchronous copying solution such as the SAN Volume Controller Global Mirror feature. IBM TotalStorage Productivity Center monitors key performance measures and alerts you when thresholds are exceeded.

Note: If your VDisk or MDisk configuration changes, restart the IBM TotalStorage Productivity Center performance report to ensure that performance is monitored for the new configuration.

Use IBM TotalStorage Productivity Center to check the following measurements:

- The *Port to Remote Node Send Response Time* measurement is less than 80 milliseconds. If this measurement is greater than 80 milliseconds during monitoring, the long-distance link has excessive latency. Ensure that the link is operating at its maximum bandwidth.
- The sum of the *Port to Local Node Send Response Time* measurement and the *Port to Local Node Send Queue* measurement is less than 1 millisecond for the primary cluster and the CPU Utilization Percentage is below 50%. A value that exceeds these amounts can indicate that an I/O group is reaching the I/O throughput limit, which can limit performance.
- The sum of the *Backend Write Response Time* measurement and the *Write Queue Time for Global Mirror MDisks* measurement of the secondary cluster is less than 100 milliseconds. A longer response time can indicate that the storage controller is overloaded.
- The sum of the *Backend Write Response Time* measurement and the *Write Queue Time for Global Mirror MDisks* measurement of the primary cluster is less than 100 milliseconds. If the response time is greater than 100 milliseconds, application hosts might see extended response times when the SAN Volume Controller cluster cache is full.
- The *Write Data Rate for Global Mirror MDisk groups* measurement of the secondary cluster indicates the amount of data that is being written by Global Mirror operations. If this value approaches either the intercluster link bandwidth or the storage controller throughput limit, further increases can cause overloading of the system. Monitor for this condition in a way that is appropriate for your network.

The gmlinktolerance feature

You can use the `svctask chcluster` CLI command or the SAN Volume Controller Console to set the `gmlinktolerance` feature. The `gmlinktolerance` feature represents the number of seconds that the primary SAN Volume Controller cluster tolerates slow response times from the secondary cluster.

If the poor response extends past the specified tolerance, a 1920 error is logged and one or more Global Mirror relationships are automatically stopped. This protects the application hosts at the primary site. During normal operation, application hosts see a minimal impact to response times because the Global Mirror feature uses asynchronous replication. However, if Global Mirror operations experience degraded response times from the secondary cluster for an extended period of time, I/O operations begin to queue at the primary cluster. This results in an extended response time to application hosts. In this situation, the `gmlinktolerance` feature stops Global Mirror relationships and the application hosts response time returns to normal. After a 1920 error has occurred, the Global Mirror auxiliary VDisks are no longer in the `consistent_synchronized` state until you fix the cause of the error and restart your Global Mirror relationships. For this reason, ensure that you monitor the cluster to track when this occurs.

You can disable the `gmlinktolerance` feature by setting the `gmlinktolerance` value to 0 (zero). However, the `gmlinktolerance` cannot protect applications from extended response times if it is disabled. It might be appropriate to disable the `gmlinktolerance` feature in the following circumstances:

- During SAN maintenance windows where degraded performance is expected from SAN components and application hosts can withstand extended response times from Global Mirror VDisks.
- During periods when application hosts can tolerate extended response times and it is expected that the `gmlinktolerance` feature might stop the Global Mirror relationships. For example, if you are testing using an I/O generator which is configured to stress the backend storage, the `gmlinktolerance` feature might detect the high latency and stop the Global Mirror relationships. Disabling `gmlinktolerance` prevents this at the risk of exposing the test host to extended response times.

Diagnosing and fixing 1920 errors

A 1920 error indicates that one or more of the SAN components are unable to provide the performance that is required by the application hosts. This can be temporary (for example, a result of maintenance activity) or permanent (for example, a result of a hardware failure or unexpected host I/O workload). To diagnose the cause of a 1920 error, it is very important that your SAN performance analysis tool (for example, IBM TotalStorage Productivity Center) is correctly configured and monitoring statistics when the problem occurs. Set your SAN performance analysis tool to the minimum available statistics collection interval. For IBM TotalStorage Productivity Center, the minimum interval is five minutes. If several 1920 errors have occurred, diagnose the cause of the earliest error first. The following questions can help you determine the cause of the error:

- Was maintenance occurring at the time of the error? This might include replacing a storage controller's physical disk, upgrading a storage controller's firmware, or performing a code upgrade on one of the SAN Volume Controller clusters. You must wait until the maintenance procedure is complete and then restart the Global Mirror relationships. You must wait until the maintenance procedure is complete to prevent a second 1920 error because the system has not yet returned to a stable state with good performance.
- Is the long distance link overloaded? If your link is not capable of sustaining the short-term peak Global Mirror workload, a 1920 error can occur. Perform the following checks to determine if the long distance link is overloaded:
 - Look at the total Global Mirror auxiliary VDisk write throughput before the Global Mirror relationships were stopped. If this is approximately equal to your link bandwidth, your link might be overloaded. This might be due to application host I/O operations or a combination of host I/O and background (synchronization) copy activities.
 - Look at the total Global Mirror source VDisk write throughput before the Global Mirror relationships were stopped. This represents the I/O operations that are being performed by the application hosts. If these operations are approaching the link bandwidth, upgrade the link's bandwidth, reduce the I/O operations that the application is attempting to perform, or use Global Mirror to copy fewer VDisks. If the auxiliary disks show significantly more write I/O operations than the source VDisks, there is a high level of background copy. Decrease the Global Mirror partnership's background copy rate parameter to bring the total application I/O bandwidth and background copy rate within the link's capabilities.

- Look at the total Global Mirror source VDisk write throughput after the Global Mirror relationships were stopped. If write throughput increases by 30% or more when the relationships are stopped, the application hosts are attempting to perform more I/O operations than the link can sustain. While the Global Mirror relationships are active, the overloaded link causes higher response times to the application host, which decreases the throughput it can achieve. After the Global Mirror relationships have stopped, the application host sees lower response times. In this case, the link bandwidth must be increased, the application host I/O rate must be decreased, or fewer VDIs must be copied using Global Mirror.
- Are the storage controllers at the secondary cluster overloaded? If one or more of the MDisks on a storage controller are providing poor service to the SAN Volume Controller cluster, a 1920 error occurs if this prevents application I/O operations from proceeding at the rate that is required by the application host. If the backend storage controller requirements have been followed, the error might have been caused by a decrease in controller performance. Use IBM TotalStorage Productivity Center to obtain the backend write response time for each MDisk at the secondary cluster. If the response time for any individual MDisk exhibits a sudden increase of 50 ms or more or if the response time is above 100 ms, this indicates a problem. Perform the following checks to determine if the storage controllers are overloaded:
 - Check the storage controller for error conditions such as media errors, a failed physical disk, or associated activity such as RAID array rebuilding. If there is an error, you should fix the problem and then restart the Global Mirror relationships.
 - If there is no error, determine if the secondary controller is capable of processing the required level of application host I/O operations. It might be possible to improve the performance of the controller by adding more physical disks to a RAID array, changing the RAID level of the array, changing the controller's cache settings and checking the cache battery to ensure it is operational, or changing other controller-specific configuration parameters.
- Are the storage controllers at the primary cluster overloaded? Analyze the performance of the primary backend storage using the same steps as for the secondary backend storage. If performance is bad, limit the amount of I/O operations that can be performed by application hosts. Monitor the backend storage at the primary site even if the Global Mirror relationships have not been affected. If bad performance continues for a prolonged period, a 1920 error occurs and the Global Mirror relationships are stopped.
- Is one of your SAN Volume Controller clusters overloaded? Use IBM TotalStorage Productivity Center to obtain the port to local node send response time and the port to local node send queue time. If the total of these two statistics for either cluster is above 1 millisecond, the SAN Volume Controller might be experiencing a very high I/O load. Also check the SAN Volume Controller node CPU utilization. If this figure is above 50%, this can also be contributing to the problem. In either case, contact your IBM service representative for further assistance. If CPU utilization is much higher for one node than for the other node in the same I/O group, this might be caused by having different node hardware types within the same I/O group. For example, a SAN Volume Controller 2145-8F4 in the same I/O group as a SAN Volume Controller 2145-8G4. If this is the case, contact your IBM service representative.
- Do you have FlashCopy operations in the prepared state at the secondary cluster? If the Global Mirror auxiliary VDIs are the sources of a FlashCopy mapping and that mapping is in the prepared state for an extended time,

| performance to those VDisks can be impacted because the cache is disabled.
| Start the FlashCopy mapping to enable the cache and improve performance for
| Global Mirror I/O operations.
|

Chapter 3. Configuring the SAN fabric

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

Table 9 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 and Global 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 and Global Mirror, significant distances might exist between the components of the local cluster and those of the remote cluster.

Table 9. Configuration terms and definitions (continued)

Term	Definition
Local/remote fabric interconnect	The SAN components that connect the local fabrics to the remote fabrics. There might be significant distances between the components in the local cluster and those in the remote cluster. These components might be single-mode optical fibers that are driven by gigabit interface converters (GBICs), or they might be other, more advanced components, such as channel extenders
SAN Volume Controller fibre-channel port fan in	The number of hosts that can see any one port. Some controllers recommend that the number of hosts using each port be limited to prevent excessive queuing at that port. If the port fails or the path to that port fails, the host might failover to another port, and the fan in requirements might be exceeded in this degraded mode.
Not valid configuration	The current SAN configuration is indicated as not being valid. An attempted operation failed and generated an error code to indicate what caused it to become not valid. The most likely cause is that either a device has failed, or a device has been added to the SAN that has caused the configuration to be marked as not valid.
Unsupported configuration	A configuration that might operate successfully, but for which IBM does not guarantee the solution for problems that might occur. Usually this type of configuration does not create an error log entry.
Valid configuration	A configuration that consists of devices and connections that are identified as a valid and supported configuration. Neither of these two conditions exist with the current configuration: <ul style="list-style-type: none"> • Not valid • Unsupported configuration.
Degraded	A valid configuration that has had a failure, but continues to be neither invalid nor unsupported. Typically, a repair action is required to restore the degraded configuration to a valid configuration.
Fibre channel extender	A device for long distance communication 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 in a loop with some of the paths short circuiting the loop. An example of this configuration is to have four switches connected in a loop with ISLs for one of the diagonals.

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 for each host type. In the disk zone, the SAN Volume Controller nodes identify the disk drives. Host systems cannot operate on the disk drives directly; all data transfer occurs through the SAN Volume Controller nodes. Figure 12 on page 67 shows several host systems that are connected in a SAN fabric.

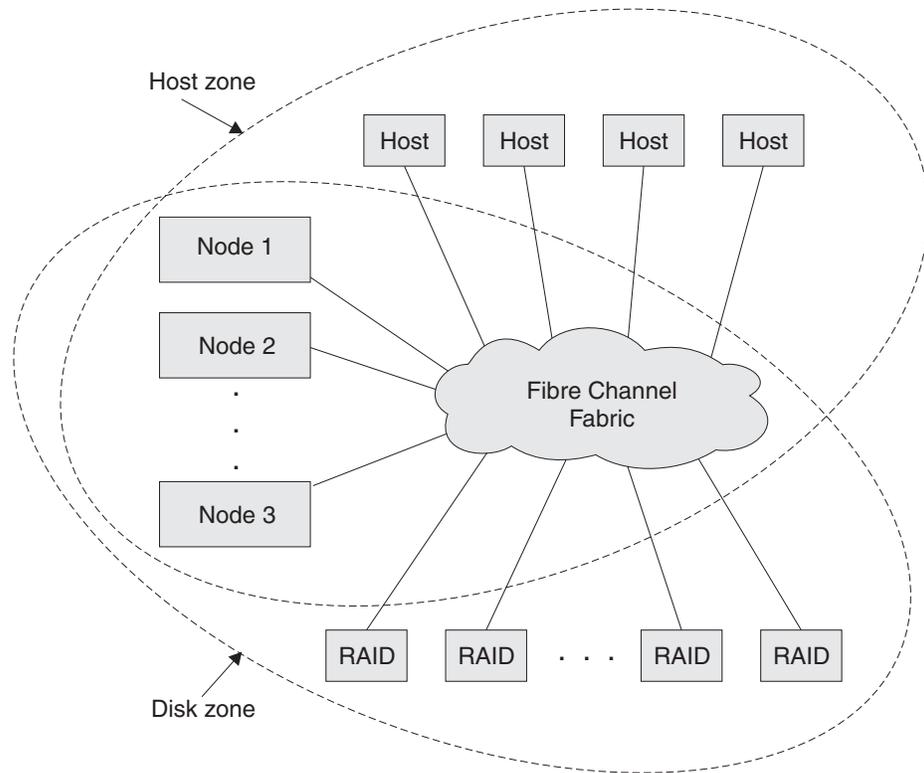


Figure 12. Example of a SAN Volume Controller cluster in a fabric

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

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

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.

Configuration rules

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

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

- Storage subsystems
- Host HBAs
- Nodes
- Fibre-channel switches
- Fabrics
- Zoning

Storage subsystems

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

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

A configuration in which SAN Volume Controller nodes bridge a separate host device and a RAID is not supported. See the following Web site for the latest support information:

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

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 JBOD (just a bunch of disks) LUs that are presented by supported RAID controllers, the SAN Volume Controller itself does not provide RAID functions, so these LUs are exposed to data loss in the event of a disk failure.

If a single RAID controller presents multiple LUs, either by having multiple RAID configured or by partitioning one or more RAID into multiple LUs, each LU can be owned by either 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 Controller nodes 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 multipathing driver for the two hosts to be the same. Figure 13 on page 70 shows that the RAID controller is an IBM DS4000, with RDAC used for pathing on the directly attached host, and SDD used on the host that is attached with the SAN Volume Controller. Hosts can simultaneously access LUs that are provided by the SAN Volume Controller and directly by the device.

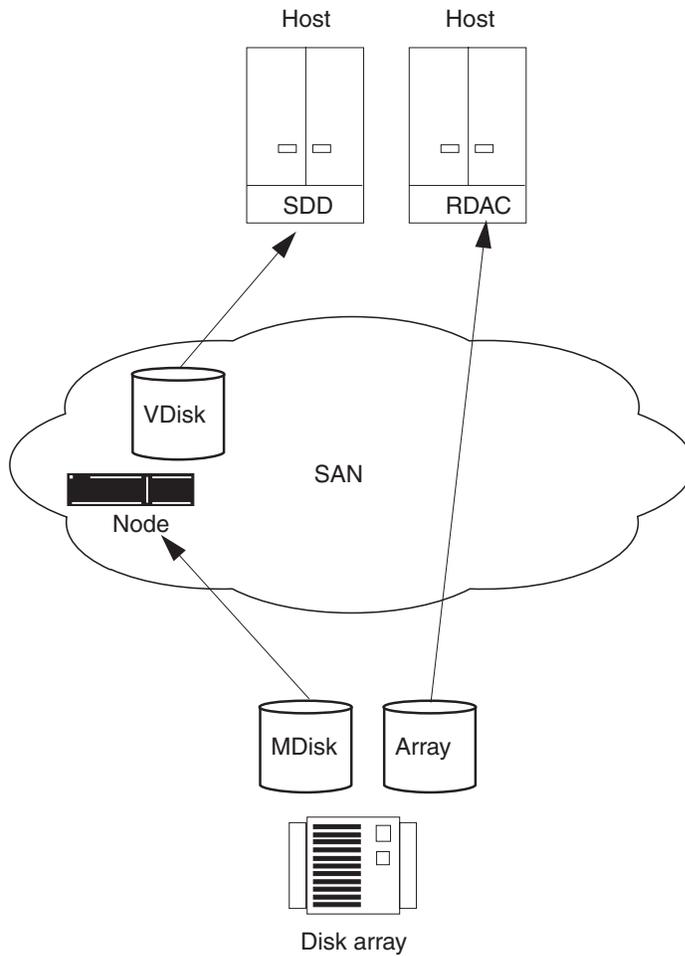


Figure 13. Disk controller system shared between SAN Volume Controller and a host

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 the SAN Volume Controller nodes multipathing software. Figure 14 on page 71 is a supported configuration because the same multipathing driver is used for both direct and VDIs.

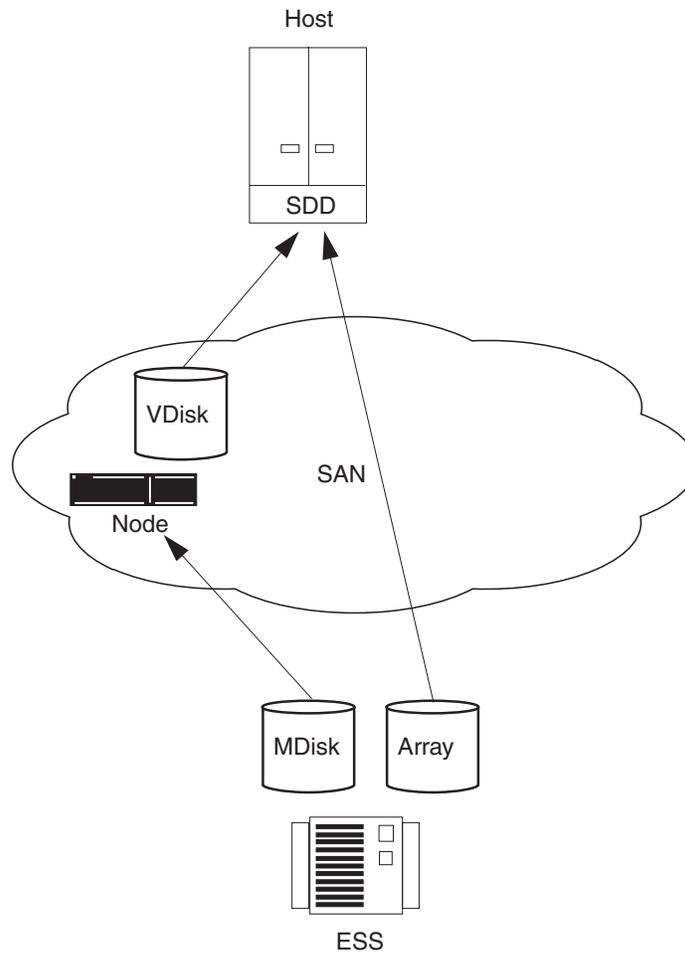


Figure 14. IBM ESS LUs accessed directly with a SAN Volume Controller

In the case where the RAID controller uses multipathing software that is compatible with SAN Volume Controller nodes multipathing software (see Figure 15 on page 72), it is possible to configure a system where some LUNs are mapped directly to the host and others are accessed through the SAN Volume Controller. An IBM TotalStorage Enterprise Storage Server (ESS) that uses the same multipathing driver as a SAN Volume Controller is one example.

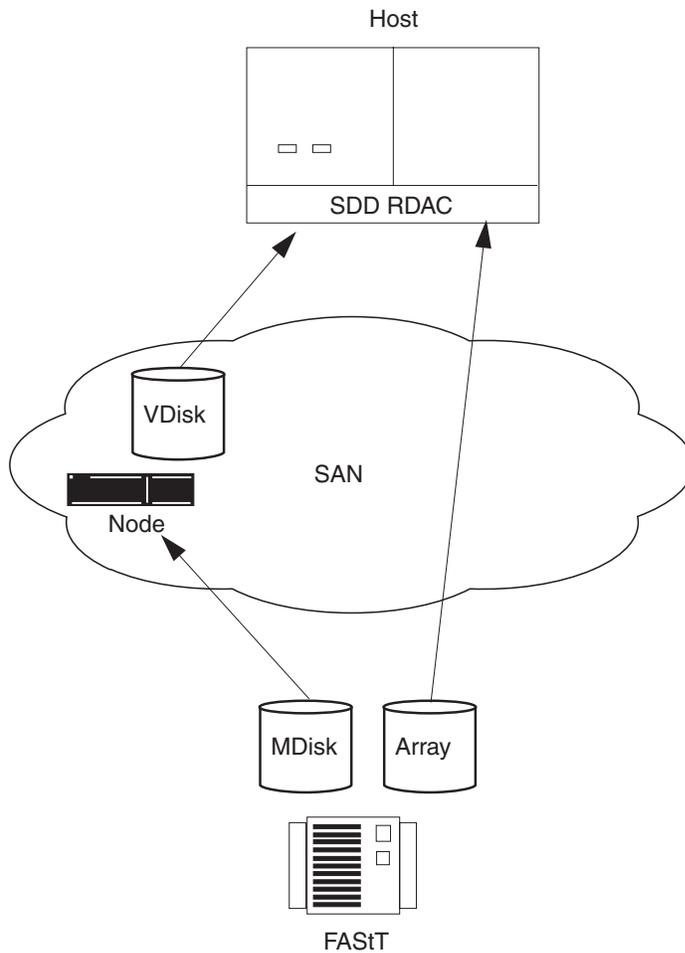


Figure 15. IBM DS4000 direct connection with a SAN Volume Controller on one host

Host 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.

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

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

Operation with other HBAs is not supported.

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

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

Nodes

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

Node HBAs

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

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

I/O groups

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

VDisks

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

Optical connections

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

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

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

The number of paths through the network from the node to a host must not exceed eight. Configurations in which this number is exceeded are not supported. Each node has four ports and each I/O group has two nodes. Therefore, without any zoning, the number of paths to a VDisk is eight × the number of host ports.

Ethernet connection

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

Port speed

You can change the operational port speed for SAN Volume Controller 2145-4F2 and SAN Volume Controller 2145-8F2 nodes to 1 Gbps or 2 Gbps. However, the optical fiber connections between the fibre-channel switches and all SAN Volume Controller 2145-4F2 and SAN Volume Controller 2145-8F2 nodes in a cluster must run at the same speed. The fibre-channel ports on SAN Volume Controller 2145-8F4 and SAN Volume Controller 2145-8G4 nodes auto negotiate the operational port speed independently, which allows these nodes to operate at different speeds. SAN Volume Controller 2145-8F4 and SAN Volume Controller 2145-8G4 nodes can operate at 1 Gbps, 2 Gbps or 4 Gbps. If these nodes are connected to a 4 Gbps capable switch, the port attempts to operate at 4 Gbps; however, if there is a large number of link error rates, the adapter negotiates a lower speed.

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/storage/support/2145>

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

Configurations with more than four SANs are not supported.

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

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

When you attach a node to a SAN fabric that contains core directors and edge switches, connect the node ports to the core directors and connect the host ports to

the edge switches. In this type of fabric, the next priority for connection to the core directors is the storage controllers, leaving the host ports connected to the edge switches.

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

Mixing manufacturer switches in a single SAN fabric

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

- 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.

Brocade core-edge fabrics

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

SAN Volume Controller software level 4.1.1 or higher

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

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

Fibre-channel switches and interswitch links

The local or remote fabric must not contain more than three interswitch link (ISL) 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 or Global 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.

If all three allowed ISL hops are used within the local and remote fabrics, the local remote fabric interconnect must be a single ISL hop between a switch in the local fabric and a switch in the remote fabric.

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

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

With ISLs between nodes in the same cluster, the ISLs are considered a single point of failure. This is illustrated in Figure 16.

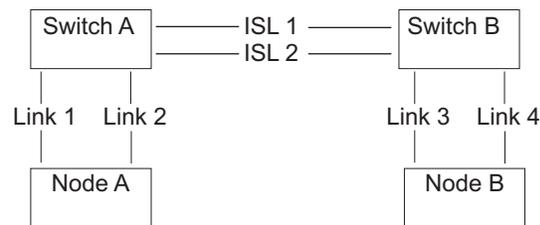


Figure 16. Fabric with ISL between nodes in a cluster

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 17 on page 77.

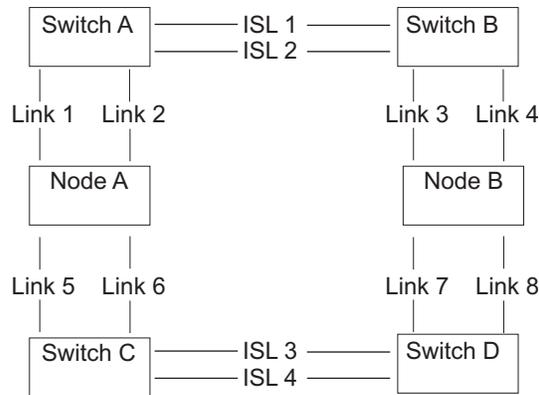


Figure 17. Fabric with ISL in a redundant configuration

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

ISL oversubscription

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

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

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

SAN Volume Controller in a SAN with director class switches

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

Port switches

The following sections provide examples of how to configure your SAN in a small environment.

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.

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.

Zoning guidelines

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

Paths to hosts

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

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

Storage controller zones

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

SAN Volume Controller zones

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

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

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

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

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

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

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

Host zones

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

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

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

Clusters with less than 64 hosts

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

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

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

Clusters with more than 64 hosts

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

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

Example SAN environments

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

Example 1

Consider the SAN environment in the following example:

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

Example 2

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

- Two nodes called A and B
- Nodes A and B have four ports each
 - Node A has ports A0, A1, A2, and A3
 - Node B has ports B0, B1, B2, and B3
- Six hosts called P, Q, R, S, T and U
- Four hosts have four ports each and the other two hosts have two ports each as described in Table 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	—	—
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 T1) of each node and host to switch X.
2. Attach ports 3 (A2, B2, P2, Q2, R2, S2 and T1) and 4 (A3, B3, P3, Q3, R3, S3 and T1) of each node and host to switch Y.
3. Attach ports 1 and 2 (I0 and I1) of the storage controller to switch X.
4. Attach ports 3 and 4 (I2 and I3) of the storage controller to switch Y.

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

Create the following host zones on switch X:

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

Create the following host zones on switch Y:

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

Create the following storage zone:

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

Zoning considerations for Metro Mirror and Global Mirror

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

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

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

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

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

1. Configure your SAN so that fibre-channel traffic can be passed between the two clusters. To configure the SAN this way, you can connect the clusters to the same SAN, merge the SANs, or use routing technologies.
2. Create a new zone in each fabric that contains all SAN Volume Controller ports for both clusters.
3. Optionally, modify the zoning so that the hosts that are visible to the local cluster can recognize the remote cluster. This allows a host to examine data in both the local and remote cluster.
4. Verify that cluster A cannot recognize any of the back-end storage that is owned by cluster B. Two clusters cannot share the same back-end storage devices.

Switch operations over long distances

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

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

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

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 total number of VDisk-to-host mappings is at least 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 (VDI) 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 VDI 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.

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 RAIDs 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 RAIDs 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 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/systems/support/storage/software/sdd/>

5. Install and configure the master console for the SAN Volume Controller. The master console includes two components that enable you to configure the SAN Volume Controller. The first is the SAN Volume Controller Console, which is a Web-based application. The second is PuTTY, which is an SSH client software that enables you to use the command-line interface (CLI). If you purchased the master console hardware option, the IBM service representative installs the master console hardware and you must configure it. If you have the software-only master console option, you must install the master console software on your hardware and then configure it.

Notes:

- a. You can install the master console software on other machines (which you provide) using the CD that contains master console software installation wizard. This CD is provided with both the master console hardware option and the software-only option.
- b. 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. If you have a Microsoft Windows host, install PuTTY as the SSH client. AIX and Linux hosts typically have an SSH client already installed.

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, Metro Mirror, and Global Mirror) objects, as required.

Supported fibre-channel extenders

The supported hardware for the SAN Volume Controller frequently changes.

See the following Web site for the latest supported hardware:

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

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.

Chapter 4. Installing master console software

If you purchased the software-only master console option, you must install the master console software on your own hardware before you can configure and use the SAN Volume Controller. If you have purchased the master console hardware option and the IBM service representative has finished installing it, you can skip the master console software installation instructions and start configuring the master console.

Before you install the master console software, you must ensure that you have met all the prerequisites for installing the software as provided by the *IBM System Storage SAN Volume Controller: Planning Guide*.

In addition, view the release notes for the master console software on the following IBM support Web site for the latest information:

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

You can use the installation wizard to help you install the master console software components instead of installing each component individually.

The following tasks are involved in the master console software installation process:

- **Task 1:** Install the master console software by using the installation wizard.
- **Task 2:** Mirror the boot drive.

Important: Other than antivirus software, ensure that the master console software is the only type of software that is installed on the hardware that you are using for the master console.

Using the master console installation wizard

You use the master console installation wizard to install all the master console components.

Before you install the master console software, ensure that your system meets the requirements that are provided in the *IBM System Storage SAN Volume Controller: Planning Guide*.

During the installation, the wizard installs the following master console components:

- SAN Volume Controller Console and CIM agent
- PuTTY

Perform the following steps to install the master console software:

1. Log on as a local administrator (for example, as the Administrator user) on the system where you want to install the master console software.
2. Insert the master console CD into the CD drive.
3. Click **Start** → **Run** to open the Run window.
4. Enter *drive*:\setup.exe, where *drive* is the letter of the drive in which you inserted the CD. Click **OK**.

The following message is briefly displayed:

```
InstallShield(R) is preparing the InstallShield Wizard, which
will guide you through the rest of the process.
Please wait .....

Preparing Java (tm) Virtual Machine
```

You are then prompted to select the language to use for the installation wizard.

5. Select the language to use, and then click **OK**.
The Welcome panel is displayed.
6. Read the information on the Welcome panel, and then click **Next**.
The License Agreement panel is displayed.
7. Read the license agreement and perform one of the following actions:
 - Click **I accept the terms in the license agreement**, and then click **Next** to continue with the installation.
 - Click **I do not accept the terms of the license agreement**, and then click **Cancel** to exit the installation.

When you click **Next**, the wizard verifies that your system meets the hardware requirements for the installation.

Note: If your system does not meet hardware requirements, the wizard opens a panel that warns you about a decrease in the performance level if these requirements are not met. Click **OK** to close the warning panel.

A panel is displayed for you to select the destination directory for your master console software installation. By default, the master console software is installed in C:\Program Files\IBM\MasterConsole

8. Click **Next** to accept the default directory. If you want to change the directory, click **Browse**, select a different directory, and then click **Next**.
The installation wizard compares the list of components to be installed with the products that are already installed on your system. If the wizard finds that any master console components are already installed, it compares the versions and uses the following logic to determine which components to install:
 - If the component is not installed or the installed version is an earlier version than the required version, the component is installed or upgraded by launching the component-specific installation program.
 - If the component is installed at the same version level as the version that is installed by the installation wizard, the wizard does not install the component.
 - If the installed component is a later version than the version to be installed, the installation wizard does not install the component, but a warning that indicates that this version is not tested with the master console is displayed. Click **OK** if you see this warning. Decide whether you want to continue the installation or exit the installation so that you can first remove the later version from your system. If you remove the later version, restart the master console installation wizard after the removal.
 - If a component is incorrectly installed on the system, you are asked to continue with the installation by reinstalling the component with the component-specific installer. If the reinstallation does not succeed, you must exit the master console installation wizard, manually remove the product from your system, and restart the master console installation wizard.

After the wizard performs the comparison, the Product List panel is displayed. This panel provides the following information:

- Versions of existing master console components
- Required versions
- Actions to be done by the installation wizard or by you

If you need to refer to the installation product list later, this information is saved on your system in the *MasterConsoleProducts.htm* file. This file is located in the directory that you specified for the master console installation.

9. After you have reviewed the information presented in the Product List panel, click **Next**.

The PuTTY installation panel is displayed.

10. Click **Next** to begin the installation of PuTTY.

11. Click **Next** when the wizard indicates that PuTTY has been installed.

The installation wizard validates the installation of PuTTY. If the validation is not successful, an error panel is displayed and you must exit the wizard, correct the errors, and restart the installation wizard.

12. Before continuing with the installation wizard, create an SSH key pair using PuTTY. Follow the instructions that are described in the “Generating an SSH key pair using PuTTY” on page 98 topic. After you have created the SSH key pair, return to these instructions.

The next part of the installation wizard uses the generated SSH key pair to install the SAN Volume Controller Console.

13. Click **Next** to launch the SAN Volume Controller Console installation program.

The SAN Volume Controller Console installation program opens in a new window.

Important: Do not close the main master console installation wizard panel while you complete the SAN Volume Controller Console installation program.

14. From the SAN Volume Controller Console installation program, perform the following steps to install the SAN Volume Controller Console:

- a. Click **Next** in the Welcome panel.
- b. Click **I accept the terms of the license agreement**, and then click **Next** in the License Agreement panel.
- c. Click **Next** in the Destination Directory panel to accept the default directory.
- d. Enter the private key that you created in step 12 (the default key name is *icat.ppk*), and then click **Next**.
- e. Click **Next** in the CIMOM ports panel to accept the default ports.
- f. Click **Next** to accept the default ports for the embedded WebSphere® Application Server.
- g. Click **Install** in the installation confirmation panel to install the SAN Volume Controller Console.
- h. (Optional) Click **View post installation tasks** to view the steps that you must later perform to access the SAN Volume Controller Console.
- i. Click **Finish** to complete the installation and close the SAN Volume Controller Console installation program.

The SAN Volume Controller Console installation program closes and control is returned to the main master console installation wizard.

15. In the main master console installation wizard, click **Next**.
The wizard validates the installation of the SAN Volume Controller Console. If the validation is successful, the Finish panel is displayed. If the validation fails, an error panel is displayed. If errors are found, exit the wizard, correct the errors, and restart the master console installation wizard.
16. Review the master console installation log (mclog.txt) to ensure that all components are properly installed. The log file is located in *installation_directory*\logs, where *installation_directory* is the directory where the master console was installed. The default installation directory is C:\Program Files\IBM\MasterConsole.
17. Click **Finish** to complete the installation.
18. If a system restart is required, accept the prompt to complete the master console installation process.

Mirroring the boot drive

You can use the Microsoft® Windows software mirroring capability to mirror the boot drive on the master console system.

Note: The master console hardware option is delivered with the boot drive already mirrored.

Before you mirror the boot drive, ensure that you meet the following prerequisites:

- You must have a second drive that is as large or larger than your original boot drive.
- If the target disk has a partition that is assigned to it (that is, it already has a drive letter), perform the following steps to remove the partition:
 1. Back up any necessary data on the existing partition. When you remove an existing partition, you lose any data that is on it.
 2. Remove the partitioning by right-clicking **My Computer**, selecting **Manage** → **Storage** → **Disk Management**, right-clicking the target disk drive, and then selecting **Delete Partition**.

Perform the following steps to mirror the boot drive:

1. Right-click **My Computer** on the desktop.
2. Click **Manage**.
3. Click **Storage** → **Disk Management**.
4. Right-click the disk icon for the system disk.
5. Convert the disk to a dynamic disk by performing the following actions:
 - For Windows 2000 systems:**
 - a. Click **Upgrade to dynamic disk** and then click **OK**.
 - b. Right-click the disk icon for the disk that you want to become the mirror of the system disk.
 - c. Click **Upgrade to dynamic disk** and then click **OK**.
 - d. If you receive a warning, click **Yes**.
 - e. If your system restarts, restart the Disk Management utility.
 - For Windows 2003 systems:**
 - a. Click **Convert to dynamic disk**.
 - b. Select both drives and click **OK**.
 - c. Click **Convert**.

- d. Click **Yes** to continue when you receive an operating system boot warning.
- e. Click **Yes** to unmount file systems and to continue.
- f. Allow the system to reboot.
- g. Restart the Disk Management utility
6. For Windows 2000 systems, right-click the disk icon for the system disk or, for Windows 2003 systems, right-click the system disk partition.
7. Click **Add Mirror**.
8. Select the disk that you want to become the mirror of the system disk, and then click **Add Mirror**.
9. Click **OK** to continue when you receive a warning indicating that you must update the boot.ini file.
10. Perform the following steps to update the boot.ini file:
 - a. Double-click **My Computer** on the desktop.
 - b. Click **Tools** → **Folder options**.
 - c. Click the **View** tab.
 - d. In the **Advanced settings** list, perform one of the following actions, depending on your operating system:
 - For Windows 2000 systems, select **Show hidden files and folders**.
 - For Windows 2003 systems, select **Show hidden files and folders** and clear the **Hide protected operating system files** option. Click **Yes** when you receive a warning.
 - e. In the My Computer window, click **Local disk (C:)**.
 - f. Open the C:\boot.ini file in a text editor.

Attention: Be cautious when you edit this file and make only the specified changes. Do not modify any other lines in this file.
 - g. In the operating system section, add **Primary** to the end of the operating description for the system disk.
 - h. Copy the line for the system disk, change **Primary** to **Secondary**, and change the system disk (for example, **rdisk(0)**) to mirrored disk (for example, **rdisk(1)**). The file looks similar to one of following examples.

Windows 2000 example:

```
[boot loader]
timeout=30 default=multi(0)disk(0)rdisk(0)partition(1)\WINNT
[operating systems]
multi(0)disk(0)rdisk(0)partition(1)\WINNT="Microsoft Windows 2000 Advanced
  Server Primary" /fastdetect
multi(0)disk(0)rdisk(1)partition(1)\WINNT="Microsoft Windows 2000 Advanced
  Server Secondary" /fastdetect
```

Windows 2003 example:

```
[boot loader]
timeout=30
default=multi(0)disk(0)rdisk(0)partition(1)\WINDOWS
[operating systems]
multi(0)disk(0)rdisk(0)partition(1)\WINDOWS="Windows Server 2003,
  Standard Primary" /fastdetect
multi(0)disk(0)rdisk(1)partition(1)\WINDOWS="Windows Server 2003,
  Standard Secondary" /fastdetect
```

- i. Save and close the file.
- j. Restart the machine.

After the machine completes the Power On Self Test (POST), the system prompts you to select the Windows operating system to use for the boot process

- k. Select the secondary operating system and press Enter to verify that the machine successfully restarts from the mirrored drive.
- l. Restart the machine again.

After the machine completes the POST, the system again prompts you to select the Microsoft Windows operating system to use for the boot process.

- m. Select your operating system and press Enter to verify that the machine successfully restarts from the system drive.

Chapter 5. Configuring the master console

You can configure the master console to access the SAN Volume Controller Console and the SAN Volume Controller command-line interface (CLI). If you installed the master console on your own hardware, you have already performed some of these steps during the installation process.

If you purchased the hardware master console and experience a problem, use the 2145 machine type and the serial number of a SAN Volume Controller node that was installed with the master console to open a hardware problem.

Perform the following steps to configure the master console:

1. Log on as a local administrator (for example, as the Administrator user) to the system where the master console software is installed.

Note: If you installed the software master console, skip to step 3 because you already performed the tasks described in step 2 before or during the installation of the master console software.

2. If you purchased a hardware master console, perform the following configuration steps:
 - a. Optionally, reconfigure the master console host name. When you receive the hardware master console, the host name is preconfigured as mannode. If you choose to change this name see “Changing the master console host name” for more information.
 - b. Configure the internal IP network connection (Local Area Network). “Configuring the internal IP network connection” on page 96 provides more details for this step.
 - c. Configure the browser. “Configuring the Web browser” on page 96 provides more details for this step.
 - d. Generate an SSH key pair using the PuTTYgen. “Generating an SSH key pair using PuTTY” on page 98 provides more details for this step.
3. For a software master console or a hardware master console, perform the following configuration steps:
 - a. Configure a default PuTTY session for command-line interface (CLI) access. “Configuring the PuTTY session for the CLI” on page 99 provides more details for this step.
 - b. Store keys in the SAN Volume Controller Console software. “Storing the private SSH key in the SAN Volume Controller Console software” on page 99 provides more details for this step.
 - c. Install your chosen antivirus software on the master console system.

Changing the master console host name

You can change the master console host name anytime. When you change the host name, you must also be sure that other master console applications are updated to use the new name.

Perform the following steps to change the host name and to update the name in other master console applications:

1. Right-click on **My Computer** from the desktop.

2. Click **Properties**.
3. Click **Computer Name**.
4. Click **Change**.
5. Type the master console host name in the **Computer name** field.
6. Click **More**.
7. Type the full path information in the **Primary DNS suffix of this computer** field.
8. Click **OK** until you return to the desktop.
9. Click **Yes** to restart the master console system so that the change to the host name is applied.

Configuring the internal IP network connection

Before you can use the master console, you must configure the internal IP network connection.

Perform the following steps to configure the Local Area Connection:

1. From the desktop, right-click **My Network Places**.
2. Click **Properties**.
3. Right-click **Local Area Connection**.
4. Click **Properties**.
5. Click **Internet Protocol (TCP/IP)**.
6. Click **Properties**.
7. Type all required information for the IP and DNS addresses.

Note: You do not have to use a static TCP/IP address. If you only want to access the master console directly, you can use a DHCP TCP/IP address. If you use a DHCP TCP/IP address, ensure that the properties are set to DHCP. To access the master console remotely, you must use a static IP address.

8. Click **OK** until you return to the desktop.
9. Connect the Ethernet port to the network.

Configuring the Web browser

You must configure your Web browser so that you can use it to access the SAN Volume Controller Console.

Before you configure the Web browser, ensure that you have one of the following required browsers:

- For Windows operating systems:
 - Netscape version 6.2
Netscape is available from <http://browser.netscape.com/ns8/download/archive.jsp>.
 - Internet Explorer Version 6 or later
You can get version 6 or later from <http://www.microsoft.com/windows/products/winfamily/ie/>.
- For AIX operating system:
 - AIX Netscape version 7.0

You can get AIX Netscape version 7.0 from <http://browser.netscape.com/ns8/download/archive.jsp>.

After you confirm that you have the correct Web browser, perform the following steps to configure it:

1. Ensure that the Web browser is not set to block or suppress popup windows.

Note: If you are using Internet Explorer 7.0 and receive a message that a pop-up has been blocked, click the Information Bar at the top of the browser and select **Always allow popups** from this site. If you receive a message that content was blocked because it was not signed by a valid security certificate, click the Information Bar at the top of the screen and select **Show blocked** content.

2. Ensure that you have not installed any 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.
3. Disable the proxy setting by performing the following steps:

For Netscape:

- a. Open your Netscape browser and click **Edit** → **Preferences**. The Preferences window displays.
- b. From the left side category, click **Advanced** to expand the suboptions. The suboption Proxies displays.
- c. Click **Proxies**. The Proxies window displays.
- d. Select **Direct connection to Internet**.

For Internet Explorer:

- a. Click **Tools** → **Internet Options** → **Connections** → **LAN Settings**.
- b. Click to clear the **Use a proxy server** box.
4. (Optional) Perform the following steps to add password protection so that your password does not display when you type it in:

For Netscape:

- 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**.

For Internet Explorer:

- 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**.

Generating an SSH key pair using PuTTY

You must generate a Secure Shell (SSH) key pair to use the SAN Volume Controller Console and the command-line interface (CLI).

Perform the following steps to generate SSH keys on the master console using the PuTTY key generator (PuTTYgen):

1. Start PuTTYgen by clicking **Start** → **Programs** → **PuTTY** → **PuTTYgen**. The PuTTY Key Generator panel is displayed.
2. Click **SSH-2 RSA** as the type of key to generate.

Note: Leave the number of bits in a generated key value at 1024.

3. Click **Generate** and then move the cursor around the blank area of the Key section to generate the random characters that create a unique key. When the key has been completely generated, the information about the new key is displayed in the Key section.

Attention: Do not modify the Key fingerprint or the Key comment fields; this can cause your key to no longer be valid.

4. (Optional) If you are generating SSH keys for a computer other than the master console, enter a passphrase in the Key passphrase and Confirm passphrase fields. The passphrase encrypts the key on the disk; therefore, it is not possible to use the key without first entering the passphrase.

Attention: If you are generating the key pair for the master console, do not enter anything in the Key passphrase or the Confirm passphrase fields.

5. Save the public key by performing the following steps:
 - a. Click **Save public key**. You are prompted for the name and location of the public key.
 - b. Type `icat.pub` as the name of the public key and specify the location where you want to save the public key.

Note: For AIX, store the key in the `$HOME/.ssh` directory.

- c. Click **Save**.
6. Save the private key by performing the following steps:
 - a. Click **Save private key**. The PuTTYgen Warning panel is displayed.
 - b. Click **Yes** to save the private key without a passphrase.
 - c. Type `icat` as the name of the private key, and specify the location where you want to save the private key.
 - d. Click **Save**.
 - e. For AIX, store the key in the `$HOME.ssh/identity` file, which is in the `$HOME/.ssh` directory. In the simplest cases, you can replace the contents of the identity file with the contents of the key file. However, when you are using multiple keys, all of these keys must appear in the identity file.
 7. Close the PuTTY Key Generator window.

After the SSH key pair is generated, you must configure the SAN Volume Controller Console to use them. For the private key, you must copy the private key to the directory that is used by the IBM CIM Object Manager. For example, `C:\Program Files\IBM\svcconsole\cimom`. For the public key, you use the Add SSH Public Key panel in the SAN Volume Controller Console when you create the cluster.

Storing the private SSH key in the SAN Volume Controller Console software

For both the software and hardware versions of the master console, when the keys that are used to communicate with the SAN Volume Controller are generated or changed, you must store a copy of the new private key in the SAN Volume Controller Console software.

Perform the following steps to store a copy of the new private key in the SAN Volume Controller Console software:

1. Open a command prompt window.
2. Type the following command:

```
copy path\icat.ppk C:\"Program Files"\IBM\succonsole\cimom
```

where *path* is the path where you stored the SSH private key when it was generated and *C:\"Program Files"\IBM\succonsole\cimom* is the location where you installed the SAN Volume Controller Console.

Important: Directory names with embedded spaces must be surrounded by double quotation marks.

3. Close the command prompt window.
4. Perform the following steps to stop and restart the IBM CIM Object Manager to save the changes:
 - a. Click **Start** → **Settings** → **Control Panel** → **Administrative Tools** → **Services**. The Services panel is displayed.
 - b. Double-click **Administrative Tools**.
 - c. Double-click **Services**.
 - d. From the list of services, right-click **IBM CIM Object Manager - SVC**. Select **Stop** and wait for Windows to stop the service. The Service Control panel is displayed to indicate that the IBM CIM Object Manager service has stopped.
 - e. From the list of services, right-click **IBM CIM Object Manager - SVC**. Select **Start**. The Service Control panel is displayed to indicate that the IBM CIM Object Manager service has started.

Configuring the PuTTY session for the CLI

You must configure the PuTTY session using the Secure Shell (SSH) key pair that you have generated before you can use the command-line interface (CLI).

Attention: Do not run scripts that create child processes that run in the background and invoke SAN Volume Controller commands. This can cause the system to lose access to data and cause data to be lost.

Perform the following steps to configure a PuTTY session for the CLI:

1. Select **Start** → **Programs** → **PuTTY** → **PuTTY**. The PuTTY Configuration window opens.
2. Click **Session** in the Category navigation tree. The Basic options for your PuTTY session are displayed.
3. Click **SSH** as the Protocol option.
4. Click **Only on clean exit** as the Close window on exit option. This ensures that connection errors are displayed.

5. Click **Connection** → **SSH** in the Category navigation tree. The options controlling SSH connections are displayed.
6. Click **2** as the Preferred SSH protocol version.
7. Click **Connection** → **SSH** → **Auth** in the Category navigation tree. The Options controller SSH authentication are displayed.
8. Click **Browse** or type the fully-qualified file name and location of the SSH client and private key in the **Private key file for authentication** field. The file that you specify in this field is the one that you stored in the SAN Volume Controller software (for example, C:\Program Files\IBM\svconconsole\cimom\icat.ppk).
9. Click **Session** in the Category navigation tree. The Basic options for your PuTTY session are displayed.
10. Click **Default Settings** and then click **Save**.
11. Type the name or IP address of the SAN Volume Controller cluster in the **Host Name (or IP Address)** field.
12. Type **22** in the **Port** field. The SAN Volume Controller cluster uses the standard SSH port.
13. Type the name that you want to use to associate with this session in the **Saved Sessions** field. For example, you can name the session SAN Volume Controller Cluster 1.
14. Click **Save**.

You have now configured a PuTTY session for the CLI.

Chapter 6. Creating a SAN Volume Controller cluster

After you generate a Secure Shell (SSH) key pair, you must complete the two phases that are required to create a cluster before you can configure the SAN Volume Controller.

The first phase to create a cluster 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.

Creating a cluster from the front panel

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

Note: Before you create a cluster, ensure that you have followed the steps in the “Verifying the SAN Volume Controller installation” topic in the *IBM System Storage SAN Volume Controller: Hardware Installation Guide*.

A pair of nodes must be installed before you can create a cluster. If you choose to have the IBM service representative or IBM Business Partner initially create the cluster, you must provide the following information 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 IBM service representative or IBM Business Partner uses the front panel of the node to enter the information that you have provided. The node generates a random password on the display panel. The IBM service representative or IBM Business Partner gives you this password. You must record the password and the IP address. The password and IP address are used to connect to the node 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 that you are creating.

Note: You can add additional nodes after you have successfully created and initialized the cluster.

2. Press and release the up or down button until **Node:** is displayed on the node service panel.
3. Press and release the left or right button until **Create Cluster?** is displayed.
4. Press and release the select button.
 - If **IP Address:** is displayed on line 1 of the screen, go to step 6 on page 102.

- If Delete Cluster? is displayed on the first line 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 no action is performed.
- 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 button.
 - b. Press and release the select button. The node restarts. After the node has restarted, you must restart this task from step 1 on page 101. IP Address: is displayed on the service display screen.

5. Press and release the select button.
6. Use the up or down button to quickly increase or decrease the value of the first field of the IP Address to the value that you have chosen.

Note: If you want to disable the fast increase/decrease function, press and hold the down button, press and release the select button, and then release the down button. The disabling of the fast increase/decrease function lasts until cluster creation is completed or until the feature is again enabled. If the up or down buttons are pressed and held while the function is disabled, the value increases or decreases once every two seconds. To again enable the fast increase/decrease function, press and hold the up button, press and release the select button, and then release the up button.

7. Use the right button to move to the next field and use the up or down button to change the value of this field.
8. Repeat step 7 for each of the remaining fields of the IP Address.
9. Press the select button when you have changed the last field of the IP Address.
10. Press the right button. Subnet Mask: is displayed.
11. Press the select button.
12. Use the up or down button to quickly increase or decrease the value of the first field of the Subnet Mask to the value that you have chosen.
13. Use the right button to move to the next field and use the up or down buttons to change the value of this field.
14. Repeat step 13 for each of the remaining fields of the Subnet Mask.
15. Press the select button when you have changed the last field of Subnet Mask.
16. Press the right button. Gateway: is displayed.
17. Press the select button.
18. Use the up or down button to quickly increase or decrease the value of the first field of the Gateway to the value that you have chosen.
19. Use the right button to move to the next field and use the up or down button to change the value of this field.
20. Repeat step 19 for each of the remaining fields of the Gateway.
21. Press the select button when you have changed the last field of Gateway.
22. Press and release the right button until Create Now? is displayed.
23. Ensure that you have a pen and paper to record the cluster password before you create the cluster.

Attention: The password is displayed for 60 seconds, or until the up, down, left or right button is pressed, which deletes it. You must be ready to record the password before you select one of the following actions:

- If you want to review your settings before you create the cluster, use the right and left buttons to review those settings. Make any necessary changes, return to *Create Now?*, then press the select button.
- If you are satisfied with your settings, 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.

Important: If you do not record the password, you must restart the cluster configuration procedure. After you have recorded the password, press the up, down, left, or right 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 18 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.

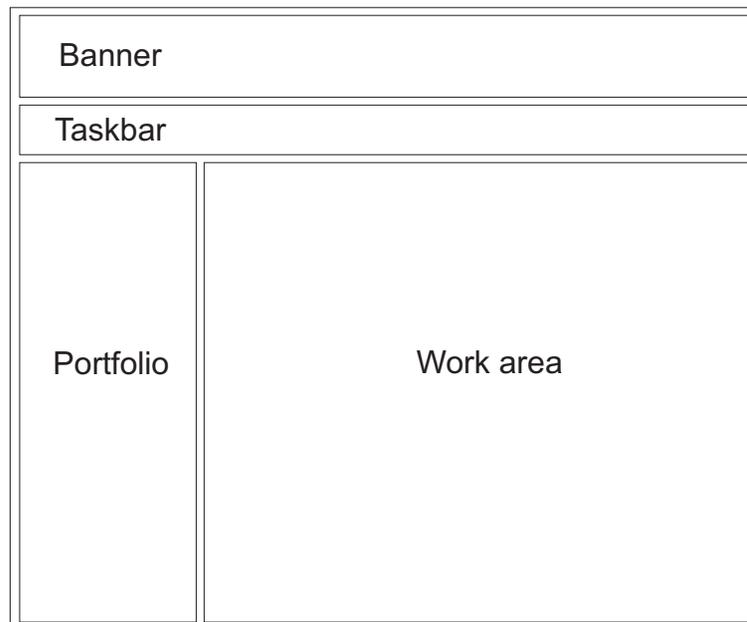


Figure 18. Basic frame layout

SAN Volume Controller Console banner

The banner of the SAN Volume Controller Console provides product or customer identification.

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 19 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 19. 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 Console:

- Welcome
- Manage Cluster
 - View Cluster Properties
 - Maintain Cluster 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
- Discovery Status
- Managed Disks
- Managed Disk Groups
- Work with Virtual Disks
 - Virtual Disks
 - Virtual Disk-to-Host Mappings
- Work with Hosts
 - Hosts
 - Fabrics
- Manage Copy Services
 - FlashCopy Mappings
 - FlashCopy Consistency Groups
 - Metro & Global Mirror Relationships
 - Metro & Global Mirror Consistency Groups
 - Metro & Global Mirror Cluster Partnership
- Service and Maintenance
 - Upgrade Software
 - Run Maintenance Procedures
 - Set Error Notification
 - Analyze Error Log
 - License Settings
 - View License Settings 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.

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. If you are using Internet Explorer 7.0 and receive a message that a popup has been blocked, click the Information Bar at the top of the browser and select **Always allow popups from this site**. If you receive a message that content was blocked because it was not signed by a valid security certificate, click the Information Bar at the top of the screen and select **Show blocked content**.

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.

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

`http://localhost:9080/ica`

If a SAN Volume Controller Console icon is present on your desktop, you can double click this icon to access the console.

Log on to the SAN Volume Controller Console using the superuser user name, which is superuser, and the superuser password, which is passw0rd. The first time that 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 the cluster using the SAN Volume Controller service panel, you can use the Add SAN Volume Controller Cluster function from the SAN Volume Controller Console to identify the cluster to the master console.

You must generate an SSH key pair before you can use the SAN Volume Controller Console 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 IBM System Storage SAN Volume Controller Signon window is displayed.

Note: If you are using Internet Explorer 7.0 and receive a message that a popup has been blocked, click the Information Bar at the top of the browser and select **Always allow popups from this site**. If you receive a message that content was blocked because it was not signed by a valid security certificate, click the Information Bar at the top of the screen and select **Show blocked content**.

2. Type superuser for the user ID and passw0rd 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 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, proceed to step 17 on page 108.
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 if you cannot access the cluster using the administrator ID and password.
 - 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.
 - f. If you want the ability 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 to collect SNMP traps, type the IP address of the master console here. The IBM SAN Volume Controller support Web site at <http://www.ibm.com/storage/support/2145> provides additional information about using the IBM Director with the SAN Volume Controller.

- c. Type the SNMP community name.
 - d. Click **Update Settings** to continue.
16. Click **Continue**. The License Settings panel is displayed. The allowed setting for each of the parameters is specified in your user's license.
- a. Enable the FlashCopy, Metro Mirror, or Global 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 License Settings**.
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 that you gave during step 14 on page 107.
 - b. Click **Browse** to locate the public key for the master console. This key should be located in the same directory as the private key.
 - c. Type an ID (label) for this key.
 - d. Select **Give Administrator role to administrator access key (default is Monitor)**.
 - e. 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.

You can now set up cluster properties, add additional nodes to the cluster, create managed disk groups, identify virtual disks, create and map host objects.

Chapter 7. 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.
 - Mirror and Mirror consistency groups.

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

Launching the SAN Volume Controller Console to manage a cluster

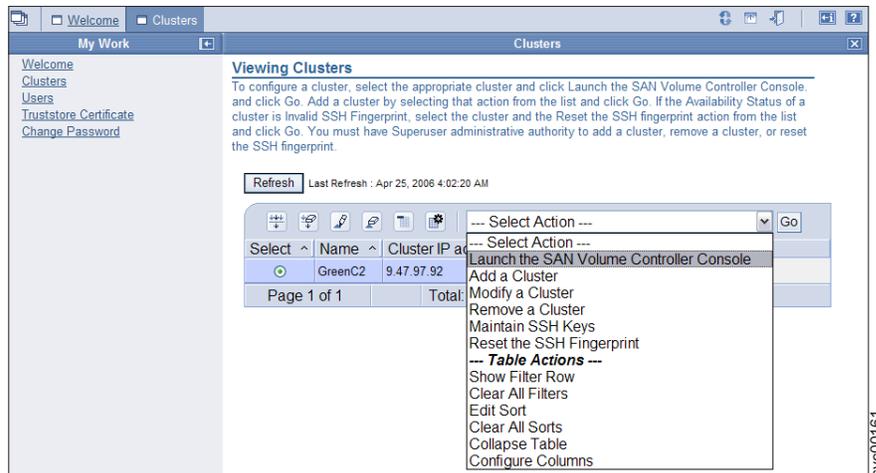
You can launch the SAN Volume Controller Console from the Viewing Clusters panel.

The SAN Volume Controller Console is the centralized Web application that is used to manage your clusters. It is preinstalled on the hardware master console or you install it when you install the software master console on your own hardware.

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 Console for a specific cluster:

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 Console** from the task list.



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

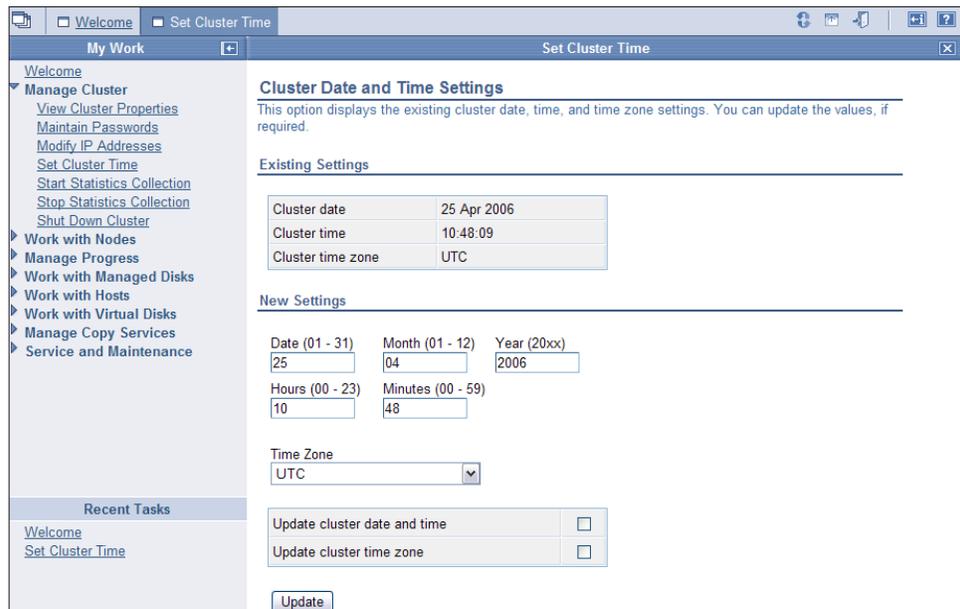
Setting the cluster date and time

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

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

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 Console.

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 Console.

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 Console.

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 & Global Mirror** to view the Metro Mirror or Global 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 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 that is provided by these SCSI and fibre-channel device drivers uses a combination of a SCSI logical unit number (LUN) and the worldwide node name (WWNN) 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 WWNN 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 WWNN 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

Attention:

1. 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 the IBM Support Center to add the node back into the cluster without corrupting the data.
2. 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 ensure that the LUNs are the same before you add the new node to the cluster.
3. LUN masking for each LUN must be identical on all nodes in a cluster. You must ensure that the LUN masking for each LUN is identical before you add the new node to the cluster.
4. You must ensure that the model type of the new node is supported by the SAN Volume Controller software level that is currently installed on the cluster. If the model type is not supported by the SAN Volume Controller software level, upgrade the cluster to a software level that supports the model type of the new node. See the following Web site for the latest supported software levels:
<http://www.ibm.com/storage/support/2145>

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 Console.

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**.

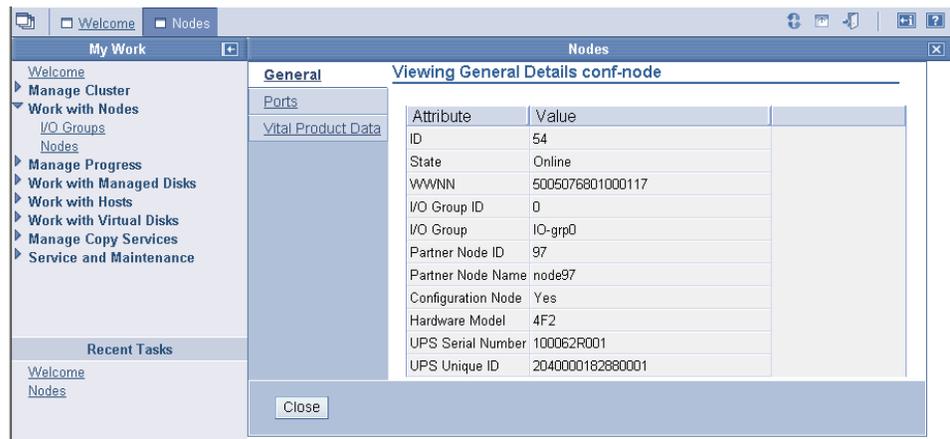
Viewing the node status

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

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

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.



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 the IBM Support Center to add the node to the cluster without corrupting data.

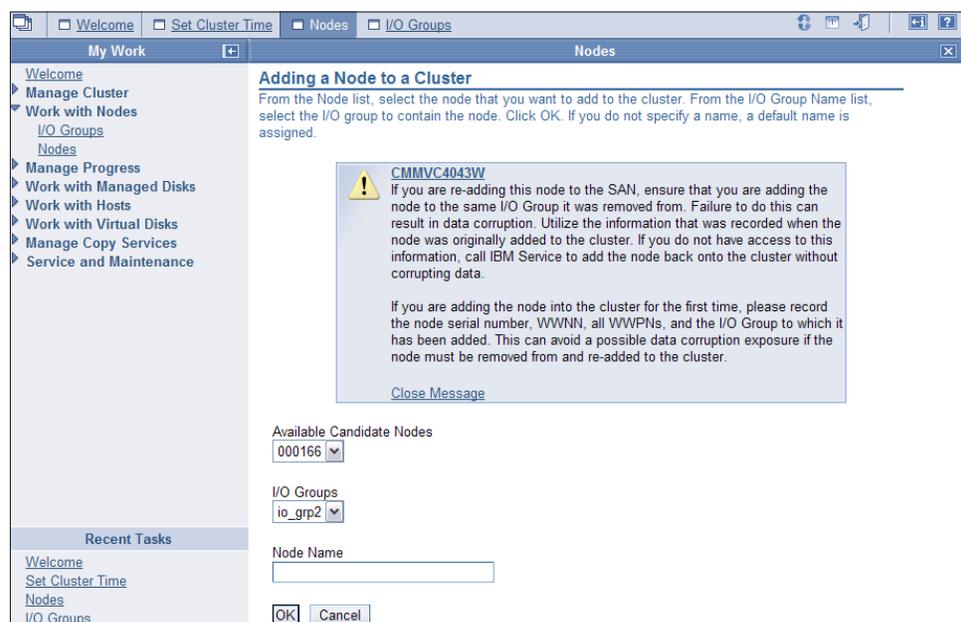
If you want to add a node that was previously removed from a cluster, you must have the following information about the node:

- 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

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

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 on page 116.

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
 - WWPNN
 - 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.

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

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 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 **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 System Storage 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 attributes	Description												
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.												
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" data-bbox="740 835 1450 997"> <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 Console.

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.

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 System Storage 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.

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, 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 System Storage 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 System Storage 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 Console.

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 Console.

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. Turn 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 turn 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.

Select	ID	Name	Status	World Wide Node Name (WWNN)	I/O Group Name	Config Node
<input type="radio"/>	1	node1	Online	500507680100018C	io_grp0	Yes
<input type="radio"/>	2	node2	Online	5005076801000173	io_grp0	No

Page 1 of 1 Total: 2 Filtered: 2 Displayed: 2 Selected: 0

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 Console.

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**.

Modifying a cluster

You can rename a cluster and change the fabric speed from the Modifying Cluster panel.

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

Perform the following steps to modify a cluster:

1. Click **Clusters** in the portfolio. The Viewing Clusters panel is displayed.
2. Select the cluster that you want to modify and select **Modify a Cluster** from the task list. Click **Go**. The Modifying Cluster panel is displayed. You can perform the following from this panel:
 - Type a new name for the cluster.
 - Select a fabric speed from the **Fabric Speed** list.
3. Click **OK** to modify the cluster.

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 VDIs 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.

Begin the following process of quiescing all I/O to the cluster by stopping the applications on your hosts that are using the VDIs that are provided by the cluster.

1. Determine which hosts are using the VDIs that are provided by the cluster.

2. Repeat the previous step for all VDisks.

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.

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

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.

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

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 task 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.

Viewing discovery status

You can view the status of a managed disk (MDisk) discovery from the Viewing Discovery Status panel.

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

Perform the following steps to view status of an MDisk discovery:

1. Click **Work with Managed Disks** → **Discovery Status**. The Viewing Discovery Status panel is displayed.
2. Click **Close** to close this 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 Console.

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 Console.

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 Console.

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 Console.

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 Console.

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. 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 Console.

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 Console.

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 Console.

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 Console.

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 a managed disk (MDisk) removal from the Viewing MDisk Removal Progress panel.

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

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 Console.

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 Console.

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 Console.

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 that 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 VDIs

You can create virtual disks (VDIs) using the Create Virtual Disks wizard.

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

Perform the following steps to create VDIs:

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 VDIs** 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 Console.

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.

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

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 Console.

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 use the SAN Volume Controller Console to shrink a virtual disk (VDisk).

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:

1. 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, Metro Mirror relationships, or Global Mirror relationships.
2. 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 Console.

Perform the following steps to shrink 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 that you want to shrink and select **Shrink a VDisk** from the task list. Click **Go**. The Shrinking Virtual Disks panel is displayed.
4. Enter the capacity that you want to reduce the size of the VDisk by in the **Reduce By Capacity** field and click **OK**.

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 Console.

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 Console.

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 VDisk that you want to view.

4. Select **Show MDisks This VDisk is Using** from the task list and click **Go**. The Viewing Managed Disks panel is displayed. This panel lists the MDisks that the selected VDisk uses.

Recovering from offline VDIsks

You can use the SAN Volume Controller Console to recover from an offline virtual disk (VDisk) 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 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 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. Delete all FlashCopy, Metro Mirror, and Global Mirror mappings and relationships that use the offline VDIsks.
3. Move all the offline VDIsks to the recovery I/O group.
4. Move all the offline VDIsks back to their original I/O group.
5. Recreate all FlashCopy, Metro Mirror, and Global Mirror mappings and relationships that use the VDIsks.

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. Delete all FlashCopy, Metro Mirror, and Global Mirror mappings and relationships that use the offline VDIsks.
2. Move all the offline VDIsks to the recovery I/O group.
3. Move both recovered nodes back into the cluster.
4. Move all the offline VDIsks back to their original I/O group.
5. Recreate all FlashCopy, Metro Mirror, and Global Mirror mappings and relationships that use the VDIsks.

Moving offline VDIsks to the recovery I/O group

After a node or an I/O group fails, you can move the offline virtual disks (VDIsks) to the recovery I/O group.

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

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 Console.

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 (VDI) from the Deleting Virtual Disk panel.

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

Perform the following steps to delete a VDI:

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 you want to delete and select **Delete a VDI** 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 MDisks that already contain data that was written directly, not through a SAN Volume Controller node. Image mode VDIs have a minimum size of 1 block (512 bytes) and always occupy at least one extent.

Image mode MDisks are members of an MDisk group but, they do not contribute to free extents. Image mode VDIs are not affected by the state of the MDisk group because the MDisk group controls image mode VDIs through the VDIs association to an MDisk. Therefore, if an MDisk that is associated with an image mode VDisk is online and the MDisk group of which they are members goes offline, the image mode VDisk remains online. Conversely, the state of an MDisk group is not affected by the state of the image mode VDIs in the group.

An image mode VDisk behaves just as a managed mode VDisk in terms of the Metro Mirror, Global Mirror, and FlashCopy Copy Services. 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 (VDIs).

Make sure that you are aware of the following before you create image mode VDIs:

- 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/storage/support/2145>

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

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**.

After the image mode VDisk is mapped to a host object, it 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 Console.

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 Console.

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 Console.

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 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 Console.

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 host details

You can view details about a host object from the Viewing General Details panel.

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

Perform the following steps to view 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 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 Console.

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 Console.

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.

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

If a large number of new VDisks are mapped to a host and a large 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.

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 Console.

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
 - Port mask
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 Console.

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 Console.

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 name (WWPN) 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 Console.

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 System Storage 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 Console.

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.

Viewing fabrics

You can view the fabrics that are associated with a cluster from the Viewing Fabrics panel.

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

Perform the following steps to view the fabrics:

1. Click **Work with Hosts** → **Fabrics**. The Viewing Fabrics panel is displayed.
2. Click **Close** to close the panel.

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 Console.

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 Console.

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 Starting FlashCopy Mappings panel.

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

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 Console.

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 Stopping FlashCopy Mappings panel.

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

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 Console.

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 Console.

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 FlashCopy 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 Console.

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 control the FlashCopy consistency groups that are displayed on the panel.

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

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 a FlashCopy consistency group from the Starting FlashCopy Consistency Group panel.

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

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 Viewing FlashCopy Consistency Groups panel is displayed.
3. Select the appropriate group'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 Console.

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 Viewing FlashCopy Consistency Groups panel is displayed.
3. Select the appropriate group's row from the table.
4. Select **Stop a Consistency Group** from the task list and click **Go**. The FlashCopy 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 Console.

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 the appropriate group row from the table.
4. 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 Console.

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 group's row from the table.
4. Select **Delete a Consistency Group** from the task list and click **Go**. The Delete FlashCopy Consistency Groups panel is displayed.

Creating Metro Mirror and Global Mirror relationships

You can use the SAN Volume Controller Console to create Metro Mirror and Global Mirror relationships.

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

Perform the following steps to create a Metro Mirror or Global Mirror relationship:

1. Click **Manage Copy Services** → **Metro & Global Mirror Relationships** in the portfolio. The Filtering Metro & Global 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 & Global Mirror Relationships panel is displayed.
3. Select **Create a Relationship** from the list and click **Go**. The Create a Metro or Global Mirror Relationship wizard. begins.
4. Complete the Create a Metro or Global Mirror Relationship wizard.

Filtering Metro Mirror and Global Mirror relationships

You can use the SAN Volume Controller Console 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 Console.

Perform the following steps to specify some filter criteria:

1. Click **Manage Copy Services** → **Metro & Global Mirror Relationships** in the portfolio. The Filtering Metro & Global 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 & Global Mirror Relationships panel is displayed.

Starting a Metro Mirror or Global Mirror copy process

You can use the SAN Volume Controller Console to start a Metro Mirror or Global Mirror copy process.

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

Perform the following steps to start a Metro Mirror or Global Mirror copy process:

1. Click **Manage Copy Services** → **Metro & Global Mirror Relationships** in the portfolio. The Filtering Metro & Global 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 & Global 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 and Global Mirror copy processes

You can use the SAN Volume Controller Console to view the progress of Metro Mirror and Global Mirror copy processes.

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

Perform the following steps to view the progress of Metro Mirror and Global Mirror copy processes:

1. Click **Manage Progress** → **View Progress**. The View Progress panel is displayed.
2. Click the **Mirror** link. The Viewing Mirror Progress panel is displayed.

Stopping a Metro Mirror or Global Mirror copy process

You can use the SAN Volume Controller Console to stop a Metro Mirror or Global Mirror copy process.

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

Perform the following steps to stop a Metro Mirror or Global Mirror copy process:

1. Click **Manage Copy Services** → **Metro & Global Mirror Relationships** in the portfolio. The Filtering Metro & Global 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 & Global 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.
5. Click **OK** to stop the copy process.

Modifying Metro Mirror and Global Mirror relationships

You can use the SAN Volume Controller Console to modify the attributes for Metro Mirror and Global Mirror relationships.

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

Perform the following steps to modify the attributes for Metro Mirror and Global Mirror relationships:

1. Click **Manage Copy Services** → **Metro & Global Mirror Relationships** in the portfolio. The Filtering Metro & Global 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 & Global 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 & Global Mirror Relationship panel is displayed.
You can change the following attributes from this panel:
 - The relationship name
 - The consistency group that contains this relationship

Switching the copy direction of a Metro Mirror or Global Mirror relationship

You can use the SAN Volume Controller Console to reverse the roles of the primary and secondary virtual disks (VDisks) in a Metro Mirror or Global Mirror relationship.

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

Perform the following steps to reverse the roles of the primary and secondary VDisks:

1. Click **Manage Copy Services** → **Metro & Global Mirror Relationships** in the portfolio. The Filtering Metro & Global 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 & Global Mirror Relationships panel is displayed.
3. Select **Switch Copy Direction** from the task list and click **Go**. The Switching the Direction of Mirror Relationship panel is displayed.

Deleting Metro Mirror or Global Mirror relationships

You can use SAN Volume Controller Console to delete a Metro Mirror or Global Mirror relationship.

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

Perform the following steps to delete a Metro Mirror or Global Mirror relationship:

1. Click **Manage Copy Services** → **Metro & Global Mirror Relationships** in the portfolio. The Filtering Metro & Global 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 & Global 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 Mirror Relationship panel is displayed.
5. Click **OK** to delete the relationship.

Creating Metro Mirror or Global Mirror consistency groups

You can create Metro Mirror or Global Mirror consistency groups using the wizard.

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

Perform the following steps to create a Metro Mirror or Global Mirror consistency group:

1. Click **Manage Copy Services** → **Metro & Global Mirror Consistency Groups** in the portfolio. The Filtering Metro & Global 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 or Global Mirror consistency groups

You can use the Filtering 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 Console.

Perform the following steps to specify filter criteria:

1. Click **Manage Copy Services** → **Metro & Global Mirror Consistency Groups** in the portfolio. The Filtering 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 or Global Mirror consistency group

You can use the SAN Volume Controller Console to rename a Metro Mirror or Global Mirror consistency group.

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

Perform the following steps to rename a Metro Mirror or Global Mirror consistency group:

1. Click **Manage Copy Services** → **Metro & Global Mirror Consistency Groups** in the portfolio. The Filtering 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 consistency group that you want to change.
4. Select **Rename a Consistency Group** from the task list and click **Go**. The Renaming Mirror Consistency Group panel is displayed.
5. Type a new name for the consistency group in the **New Name** field.
6. Click **OK**.

Starting a Metro Mirror or Global Mirror consistency group copy

You can use the SAN Volume Controller Console to start a Metro Mirror or Global Mirror copy process.

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

Perform the following steps to start a Metro Mirror or Global Mirror copy process:

1. Click **Manage Copy Services** → **Metro & Global Mirror Consistency Groups** in the portfolio. The Filtering 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 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 or Global Mirror consistency group copy process

You can use the SAN Volume Controller Console to stop a Metro Mirror or Global Mirror copy process.

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

Perform the following steps to stop a Metro Mirror or Global Mirror copy process:

1. Click **Manage Copy Services** → **Metro & Global Mirror Consistency Groups** in the portfolio. The Filtering 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 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.
5. Follow the directions that are displayed on this panel.

Deleting Metro Mirror and Global Mirror consistency groups

You can use the SAN Volume Controller Console to delete Metro Mirror and Global Mirror consistency groups.

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

Perform the following steps to delete a Metro Mirror or Global Mirror consistency group:

1. Click **Manage Copy Services** → **Metro & Global Mirror Consistency Groups** in the portfolio. The Filtering 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 Mirror Consistency Group panel is displayed.
5. Click **OK** to delete the consistency group.

Creating Metro Mirror and Global Mirror partnerships

You can use the SAN Volume Controller Console to create Metro Mirror and Global Mirror partnerships.

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

Perform the following steps to create a Metro Mirror or Global Mirror partnership:

1. Click **Manage Copy Services** → **Metro & Global Mirror Cluster Partnership** in the portfolio. The Mirror Cluster Partnership panel is displayed.
2. Click **Create**. The Create Cluster Partnerships panel is displayed.
3. Follow the instructions that are displayed on this panel to create the cluster partnership.

Modifying Metro Mirror or Global Mirror partnerships

You can change the bandwidth that is available for background copies from the Modify Cluster Partnership panel.

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

Perform the following steps to change a Metro Mirror or Global Mirror partnership:

1. Click **Manage Copy Services** → **Metro & Global Mirror Cluster Partnerships** in the portfolio. The Mirror Cluster Partnership panel is displayed.
2. Click **Modify**. The Modify Cluster Partnership panel is displayed.
3. Type the new rate for the background copy.

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.

4. Click **OK**.

Deleting Metro Mirror or Global Mirror partnerships

You can use the SAN Volume Controller Console to delete a Metro Mirror or Global Mirror partnership on the local cluster.

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

The Metro Mirror or Global 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 or Global Mirror partnership on the local cluster:

1. Click **Manage Copy Services** → **Metro & Global Mirror Cluster Partnerships** in the portfolio. The 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 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 Console.

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 license settings

You can use the SAN Volume Controller Console to view and update your license settings.

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

Perform the following steps to view and update the feature settings:

1. Click **Service and Maintenance** → **Set Features** in the portfolio. The License Settings panel is displayed.
2. Disable or Enable your licensed options.
3. Click **Set License 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 Console.

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 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.

Modifying error notification settings

You can use the SAN Volume Controller Console to change error notification settings for the cluster.

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

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 IBM 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:

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

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 Console.

Note: Log files that are copied to the configuration node are *not* automatically deleted by the SAN Volume Controller.

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 Console.

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, you must follow your multipathing device driver's procedure to discover the new paths and to check that each device identifier is now presenting the correct number of paths. If you are using the subsystem device driver (SDD), the device identifiers are referred to as virtual paths (vpaths). See the *IBM System Storage Multipath Subsystem Device Driver: User's Guide* or the documentation that is provided with your multipathing device driver for more information.
 - 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 the IBM Support Center 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. Use the PuTTY secure copy function to copy these public keys from the master console to the cluster.
4. Repeat 3 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 SSH key pair

You can use the SAN Volume Controller Console to replace the Secure Shell (SSH) key pair.

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.

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\svconconsole\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 public SSH key for a SAN Volume Controller cluster

There are times when you must replace the SSH public key used by the SAN Volume Controller cluster. For example, if you change the SSH keys that are used by the master console to communicate with the SAN Volume Controller Console or if you change the IP address of your SAN Volume Controller cluster, you must replace the cluster's SSH public key.

Perform the following steps to replace the public key used by the cluster:

1. Start the SAN Volume Controller Console by clicking on the desktop icon or by using your Web browser to go to `http://IP_address:9080/ica`, where *IP_address* is the IP address of the master console. The Signon window is displayed. This might take a few moments to open.

2. Enter the user ID superuser and the password passw0rd. The Welcome window is displayed.
3. Click **Clusters** from the portfolio.
4. Check the **Select** box for the cluster for which you wish to replace the key.
5. Click **Maintain SSH Keys** from the task list and click **Go**. The SSH Public Key Maintenance panel is displayed.
6. Type your user name and password.
7. 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.
8. Add the SSH client key by performing one of the following actions:
 - 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.
9. Click **Administrator**.
10. Type a name of your choice in the **ID** field that uniquely identifies the key to the cluster.
11. Click **Add Key**.
12. Click **Maintain SSH Keys**.
13. 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.

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 8. 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
 - Asynchronous Global Mirror and Global Mirror consistency groups

Preparing the SSH client system for the CLI

Before you can issue command-line interface (CLI) commands from the host to the cluster, you must prepare the Secure Shell (SSH) client system.

Microsoft Windows operating systems

The master console for the SAN Volume Controller includes the PuTTY client program, which is a Windows SSH client program. The PuTTY client program can be installed on your master console system in one of the following ways:

- If you purchased the master console hardware option from IBM, the PuTTY client program has been preinstalled on the hardware.

- You can use the master console software installation CD to install the PuTTY client program. The master console hardware option and the software-only master console both provide this CD.
- You can use the separate PuTTY client program installation wizard, **putty-<version>-installer.exe**, which is included as part of the SAN Volume Controller Console zip file. You can download the SAN Volume Controller Console zip file from the following Web site:
<http://www.ibm.com/storage/support/2145>

Note: Before you install the PuTTY client program, ensure that your Windows system meets the system requirements. See the *IBM System Storage SAN Volume Controller: Planning Guide* for system requirements.

If you want to use an SSH client other than the PuTTY client, the following Web site offers SSH client alternatives for Windows:

<http://www.openssh.org/windows.html>

AIX operating systems

For AIX 5L™ 5.1 and 5.2 on Power architecture, you can obtain the OpenSSH client from the Bonus Packs, but you also must obtain its prerequisite, OpenSSL, from the AIX toolbox for Linux applications for Power Systems. For AIX 4.3.3, you can obtain the software from the AIX toolbox for Linux applications.

You can also obtain the AIX installation images from IBM DeveloperWorks at the following Web site:

<http://oss.software.ibm.com/developerworks/projects/openssh>

Linux operating systems

The OpenSSH client is installed by default on most Linux distributions. If it is not installed on your system, consult your Linux installation documentation or visit the following Web site:

<http://www.openssh.org/portable.html>

The OpenSSH client can run on a variety of additional operating systems. For more information about the openSSH client, visit the following Web site:

<http://www.openssh.org/portable.html>

Preparing the SSH client system to issue CLI commands

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.

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 that points to the SAN Volume Controller Console).
 - c. Configure the PuTTY session for the CLI.
2. For other types of hosts:
 - a. Follow the instructions that are specific to the SSH client to generate an SSH key pair.
 - b. Store the SSH clients public key on the cluster (using a Web browser to point 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@svcconsoleip
```

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 **svcinfolssshkeys** 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>svcinfolssshkeys -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 svcinfolssshkeys
-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
```

Starting a PuTTY session for the CLI

You must start a PuTTY session to connect to the command-line interface (CLI).

This task assumes that you have already configured and saved a PuTTY session using the Secure Shell (SSH) key pair that you created for the CLI.

Perform the following steps to start a PuTTY session:

1. Select **Start** → **Programs** → **PuTTY** → **PuTTY**. The PuTTY Configuration window opens.
2. Select the name of your saved PuTTY session and click **Load**.
3. Click **Open**.

Note: If this is the first time that the PuTTY application is being used since you generated and uploaded the SSH key pair, a PuTTY Security Alert window is displayed. Click **Yes** to accept the change and trust the new key.

4. Type `admin` in the **login as:** field and press Enter.

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 **svcinfolstimezones** 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 **svcinfolicense** 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 step to display cluster properties:

Issue the **svcinfolcluster** command to display the properties for a cluster.

The following is an example of the command you can issue:

```
svcinfolcluster -delim : ITSOSVC42A
```

Where *ITSOSVC42A* is the name of the cluster.

```
IBM_2145:ITSOSVC42A:admin>svcinfolcluster -delim : ITSOSVC42A
id:0000020060806FB8
name:ITSOSVC42A
location:local
partnership:
bandwidth:
cluster_IP_address:9.34.91.112
cluster_server_IP_address:9.34.91.113
total_mdisk_capacity:147.5GB
space_in_mdisk_grps:136.0GB
space_allocated_to_vdisks:101.3GB
total_free_space:42.1GB
statistics_status:off
statistics_frequency:15
required_memory:8192
cluster_locale:en_US
SNMP_setting:all
SNMP_community:public
SNMP_server_IP_address:9.1.1.1
subnet_mask:255.255.252.0
default_gateway:9.43.85.1
time_zone:520 US/Pacific
email_setting:none
email_id:
code_level:4.2.0.0
FC_port_speed:2GB
console_IP:9..43.86.83:9080
id_alias:000020060806FB8
gm_link_tolerance:200
gm_inter_cluster_delay_simulation:40
gm_inter_cluster_delay_simulation:20
```

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.

Attention:

1. 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 the IBM Support Center to add the node back into the cluster without corrupting the data.
2. 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 ensure that the LUNs are the same before you add the new node to the cluster.
3. LUN masking for each LUN must be identical on all nodes in a cluster. You must ensure that the LUN masking for each LUN is identical before you add the new node to the cluster.
4. You must ensure that the model type of the new node is supported by the SAN Volume Controller software level that is currently installed on the cluster. If the model type is not supported by the SAN Volume Controller software level, upgrade the cluster to a software level that supports the model type of the new node. See the following Web site for the latest supported software levels:
<http://www.ibm.com/storage/support/2145>

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 node name (WWNN) 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 WWNN 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:hardware  
1:node1:10L3ASH:500507680100002C:online:0:io_grp0:yes:202378101C0D18D8:other  
....
```

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:hardware  
5005076801000001:000341:10L3ASH:202378101C0D18D8:other  
5005076801000009:000237:10L3ANF:202378101C0D1796:other  
50050768010000F4:001245:10L3ANF:202378101C0D1796:other  
....
```

3. Issue the **svctask addnode** CLI command to add a node to the cluster.

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. If you did not specify a name when you issued the **svctask addnode** command, 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:IO_group_id:
IO_group_name:config_node:UPS_unique_id:hardware
1:group1node1:10L3ASH:500507680100002C:online:0:io_grp0:yes:202378101C0D18D8:other
2:group1node2:10L3ANF:5005076801000009:online:0:io_grp0:no:202378101C0D1796:other
3:group2node1:10L3ASH:5005076801000001:online:1:io_grp1:no:202378101C0D18D8:other
4:group2node2:10L3ANF:50050768010000F4:online:1:io_grp1:no:202378101C0D1796:other
....
```

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.

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 **svcinfo lsnode** 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:

```
svcinfo lsnode -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:hardware  
1:group1node1:10L3ASH:500507680100002C:online:0:io_grp0:yes:202378101C0D18D8:8G4  
2:group1node2:10L3ANF:5005076801000009:online:0:io_grp0:no:202378101C0D1796:8G4  
3:group2node1:10L3ASH:5005076801000001:online:1:io_grp1:no:202378101C0D18D8:8G4  
4:group2node2:10L3ANF:50050768010000F4:online:1:io_grp1:no:202378101C0D1796:8G4
```

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_speed:2GB  
port_id:500507680120002C  
port_status:active  
port_speed:2GB  
port_id:500507680130002C  
port_status:active  
port_speed:2GB  
port_id:500507680140003C  
port_status:active  
port_speed:2GB  
hardware:8G4
```

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).

Note: When you are adding MDisks to an MDisk group using the command `svctask addmdisk` or when you are creating an MDisk group using the command `svctask mkmdiskgrp -mdisk`, the SAN Volume Controller performs tests on the MDisks in the list before the MDisks are allowed to become part of an MDisk group. These tests include checks of the MDisk identity, capacity, status, and the ability to perform both read and write operations. If these tests fail, or exceed the time allowed, the MDisks are not added to the group. However, with the command `svctask mkmdiskgrp -mdisk`, the MDisk group is still created even if the tests fail, but it does not contain any MDisks. If tests fail, confirm that the MDisks are in the correct state and that they have been correctly discovered.

The following reasons contribute to a typical MDisk test failure:

- The MDisk is not visible to all SAN Volume Controller nodes in the cluster.
- The MDisk identity has changed from a previous discovery operation.
- The MDisk cannot perform read or write operations.
- The status of the MDisk is degraded, excluded, or offline.
- The Mdisk does not exist.

The following reasons contribute to a typical MDisk test timeout:

- The disk controller subsystem on which the MDisk resides is failing.
- A SAN fabric or cable fault condition exists that is preventing reliable communication with the MDisk.

Perform the following steps to discover and then view a list of 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.

Notes:

- a. Only issue the `svctask detectmdisk` command when you are sure that all of the disk controller ports are working and correctly configured in the controller and the SAN zoning. Failure to do this can result in errors that are not reported.

- b. Although it might appear that the **detectmdisk** command has completed, extra time might be required for it to run. The **detectmdisk** is asynchronous and returns a prompt while the command continues to run in the background. You can use the **lsdiscoverystatus** command to list the discovery status.
2. When the detection is complete, issue the **svcinfolismdiskcandidate** CLI command to show the unmanaged MDisks. These MDisks have not been assigned to an MDisk group.
3. Issue the **svcinfolismdisk** CLI command to view all of the MDisks.

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.

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 therefore 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
128 MB	512 TB
256 MB	1 PB
512 MB	2 PB
2048 MB	8 PB

Important: You can specify different extent sizes for different MDisk groups; however, you cannot migrate VDIs 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 **svcinflsmdiskgrp** 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:

```
svcinflsmdiskgrp -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 VDIs

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 can establish an empty MDisk group to hold the MDisks used for image mode VDIs.

Note: If you want to keep the data on an MDisk, you should instead create image mode VDIs. This task describes how to create a VDisk with striped virtualization.

Perform the following steps to create VDIs:

1. Issue the **svcinflsmdiskgrp** 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:

```
svcinflsmdiskgrp -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 **svcinflsiogrp** CLI command to show the I/O groups and the number of VDIs 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 VDIs in different I/O groups. You can use FlashCopy to make copies of VDIs regardless of whether the source and target VDisk are in the same I/O group. If you plan to use intracluster Metro Mirror or Global Mirror, both the master and auxiliary VDisk must be in the same I/O group.

The following is an example of the CLI command you can issue to list I/O groups:

```
svcinflsiogrp -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
```

- 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 VDIsks between the I/O groups so that the I/O workload is shared evenly between all SAN Volume Controller nodes.

- 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.

- Issue the **svcinfolsvdisk** CLI command to list all the VDIsks that have been created.

The following is an example of the CLI command you can issue to list VDIsks:

```
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:

- 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 addhostport -hbawwpn 210100e08b251dd5 demohost1`

This command adds another HBA WWPN called *210100e08b251dd5* to the host that was created in step 1 on page 182.

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.

Modifying the amount of memory available for Copy Service features using the CLI

You can use the command-line interface (CLI) to modify the amount of memory that is available for the FlashCopy, Metro Mirror, or Global Mirror Copy Services.

The following table provides an example of the amount of memory that is required for each Copy Service feature:

Copy Service feature	Grain size	1 MB of memory provides the following VDisk capacity for the specified I/O group
Metro Mirror or Global Mirror	256 KB	2 TB of total Metro Mirror and Global Mirror VDisk capacity
FlashCopy	256 KB	2 TB of total FlashCopy source VDisk capacity
FlashCopy	64 KB	512 GB of total FlashCopy source VDisk capacity
Incremental FlashCopy	256 KB	1 TB of total incremental FlashCopy source VDisk capacity
Incremental FlashCopy	64 KB	256 GB of total incremental FlashCopy source VDisk capacity

Notes:

1. For multiple FlashCopy targets, you must consider the number of mappings. For example, for a mapping with a grain size of 256 KB, 8 KB of memory allows one mapping between a 16 GB source VDisk and a 16 GB target VDisk. Alternatively, for a mapping with a 256 KB grain size, 8 KB of memory allows two mappings between one 8 GB source VDisk and two 8 GB target VDIs.
2. When creating a FlashCopy mapping, if you specify an I/O group other than the I/O group of the source VDisk, the memory accounting goes towards the specified I/O group, not towards the I/O group of the source VDisk.

Perform the following steps to modify and verify the amount of memory that is available:

1. Issue the following command to modify the amount of memory that is available for a Copy Service feature:
`svctask chiogrp -feature flash|remote -size memory_size io_group_id | io_group_name`
 where *flash|remote* is the type of Copy Service you want to modify, *memory_size* is the amount of memory you want to be available, and *io_group_id* | *io_group_name* is the ID or name of the I/O group for which you want to modify the amount of available memory.
2. Issue the following command to verify that the amount of memory has been modified:
`svcinfolsiogrp object_id | object_name`
 where *object_id* | *object_name* is the ID or name of the I/O group for which you have modified the amount of available memory.

The following information is an example of the output that is displayed.

```
id 0
name io_grp0
node_count 2
vdisk_count 236
host_count 4
flash_copy_total_memory 20.0MB
flash_copy_free_memory 15.1MB
remote_copy_total_memory 30.5MB
remote_copy_free_memory 2.0MB
```

Creating FlashCopy mappings using the CLI

You can use the command-line interface (CLI) to create FlashCopy mappings.

A FlashCopy mapping specifies the source and target virtual disk (VDisk). Source VDIs and target VDIs must meet the following requirements:

- They must be the same size.
- They must be managed by the same cluster.

A VDisk can be the source in up to 16 mappings. A mapping is started at the point in time when the copy is required.

Perform the following steps to create FlashCopy mappings:

1. The source and target VDisk must be the exact same size. Issue the **svcinfolsvdisk -bytes** CLI command to find the size (capacity) of the VDisk in bytes.
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 VDisk that you want to make the target VDisk, *main1copy* is the name that you want to call the FlashCopy mapping and 75 is the priority that you want to give the 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 VDisk that you want to make the target VDisk, *main2copy* is the name that you want to call the FlashCopy mapping.

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:clean_progress:incremental
0:main1copy:77:vdisk77:78:vdisk78:::idle_or_copied:0:75:100:off
1:main2copy:79:vdisk79:80:vdisk80:::idle_or_copied:0:50:100:off
```

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 start 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
```

Note: For any group that has just been created the status reported is empty

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 Flash Copy 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:status: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 **svcinfolsfconsistgrp** 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:

```
svcinfolsfconsistgrp -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_or_copied  
FC_mapping_id:0  
FC_mapping_name:main1copy  
FC_mapping_id:1  
FC_mapping_name:main2copy
```

Preparing and starting a FlashCopy mapping using the CLI

Before you start the FlashCopy process using the command-line interface (CLI), you must prepare and start a FlashCopy mapping.

Starting 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 start 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 **svcinfolsfcmmap** CLI command to check the state of the mapping.

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:status:progress:copy_rate  
0:main1copy:0:mainvdisk1:1:bkpvdisk1:::prepared:0:50
```

3. Issue the **svctask startfcmmap** CLI command to start the FlashCopy mapping.

The following is an example of the CLI command you can issue to start the FlashCopy mapping:

```
svctask startfcmmap main1copy
```

Where *main1copy* is the name of the FlashCopy mapping.

4. Issue the **svcinfolsfcmmapprogress** CLI command to check the progress of the FlashCopy mapping.

The following is an example of the output that is displayed:

```
svcinfolsfcmmapprogress -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 starting a FlashCopy consistency group using the CLI

You can use the command-line interface (CLI) to prepare and start 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 start command for the whole group to prepare or start all the mappings at once.

Perform the following steps to prepare and start a FlashCopy consistency group:

1. Issue the **svctask prestartfcconsistgrp** CLI command to prepare the FlashCopy consistency group before the copy process can be started.

Remember: You only have to issue a single prepare command for the whole group to prepare all the mappings simultaneously.

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 **svcinfolsfconsg** 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:

```
svcinfo lsfcconsistgrp -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 the FlashCopy consistency group to make the copy.

Remember: You only have to issue a single start command for the whole group to start all the mappings simultaneously.

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 **svcinfo lsfcconsistgrp** 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:

```
svcinfo lsfcconsistgrp -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 **svcinfo lsnode** 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 **svcinfo lsnode** 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:

```
svcinfo lsnode 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:

```
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).

Select one or more of the following options to determine the relationship between VDIs and MDIs:

- To display a list of the IDs that correspond to the MDIs that comprise the VDisk, issue the following CLI command:

```
svcinfolsvdiskmember vdiskname/id
```

where *vdiskname/id* is the name or ID of the VDisk.

- To display a list of IDs that correspond to the VDIs that are using this MDI, issue the following CLI command:

```
svcinfolsmdiskmember mdiskname/id
```

where *mdiskname/id* is the name or ID of the MDI.

- To display a table of VDisk IDs and the corresponding number of extents that are being used by each VDisk, issue the following CLI command:

```
svcinfolsmdiskextent mdiskname/id
```

where *mdiskname/id* is the name or ID of the MDI.

- To display a table of MDI IDs and the corresponding number of extents that each MDI provides as storage for the given VDisk, issue the following CLI command:

```
svcinfolsvdiskextent vdiskname/id
```

where *vdiskname/id* is the name or ID of the VDisk.

Determining the relationship between MDIs and RAID arrays or LUNs using the CLI

You can determine the relationship between managed disks (MDIs) and RAID arrays or LUNs using the command-line interface (CLI).

Each MDI 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 MDIs and RAID arrays or partitions.

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

1. Issue the following command to display a detailed view of the MDI:

```
svcinfolsmdisk mdiskname
```

Where *mdiskname* is the name of the MDI 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:

```
svcinfolscontroller 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 on page 190. 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 the IBM Support Center 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:

```
svcinfolsnode
```

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 VDIs as required.

Attention: This is a disruptive procedure. Access to the VDisk is lost while you follow this procedure. Under no circumstances should VDIs 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 VDIs.

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 System Storage 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 System Storage 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.

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> <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 **svcinfolnode** 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 **svcinfolnode** 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:

```
svcinfolnodevpd nodename
```

Where *nodename* is the name that you recorded in step 1.

- 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 System Storage 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 194.

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 194, 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 WWPNNs
 - 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 System Storage 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 System Storage 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 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.

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:
`svcinfolnode`
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:
`svcinfolnodecandidate`

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 System Storage 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 the IBM Support Center 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:

```
svcinfolnode
```

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:

```
svcinfolsvdisk -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 VDisk.

2. Issue the following command to verify that the VDIsks are now online:

```
svcinfolsvdisk -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 System Storage 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	-
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:

```
svcinfo lsvdisk -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 VDIsks are being heavily utilized as you can spread the data they contain more evenly across all the MDisks in the group using the procedure below.
5. Stop the statistics collection again by issuing the following command:

```
svctask stopstats
```

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 on page 203), does not already own a large set of extents on the target MDisk.

4. For each set of extents, issue the following command to move them to another MDisk:

```
svctask migrateextents -source <mdiskname/ID> -exts  
<num_extents_from_step1> -target <newmdiskname/ID>  
-threads 4 <vdiskid_returned_from_step1>
```

where *<num_extents_from_step1>* is the number of extents on the *<vdiskid_returned_from_step1>*, that is, the data that is returned from the command issued in step 1 on page 203. *<newmdiskname/ID>* is the name or ID of the MDisk to which you want to migrate this set of extents.

5. Repeat steps 2 on page 203 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:

```
svcinfolsmigrate
```

Migrating VDIs between MDisk groups using the CLI

You can migrate virtual disks (VDIs) between managed disk (MDisk) groups using the command-line interface (CLI).

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

When a migrate command is issued, a check is made to ensure that the destination of the migrate has enough free extents to satisfy the command. If it does, the command proceeds, but will take some time to complete.

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 MDIs 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:

```
svcinfl 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/iosstats/v_*
```

This will copy all the VDisk statistics files to the AIX host in the current directory.

4. 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.
5. Stop the statistics collection again. Issue the following command:

```
svctask stopstats
```

Once you have analyzed the I/O statistics data, you can determine which 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.

6. After having determined which VDisk you wish to migrate, and the new MDisk group you wish to migrate it to, issue the following command:

```
svctask migratevdisk -vdisk <vdiskname/ID> -mdiskgrp  
<newmdiskgrname/ID> -threads 4
```

7. You can check the progress of the migration by issuing the following command:

```
svcinfl 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:

1. Synchronize all file systems that are mounted on the given VDisk.
2. Stop all I/O operations to the VDisk.
3. 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.

4. Resynchronize the VDisk to host mapping. See the *IBM System Storage Multipath Subsystem Device Driver: User's Guide* or the documentation that is provided with your multipathing driver for more information.
5. 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 use 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/storage/support/2145>

For complete instructions on the CLI commands, see the *IBM System Storage 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. 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.
4. Issue the **svcinfolsmdisk** 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.

5. Convert the unmanaged-mode MDisk to an image mode virtual disk. Issue the **svctask mkvdisk** command to create an image mode virtual disk object.
6. 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 System Storage 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 **svcinfolcluster** 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 **svcinfolcluster** 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 **svcinfolcluster** 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 **svcinfolsshkeys** 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 SNMP 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. For example, you can issue one of the following CLI commands to set up error notification:

```
svctask setevent -snmptrap all or hardware_only
-snmppip 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.

Setting up e-mail error notifications using the CLI

You can use the command-line interface (CLI) to set up your system to enable e-mail and inventory event notification.

You must have the IP address of the server that receives the notifications and you must ensure that the server is up and running.

See the *IBM System Storage SAN Volume Controller: Command-Line Interface User's Guide* for information about the parameters that are used for these commands.

Perform the following steps to set up e-mail error notification:

1. Issue the **svctask setemail** command to enable your system to use the e-mail notification function.
2. Issue the **svctask mkemailuser** command to add recipients. You can have up to twelve recipients of the notifications.
3. Issue the **svctask lsemailuser** to generate a report that allows you to view the e-mail notification settings for recipients.
 - If you want to edit settings for a recipient, issue the **svctask chemailuser** command.
 - If you want to delete a recipient, issue the **svctask rmemailuser** command.
 - If the settings are correct, proceed to the next step.
4. Issue the **svctask startemail** command to activate the e-mail and inventory notification function.
5. Optionally, issue the **svctask testemail** command to send an e-mail notification to one or more recipients to ensure that the function is operating correctly.
6. Optionally, issue the **svctask sendinventoryemail** command to send an inventory e-mail notification to all recipients that are enabled to receive inventory e-mail notifications.

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**
- **svcinfolerrlogbydisk**

- **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, Global 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.

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 9. 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
- Global Mirror relationships
- Global 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 Console.

The backup function is designed to back up information about your cluster configuration, such as virtual disks (VDisks), local Metro Mirror information, local Global 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 **svconfig** CLI command is designed to back up information about your cluster configuration, such as virtual disks (VDisks), local Metro Mirror information, local Global 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:

```
svconfig clear -all
```

5. Issue the following CLI command to backup your cluster configuration:

```
svconfig 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 **svconfig 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.

- Issue the following command to exit the cluster:

```
exit
```

- 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.

- 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 Console.

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 unique identifier (UID) of the VDisks that are within the I/O groups can change. This can affect IBM TotalStorage Productivity Center 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.
 - If you need assistance, contact the IBM Support Center.
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.

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 Console.

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 10. 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. Only the following commands can be issued during the software upgrade:

- All svcinfo commands
- svctask rmnode

Note: Applying a software upgrade takes a varying length of time. Plan for at least 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 and Global 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 or Global Mirror partnerships between two clusters with different software levels.

Installing or upgrading the SAN Volume Controller software

The SAN Volume Controller software can be installed or upgraded after you download 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.

- All `svcinfo` commands
- `svctask rmnode`

To determine when your software upgrade process has completed, you will be notified through the SAN Volume Controller Console or, if you are using the command-line interface, examine the error log.

Because of the operational limitations that occur during the software upgrade process, the software installation is a user task.

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 the `pscp` application, 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 the `pscp` application:

1. Start a PuTTY session.
2. Configure your PuTTY session to access your SAN Volume Controller 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 to provide sufficient space:

- Issue the **svctask cleardumps** CLI command to free space on the cluster and repeat step 6 on page 226.
- 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/elogs/* 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 the error logs.

After you have transferred the error logs to the master console, repeat step 6 on page 226.

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.

If the new node requires a level of software that is higher than the software level that is available on the cluster, the entire cluster must be upgraded before the new node can be added to the cluster.

Error counts

During the software upgrade if you are using IBM Subsystem device driver (SDD) as the multipathing software on the host, 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 during a software upgrade. See the *IBM System Storage 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 fails. 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.

Upgrading the SAN Volume Controller software using the SAN Volume Controller Console

You can upgrade the cluster software using the SAN Volume Controller Console.

Attention: Before you start a software upgrade, you must check for offline or degraded VDisks. An offline VDisk can cause write data that has been modified to be pinned in the SAN Volume Controller cache. This prevents VDisk failover and causes a loss of I/O access during the software upgrade. If the `fast_write_state` is empty, a VDisk can be offline and not cause errors during the software upgrade.

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 Web site:
<http://www.ibm.com/storage/support/2145>
 - 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.
 6. Click **Upload**. The Software upgrade - file upload panel is displayed.
 7. Click **Browse** and select the SAN Volume Controller software file that you downloaded in step 1.
 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.
 - 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.

Upgrading the SAN Volume Controller software using the CLI

You can use the command-line interface (CLI) to install software upgrades.

Attention: Before you start a software upgrade, you must check for offline or degraded VDisks. An offline VDisk can cause write data that has been modified to be pinned in the SAN Volume Controller cache. This prevents VDisk failover and causes a loss of I/O access during the software upgrade. If the `fast_write_state` is empty, a VDisk can be offline and not cause errors during the software upgrade.

Perform the following steps to upgrade the software:

1. Download the SAN Volume Controller code from the following Web site:
<http://www.ibm.com/storage/support/2145>
 - 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 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.
- 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 **svcinfo lsnodevpd** 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.

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 Controller from issuing I/O to backend controllers and flushes data from the SAN Volume Controller nodes 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 Controller nodes.
8. Remount file systems and start host applications.

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 use the node rescue procedure to reinstall the SAN Volume Controller software.

Similarly, if you have replaced the service controller, you should use the node rescue procedure to ensure that the service controller has the correct software.

Attention: If you recently replaced both the service controller and the disk drive as part of the same repair operation, node rescue fails.

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 node cannot boot and the hardware boot indicator remains on the front panel display or the boot operation does not progress. If this occurs, use the node rescue procedure to reinstall the SAN Volume Controller software.

Node rescue works by booting the operating system from the service controller and running a program that copies all the SAN Volume Controller software from any other node that can be found on the fibre-channel fabric.

Attention: When running node rescue operations, only run one node rescue operation on the same SAN, at any one time. Wait for one node rescue operation to complete before starting another.

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 node is connected to the fibre-channel fabric.
3. Ensure that the SAN zoning allows a connection between at least one port of this node and one port of another node. It is better if multiple ports can connect. This is particularly important if the zoning is by worldwide port name (WWPN) and you are using a new service controller. In this case, you might need to use SAN monitoring tools to determine the WWPNs of the node. If you need to change the zoning, remember to set it back when the service procedure is complete.
4. Turn off the node.
5. Press and hold the left and right buttons on the front panel.
6. Press the power button.
7. Continue to hold the left and right buttons until the node-rescue-request symbol is displayed on the front panel (Figure 20).



Figure 20. Node-rescue-request display

The node rescue request symbol displays on the front panel display until the node 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 can take up to four hours for an eight-node cluster.

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. If a node fails to restart during this downgrade process, the process is suspended. The following scenarios can cause the downgrade process to suspend:

- A node (other than the node that is currently upgrading) is offline, restarted or asserted

- A node fails to update to the new software level
- A node is deleted while it is in the process of updating

You must check the error log to determine the reason for the failure before you attempt to upgrade the cluster again.

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.

Chapter 11. Configuring and servicing storage subsystems

To avoid performance issues, you must ensure that your storage subsystems and switches are correctly configured to work with the SAN Volume Controller.

Virtualization provides many benefits over direct-attached or direct SAN-attached. 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.

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 for storage subsystems

You must follow the guidelines and procedures for your storage subsystem to maximize performance and to avoid potential I/O problems.

General guidelines

- Avoid splitting arrays into multiple logical disks at the storage subsystem level. Where possible, create a single logical disk from the entire capacity of the array.
- Depending on the redundancy that is required, create RAID-5 arrays 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 MDisks 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 MDisks 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. To optimize the benefits of virtualization, migrate this data across the other MDisks in the group as soon as possible.

- Follow the FlashCopy requirements before you set up the storage. Balance the spread of the FlashCopy VDisks across the MDisk groups and then across the storage subsystems. The I/O characteristics of the application that is writing to the source VDisk also affects 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.

Logical disk configuration guidelines for storage subsystems

Most storage subsystems provide some mechanism to create multiple logical disks from a single array. This is useful when the storage subsystem presents storage directly to the hosts.

However, in a virtualized SAN, use a one-to-one mapping between arrays and logical disks so that 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 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 three extents. As a result, when the system reads and writes 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 guidelines for storage subsystems

With 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 can 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 a 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 mean that it takes a longer time 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 (a redundant disk) and more data is exposed if a second disk fails before the rebuild operation 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.

Notes:

1. For optimal performance, use arrays with between 6 and 8 member disks.
2. When creating RAID arrays with mirroring, the number of component disks in each array does not affect redundancy or performance.

Optimal MDisk group configuration guidelines for storage subsystems

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.

Notes:

1. The performance of an MDisk group is generally governed by the slowest MDisk in the group.
2. The reliability of an MDisk group is generally governed by the weakest MDisk in the group.
3. If a single MDisk in a group fails, access to the entire group is lost.

Use the following guidelines when you group 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 TotalStorage Enterprise Storage Server (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 IBM ESS arrays, even though they are online. The performance is limited to the performance of the JBOD in the IBM FAStT storage subsystem, therefore slowing down the IBM ESS arrays.

To fix this problem, you can create three groups. The first group must contain the IBM ESS arrays, MDisks 0 through 9, the second group must contain the RAID-1 array, and the third group must contain the large RAID-5 array.

FlashCopy mapping guidelines for storage subsystems

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 operations perform 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 operation 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 that your application performs during a FlashCopy operation. Ensure that you do not overload the storage. The calculations contain a heavy weighting when the FlashCopy feature 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 target 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.

Image mode VDisks and data migration guidelines for storage subsystems

Image mode virtual disks (VDisks) enable you to import and then migrate existing data that is managed by the SAN Volume Controller cluster.

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 they are connected through the cluster.

If the existing storage subsystems do not follow the configuration guidelines, consider stopping I/O operations on the host systems while you migrate 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 can occur.

The procedure for importing managed disks (MDisks) that contain existing data depends on the amount of free capacity you have in the cluster. You must 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 available capacity, the migration causes the MDisk group to have an uneven distribution of data because some MDisks are more heavily loaded than others. Further migration operations are required to ensure an even distribution of data and subsequent I/O loading.

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. Unmap 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 or use the SAN Volume Controller Console to discover MDisks. 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 steps 4a through 4c 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:
 - 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 that you created in step 2. 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 that you created in step 2 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.

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.

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 underused 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 task 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 or use the SAN Volume Controller Console to discover MDisks. 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 steps 4a through 4c 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 cluster.
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 that you created in step 1. 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 that you selected in step 1:
 - a. Remove the MDisks from the temporary MDisk group.
 - b. At the storage subsystem, unmap the logical disks that have been migrated from the SAN Volume Controller cluster and delete them 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 or use the SAN Volume Controller Console to discover MDisks. 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 that you selected in step 1 on page 240.
8. Repeat steps 3 on page 240 through 7 on page 240 for the next array.

Configuring a balanced storage subsystem

The attachment of a storage subsystem to a SAN Volume Controller requires that specific settings are applied to the device.

There are two 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 storage connections that allow the SAN Volume Controller to access the logical units

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 to the applications.

Scenario: You have 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 per second
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 241, 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 mappings. 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 operations generate because it reduces the rate at which I/O from host systems can be processed. When a FlashCopy mapping 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 the FlashCopy feature depends on the type of I/O workload that is generated by an application.

Table 14. Calculate the impact of FlashCopy mappings

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 mapping, 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 mapping is used to provide point-in-time backups. During the FlashCopy process, a host application generates an I/O workload of random read and write operations 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 target 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 mapping.
 - 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 mappings. The I/O rate per VDI 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 mappings that have source VDIs in the MDisk group. Both source VDIs are accessed by applications that issue random read and write operations. 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 242 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 that 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 that 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 you can 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 previous calculation is too simplistic to model your storage use after. 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 MDIs. 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 that 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 MDIs without having to take the storage offline.
- Stop unnecessary FlashCopy mappings to reduce the amount of I/O operations that are submitted to the backend storage. If you perform FlashCopy operations in parallel, consider reducing the amount of FlashCopy mappings 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 VDisk or limit the queue depth per HBA, or both. 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 that you have.

Discovering logical units

The SAN Volume Controller initialization includes a process called discovery.

The discovery process systematically recognizes all visible ports on the SAN for devices that identify themselves as storage subsystems and the number of logical units (LUs) that they export. The LUs can contain new storage or a new path for previously discovered storage. 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.
- Do not present new LUs to the SAN Volume Controller until the array initialization and format is complete. If you add a LUN to an MDisk group before the array initialization format is complete, the MDisk group goes offline. While the MDisk group is offline, you cannot access the VDIsks that are in the MDisk group.
- When you export an LU to the SAN Volume Controller, the LU *must* be accessible through all ports on the storage subsystem that are 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 rmmdisk** CLI command to remove the MDisk from the MDisk group.
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 **svcinfo 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 246:

```
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 and any SCSI reservation that the host has on the VDisk:

```
svctask rvmdiskhostmap -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 on the storage subsystem.

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 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 storage subsystem's 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.

Renaming a storage subsystem using the CLI

You can use the command-line interface (CLI) to rename a storage subsystem.

Perform the following step to rename a storage subsystem:

Issue the `svctask chcontroller -name new_name controller_id` command.

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 rmdisk -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
```

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.

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.

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.

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.

Removing a storage subsystem

You can replace or decommission a storage subsystem.

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

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 VDIs 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 VDIs that are using the MDisks.

4. Check the progress of the migration process by issuing the following command from the command-line interface (CLI): `svcinfo lsmigrate`
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. Perform the following steps to have the cluster rescan the fibre-channel network:
 - a. Click **Work with Managed Disks** → **Managed Disks**.
 - b. Select **Discover MDisks** from the task list and click **Go**. The Discovering Managed Disks panel is displayed. The rescan discovers that the MDisks have been removed from the cluster and also rebalances MDisk access across the available controller device ports.

8. Verify that there are no MDisks for the storage subsystem that you want to decommission.
9. Remove the storage subsystem from the SAN so that the SAN Volume Controller ports can no longer access the storage subsystem.

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 VDIs 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 VDIs 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:

```
svcinfolismigrate
```

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.

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.

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 configuration tools for the storage subsystem 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 Bull FDA subsystems

This section provides information about configuring Bull StoreWay FDA subsystems for attachment to a SAN Volume Controller.

Supported firmware levels for the Bull FDA

The Bull FDA subsystem must use a firmware level that is supported by the SAN Volume Controller.

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

Logical unit creation and deletion for Bull FDA

You can create or delete logical units for the Bull FDA. See the storage configuration guidelines that are specified in the Bull FDA documentation that is provided for this subsystem.

Platform type for Bull FDA

You must set all logical units that the SAN Volume Controller accesses to platform type AX (AIX).

Access control methods for Bull FDA

You can use access control to restrict access from hosts and SAN Volume Controller clusters. You do not need to use access control to allow a SAN Volume Controller cluster to use all of the defined logical units on the subsystem.

The following table lists the access control methods that are available:

Method	Description
Port Mode	Allows access to logical units that you want to define on a per storage controller port basis. SAN Volume Controller visibility (through switch zoning, physical cabling, etc.) must allow the SAN Volume Controller cluster to have the same access from all nodes and the accessible controller ports have been assigned the same set of logical units with the same logical unit number. This method of access control is not recommended for SAN Volume Controller connection.
WWN Mode	Allows access to logical units using the WWPN of each of the ports of an accessing host device. All WWPNs of all the SAN Volume Controller nodes in the same cluster must be added to the list of linked paths in the controller configuration. This becomes the list of host (SAN Volume Controller) ports for an LD Set or group of logical units. This method of access control allows sharing because different logical units can be accessed by other hosts.

Setting cache allocations for Bull FDA

Cache allocations can be set manually; however, changes to the default settings can adversely effect performance and cause you to lose access to the subsystem.

Snapshot Volume and Link Volume for Bull FDA

You cannot use Copy Services logical volumes with logical units that are assigned to the SAN Volume Controller.

Configuring the EMC CLARiiON subsystem

This section provides information about configuring the EMC CLARiiON storage system for attachment to a SAN Volume Controller.

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 EMC Navisphere GUI 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.

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 [worldwide port name (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

Option	EMC CLARiiON default setting	SAN Volume Controller required setting
Failover mode	0	2
Unit Serial Number	Array	Array

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 SPTM 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. Select **Array** in the drop down list of the **Unit Serial Number** field.
16. Assign a host name in the Host Name field.
17. Click **OK**.
18. 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.
19. 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
20. Repeat steps 1 to 19 to register the rest of your SAN Volume Controller WWPNs.

All your WWPNs are registered against the host name that you specified.

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.

Supported models of the EMC CLARiiON

The SAN Volume Controller supports models of the EMC CLARiiON.

See the following Web site for the latest supported models:

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

Supported firmware levels for the EMC CLARiiON

The EMC CLARiiON must use a firmware level that is supported by the SAN Volume Controller.

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

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

Concurrent maintenance on the EMC CLARiiON

Concurrent maintenance 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.

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.

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.

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.

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

Global settings for the EMC CLARiiON

Global settings apply across an EMC CLARiiON subsystem.

Table 17 lists the global settings that are supported by the SAN Volume Controller.

Table 17. 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

Controller settings for the EMC CLARiiON

The controller settings for the EMC CLARiiON are the settings that apply across one EMC CLARiiON subsystem.

Table 18 lists the options that can be set by the EMC CLARiiON.

Table 18. 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	Default recommended
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.

Port settings for the EMC CLARiiON

Port settings are configurable at the port level.

Table 19 lists port settings, the EMC CLARiiON defaults, and the required settings for the SAN Volume Controller.

Table 19. 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.

Logical unit settings for the EMC CLARiiON

Logical unit (LU) settings are configurable at the LU level.

Table 20 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 20. 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

Table 20. EMC CLARiiON LU settings supported by the SAN Volume Controller (continued)

Option	EMC CLARiiON default setting	SAN Volume Controller required setting
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.

Configuring the EMC Symmetrix and Symmetrix DMX subsystems

This section provides information about configuring the EMC Symmetrix and Symmetrix DMX for attachment to a SAN Volume Controller.

Supported models of the EMC Symmetrix and Symmetrix DMX controllers

The SAN Volume Controller supports models of the EMC Symmetrix and Symmetrix DMX controllers.

See the following Web site for the latest supported models:

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

Supported firmware levels for the EMC Symmetrix and Symmetrix DMX

The EMC Symmetrix and Symmetrix DMX must use a firmware level that is supported by the SAN Volume Controller.

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

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

Concurrent maintenance on the EMC Symmetrix and Symmetrix DMX

Concurrent maintenance is the capability to perform I/O operations to the EMC Symmetrix or Symmetrix DMX 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 and Symmetrix DMX are Enterprise class devices that support 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/Symmetrix DMX support concurrent upgrade of the EMC Symmetrix/Symmetrix DMX firmware.

User interfaces on EMC Symmetrix and Symmetrix DMX

Ensure that you are familiar with the user interface applications that support the EMC Symmetrix and Symmetrix DMX subsystems.

EMC Control Center

A basic EMC Symmetrix or Symmetrix DMX 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 and Symmetrix DMX subsystems.

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 the EMC Symmetrix and Symmetrix DMX.

Sharing the EMC Symmetrix or Symmetrix DMX between a host and a SAN Volume Controller

There are restrictions for sharing EMC Symmetrix and Symmetrix DMX subsystems between a host and a SAN Volume Controller.

An EMC Symmetrix or Symmetrix DMX can be shared between a host and a SAN Volume Controller under the following conditions:

- 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 or Symmetrix DMX because the multipathing drivers (for example, subsystem device driver (SDD) and PowerPath) cannot coexist.
- Other hosts can be directly connected to an EMC Symmetrix or Symmetrix DMX 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 or Symmetrix DMX must be configured such that other hosts cannot access the LUs that are managed by the SAN Volume Controller.

Switch zoning limitations for the EMC Symmetrix and Symmetrix DMX

There are limitations in switch zoning for the SAN Volume Controller and the EMC Symmetrix and Symmetrix DMX subsystems.

Switch zoning

The SAN Volume Controller switch zone must include at least one target port on two or more fibre-channel adapters to avoid a single point of failure.

The EMC Symmetrix and Symmetrix DMX 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 or Symmetrix DMX controller should be present in the fabric zone.

Connecting to the SAN

The EMC Symmetrix and Symmetrix DMX connect 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. The fibre-channel adapter provides 2 - 12 target ports. The EMC Symmetrix and Symmetrix DMX assign 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, perform the following steps:

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.

Quorum disks on EMC Symmetrix and Symmetrix DMX

The SAN Volume Controller chooses managed disks (MDisks) that are presented by the EMC Symmetrix or Symmetrix DMX as quorum disks.

The SAN Volume Controller uses a logical unit (LU) that is presented by an EMC Symmetrix or Symmetrix DMX as a quorum disk. The SAN Volume Controller provides a quorum disk even if the connection is through a single port.

Advanced functions for EMC Symmetrix and Symmetrix DMX

SAN Volume Controller cache-disabled virtual disks (VDisks) can be used as the source or target in Symmetrix advanced copy functions (for example, SRDF and TimeFinder).

LU creation and deletion on EMC Symmetrix and Symmetrix DMX

A logical unit (LU) that is exported by an EMC Symmetrix or Symmetrix DMX, meaning it is visible to a host, is either a *Symmetrix device* or a *Meta device*.

Symmetrix device

Restriction: An LU with a capacity of 32 MB or less is ignored by the SAN Volume Controller.

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 to 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. This allows an extremely large LU to be created, however, if exported to the SAN Volume Controller, only the first 2 TB is used.

Do not extend or reduce meta devices that are used for managed disks (MDisks). Reconfiguration of a meta device that is used for an MDisk causes unrecoverable data-corruption.

Configuring settings for the EMC Symmetrix and Symmetrix DMX

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)

Global settings for the EMC Symmetrix and Symmetrix DMX

Global settings apply across the EMC Symmetrix and Symmetrix DMX subsystems.

EMC Symmetrix and Symmetrix DMX 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 21 lists the EMC Symmetrix global settings that are supported by the SAN Volume Controller.

Table 21. EMC Symmetrix and Symmetrix DMX global settings supported by the SAN Volume Controller

Option	EMC Symmetrix and Symmetrix DMX 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

Port settings for the EMC Symmetrix and Symmetrix DMX

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 22 lists the EMC Symmetrix and Symmetrix DMX port settings that are supported by the SAN Volume Controller.

Table 22. EMC Symmetrix and Symmetrix DMX port settings supported by the SAN Volume Controller

Option	EMC Symmetrix and Symmetrix DMX default setting	SAN Volume Controller required setting
Disk_Array	Enabled	Disabled
Volume_Set_Addresssing	Enabled	Disabled

Table 22. EMC Symmetrix and Symmetrix DMX port settings supported by the SAN Volume Controller (continued)

Option	EMC Symmetrix and Symmetrix DMX 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

Logical unit settings for the EMC Symmetrix and Symmetrix DMX

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 23 lists the options that must be set for each LU that is accessed by the SAN Volume Controller.

Table 23. EMC Symmetrix and Symmetrix DMX LU settings supported by the SAN Volume Controller

Option	EMC Symmetrix and Symmetrix DMX default setting	SAN Volume Controller required setting
emulation	-	FBA
attribute	-	Set all attributes to disabled.

Mapping and virtualization settings for the EMC Symmetrix and Symmetrix DMX

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.

Volume Logix and masking

Volume Logix allows you to restrict access to particular WWPNs on the fabric for Symmetrix Volumes.

This function can be switched on and off by changing the VMC_State port setting. The SAN Volume Controller requires that you do not share target ports between a host and a SAN Volume Controller. However, you can still use Volume Logix to protect the subsystem from errors that can occur if the SAN is not correctly configured.

To mask a volume to the SAN Volume Controller, you must first identify the SAN Volume Controller ports that are connected to each subsystem. This can be done using the EMC Symmetrix symmask command.

The SAN Volume Controller automatically logs into any EMC Symmetrix subsystem it sees on the fabric. You can use the SAN Volume Controller svcinfo lsnod CLI command to find the correct port identifiers.

After you have identified the ports, you can map each volume on each port to each WWPN. The EMC Symmetrix stores the LUN masking in a database, so you must apply the changes you have made to refresh the contents of the database to view the changes.

Configuring the Fujitsu ETERNUS subsystems

This section provides information about configuring the Fujitsu ETERNUS subsystems for attachment to a SAN Volume Controller.

Supported models of the Fujitsu ETERNUS

The SAN Volume Controller supports models of the Fujitsu ETERNUS series of subsystems.

See the following Web site for the latest supported models: <http://www.ibm.com/storage/support/2145>

Supported firmware levels for the Fujitsu ETERNUS

The Fujitsu ETERNUS must use a firmware level that is supported by the SAN Volume Controller.

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

User interfaces on the Fujitsu ETERNUS

Ensure that you are familiar with the user interface application that is used by the Fujitsu ETERNUS.

You can use the ETERNUSmgr web-based configuration utility. See the documentation that is provided with the Fujitsu ETERNUS subsystem for more information.

Configuring the Fujitsu ETERNUS to use with the SAN Volume Controller

Ensure that you use the settings that are required to use the Fujitsu ETERNUS with the SAN Volume Controller. It is important that you use the correct settings to avoid data access problems.

Use the following sequence of steps to configure the Fujitsu ETERNUS subsystem:

1. Configure the SAN Volume Controller host response pattern.
2. Register the host world wide names (WWNs) and associate them with the host response pattern.
3. Setup the affinity group for SAN Volume Controller volumes or setup LUN mapping.
4. Create or reassign storage to the SAN Volume Controller.

For all other settings and procedures, consider the SAN Volume Controller a host. See the documentation that is provided with the Fujitsu ETERNUS subsystem.

CA parameters

The following table lists the port settings that are required. See the documentation that is provided with your Fujitsu ETERNUS subsystem for more information because some options are only available on certain models.

Option	Fujitsu ETERNUS default setting	SAN Volume Controller required setting
Connection Topology/FC Connection Settings	FC-AL Connection	Fabric Connection
Service Class	Class 3	Class 3
FC Transfer Rate	Auto Setting	Any
Reset Scope/Scope of LUR Actions	T_L	T_L Note: If this option is not set correctly, data corruption can occur.
Release Reservation upon Chip Reset	Enable/valid	Enable/valid
HP-UX Connection Setting	Disable	Disable
Frame Size Setting	2048	Any
Affinity/Addressing Mode	Off	Any

Host response pattern

The SAN Volume Controller requires that a new host response pattern is created. If the Host Affinity/Host Table Settings Mode is used, this host response pattern must be associated with each WWN. If the Host Affinity/Host Table Settings Mode is not used, this host response pattern must be associated with the target port.

The following table lists the settings that are required. See the documentation that is provided with your Fujitsu ETERNUS subsystem for more information because some options are only available on certain models.

Option	Fujitsu ETERNUS default setting	SAN Volume Controller required setting
Command timeout interval	Depends on the Fujitsu ETERNUS model	Default
Response status in overload	Unit Attention	Unit Attention
Byte 0 of Inquiry response/Response to inquiry commands	Default	Default

Option	Fujitsu ETERNUS default setting	SAN Volume Controller required setting
Inquiry Standard Data NACA Function	Disable	Disable
Inquiry Standard Data Version	Depends on the Fujitsu ETERNUS model	Default
Inquiry Command Page 83/Inquiry VPD ID Type	Depends on the Fujitsu ETERNUS model	Type 01
Reservation Conflict Response to Test Unit Ready Commands	Disable/Normal Response	Enable/Conflict Response
Target Port Group Access Support	Disable	Enable
Host Specific Mode	Normal Mode	Normal Mode
Response Sense at Firmware Hot Switching	Enable	Enable
Change LUN mapping	No Report	Report
LUN Capacity Expansion	No Report	Report
Aymmetric / Symmetric Logical Unit Access	Active/Active	Active/Active
Pattern of Sense Code Conversion	No Conversion	No Conversion

Notes:

1. If you set Inquiry VPD ID Type option to Type 3 on E4000 or E8000 range, the MDisks go offline.
2. If you set the Target Port Group Access Support option to Disabled on E3000 range, a 1370 error is shown in the error log.

Host WWNs

After the SAN Volume Controller is zoned on the fabric to see the Fujitsu ETERNUS, the subsystem might not initially appear in the list of controllers when you issue the **lscontroller** CLI command. This is normal and expected behavior.

See the documentation that is provided with the Fujitsu ETERNUS subsystem to add all SAN Volume Controller WWPNs as host WWNs. The following restrictions apply:

- The SAN Volume Controller WWNs must be associated with a host response pattern. The host response pattern must be defined prior to registration. If you use an incorrect or default host response pattern, you can lose access to data.
- All SAN Volume Controller WWNs must be registered on all Fujitsu ETERNUS ports on the same fabric. If the WWNs are not registered, you can lose access to data.

Affinity groups/zones

Use the affinity groups/zones mode to protect the SAN Volume Controller LUs if the SAN is incorrectly configured. The affinity group mode is setup in the CA

configuration. See the documentation that is provided with your Fujitsu ETERNUS subsystem for more information about using the affinity groups/zones mode. The following restrictions apply:

- Each SAN Volume Controller must have exactly one affinity group/zone.
- The SAN Volume Controller affinity group/zone must be associated with all SAN Volume Controller WWNs.

LUN mapping

You can use the LUN mapping mode (also called the zone settings mode for some models) with the following restrictions:

- The SAN zoning must only allow a single SAN Volume Controller access to this target port.
- The host response pattern must be set in CA configuration using the required SAN Volume Controller settings.

Note: If you use the LUN mapping mode, you cannot use the host affinity mode. The host affinity mode is set to OFF.

Assigning storage to the SAN Volume Controller

Ensure that you understand all SAN Volume Controller and Fujitsu ETERNUS restrictions before you assign storage to the SAN Volume Controller. See the documentation that is provided with the Fujitsu ETERNUS subsystem for more information.

Zoning configuration for the Fujitsu ETERNUS

If LUN mapping mode is used for a Fujitsu ETERNUS port, you must exclusively zone the SAN Volume Controller with this target port.

Migrating logical units from the Fujitsu ETERNUS to the SAN Volume Controller

You can use the standard migration procedure with the following restrictions:

- The SAN Volume Controller must have software level 4.2.0 or higher installed before you start migration. Upgrades from previous SAN Volume Controller software levels to software level 4.2.0 or higher causes all Fujitsu ETERNUS subsystems that are attached to be excluded.
- You must configure the Fujitsu ETERNUS subsystem to work with the SAN Volume Controller before you start migration.
- The subsystem device driver (SDD) and Fujitsu Multipath driver cannot coexist.
- The SAN Volume Controller must support all host code levels.

Concurrent maintenance on the Fujitsu ETERNUS

Concurrent maintenance is the capability to perform I/O operations to a Fujitsu ETERNUS while simultaneously performing maintenance operations on it.

You can perform nondisruptive maintenance procedures concurrently on the following components:

- Fujitsu ETERNUS controller module
- Fujitsu ETERNUS controller cache
- Fujitsu ETERNUS cache battery pack

- Fan
- Power supply
- Disk drive
- SFP transceiver

See the documentation that is provided with the Fujitsu ETERNUS subsystem for more information.

Advanced functions for the Fujitsu ETERNUS

The Fujitsu ETERNUS subsystem provides several Advanced Copy functions. Do not use these Advanced Copy functions for storage that is managed by the SAN Volume Controller, even if the VDisk cache is disabled.

Configuring the IBM TotalStorage ESS subsystem

This section provides information about configuring the IBM TotalStorage Enterprise Storage Server (ESS) for attachment to a 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 this option 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. Click **Perform Configuration Update**.

Important: If you are adding SAN Volume Controller ports to a volume that is already assigned to other SAN Volume Controller ports, you must select the **Use same ID/LUN in source and target** check box.

Supported models of the IBM ESS

The SAN Volume Controller supports models of the IBM Enterprise Storage Server (ESS).

See the following Web site for the latest supported models:

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

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/storage/support/2145>

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.

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.

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/storage/support/2145>

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.

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.

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).

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 additional SAN Volume Controller ports to an existing 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. If this checkbox is not available, the option is not required. The detect MDisks task in the SAN Volume Controller Console or the `svctask detectmdisk` command-line interface (CLI) command must be run for the SAN Volume Controller to detect the new disks.

Configuring the IBM System Storage DS4000 (formerly FAStT) series subsystem

This section provides information about configuring the IBM System Storage DS4000 series subsystem for attachment to a SAN Volume Controller. Certain models of the IBM 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.

Configuring IBM DS4000 series disk controllers for the storage server

The IBM DS4000 series of disk controllers are compatible with the SAN Volume Controller.

Attention: The SAN Volume Controller does not concurrently support I/O operations with the download of Environmental Services Monitor (ESM) 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:

- Do not use the IBM DS4000 series Copy Services when the SAN Volume Controller is attached to the IBM DS4000 series.
- You can use partitioning to allow IBM DS4000 series Copy Services usage for 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 IBM DS4000 series 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.

Supported options of the IBM DS4000 series controller

The IBM DS4000 series disk controllers provide functions 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 are 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

You can redistribute logical drives or change ownership of the preferred path; however, these options might not take effect until a discovery is started on the SAN Volume Controller cluster. You can use the **svctask detectmdisk** command-line interface (CLI) command to restart a cluster discovery process. 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 must 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 is automatically recovered by the SAN Volume Controller software.

Check your MDisks to ensure that they have not been set to the degraded state during the controller reset process. You can issue the **svctask includemdisk** CLI command to repair degraded MDisks.

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 Sun StorageTek and StorageTek models; therefore, the SAN Volume Controller also supports models of Sun StorageTek's series, and StorageTek's FlexLine series and D series.

See the following Web site for the latest supported models:

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

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:

- EXP420
- EXP500
- DS4000 EXP100 R2

- DS4000 EXP700
- EXP710

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/storage/support/2145>

The Web site includes the maximum number of LUNs per partition that are supported by the firmware level.

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.

See your IBM DS4000 series documentation for information about concurrent maintenance.

Attention: The SAN Volume Controller does not concurrently support I/O operations with the download of Environmental Services Monitor (ESM) 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.

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.

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.

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.

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).

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 using 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 identified to the assigned host fibre ports (no subpartition 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 make the LUNs known to the SAN Volume Controller, which then configures these LUNs as image-mode LUNs and identifies 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 one storage partition, so that you can move all the LUNs from all the other partitions into one 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 eight 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 four partitions on the IBM DS4000 series such that each partition is mapped to one 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 four 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.

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.

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 options 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.

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 the following scope:

- 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.
 - You must have the following options 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.

Global settings for the IBM DS4000 series

Global settings apply across an IBM DS4000 series controller.

Table 24 on page 286 lists the IBM DS4000 series global settings that are supported by the SAN Volume Controller.

Table 24. IBM DS4000 series controller global settings supported by the SAN Volume Controller

Option	IBM 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 the IBM Support Center.

Use a host type of IBM TS SAN VCE 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.

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 the IBM Support Center.

The following options 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.

Miscellaneous settings for the IBM DS4000 series

The SAN Volume Controller supports all media scan settings that are provided by the IBM DS4000. Set the background media scan to enabled and set the frequency to 30 days. These settings are enabled at both the subsystem level and the individual logical drive level.

See your IBM DS4000 series documentation for information about other settings.

Configuring the IBM System Storage DS6000 subsystem

This section provides information about configuring the IBM System Storage DS6000 subsystem for attachment to a SAN Volume Controller.

Configuring the IBM DS6000

The IBM DS6000 provides functions that are compatible with the SAN Volume Controller.

After you have defined at least one storage complex, storage unit, and I/O port, you can define the SAN Volume Controller as a host and create host connections. If you have not defined all of these required storage elements, use the IBM System Storage DS6000 Storage Manager or the IBM DS6000 command-line interface (CLI) to define these elements and return to this topic after they are configured.

This task assumes that you have already launched the IBM System Storage DS6000 Storage Manager.

Perform the following steps to configure the IBM DS6000:

1. Click **Real-time manager** → **Manage hardware** → **Host systems**.
2. Select **Create** from the **Select Action** list. The Create Host System wizard is displayed.
3. Perform the following steps to select a host type:
 - a. Select **IBM SAN Volume Controller (SVC)** from the **Host Type** list.
 - b. Enter a unique name of up to 16 characters for each port in the **Nickname** field. The value that you enter in this field appears in other panels when you select defined hosts. This is a required field.
 - c. Optionally, enter a detailed description of up to 256 characters in the **Description** field.
 - d. Click **Next**. The Define host wizard panel is displayed.
4. Perform the following steps in the Define host panel:
 - a. Enter the number of ports that you are defining for the SAN Volume Controller node in the **Quantity** field.

Note: You must add all of the SAN Volume Controller node ports.
 - b. Select **FC Switch fabric (P-P)** from the **Attachment Port Type** list.
 - c. Click **Add**.
 - d. Select **Group ports to share a common set of volumes**.
 - e. Click **Next**. The Define host WWPN panel is displayed.
5. Specify a WWPN for each SAN Volume Controller node port that you are configuring. After you have defined all SAN Volume Controller node port WWPNs, click **Next**.
6. Perform the following steps in the Specify storage units panel:
 - a. Select all the available storage units that use the ports that you defined in step 5.
 - b. Click **Add** to move the selected storage units to the **Selected storage units** field.
 - c. Click **Next**. The Specify storage units parameters panel is displayed.
7. Perform the following steps in the Specify storage units parameters panel:
 - a. Select a host attachment identifier from the table.
 - b. Click **the following specific storage unit I/O ports** in the **This host attachment can login to** field. The available ports are displayed in the Available storage unit I/O ports table.
 - c. Select each port in the Available storage unit I/O ports table.

Note: The **Type** for each port should be **FcSf**. If the listed type is not **FcSf**, click **Configure I/O Ports**. The Configure I/O Ports panel is displayed. Click the port that you want to configure and select **Change to FcSf** from the **Select Action** list.

- d. Click **Apply assignment**.
- e. Click **OK**. The Verification panel is displayed.
8. Verify that the attributes and values that are displayed in the table are correct.
9. Click **Finish** if the values that are displayed in the table are correct. Otherwise, click **Back** to return to the previous panels and change the values that are not correct.

Supported firmware levels for the IBM DS6000

The IBM DS6000 must use a firmware level that is supported by the SAN Volume Controller.

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

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

Supported models of the IBM DS6000 series

The SAN Volume Controller supports models of the IBM DS6000 series of controllers.

See the following Web site for the latest supported models:

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

User interfaces on the IBM DS6000

Ensure that you are familiar with the user interfaces that support the IBM DS6000.

Web server

You can manage, configure, and monitor the IBM DS6000 through the IBM System Storage DS6000 Storage Manager.

CLI

You can also manage, configure, and monitor the IBM DS6000 through the IBM System Storage DS command-line interface.

Concurrent maintenance on the IBM DS6000

Concurrent maintenance is the capability to perform I/O operations to an IBM DS6000 while simultaneously performing maintenance operations on it.

All IBM DS6000 concurrent maintenance procedures are supported.

Target port groups on the IBM DS6000

The IBM DS6000 uses the SCSI Target Port Groups feature to indicate a preferred path for each logical unit (LU).

Configuring the IBM System Storage DS8000 subsystem

This section provides information about configuring the IBM System Storage DS8000 subsystem for attachment to a SAN Volume Controller.

Configuring the IBM DS8000

The IBM DS8000 provides functions that are compatible with the SAN Volume Controller.

After you have defined at least one storage complex, storage unit, and I/O port, you can define the SAN Volume Controller as a host and create host connections. If you have not defined all of these required storage elements, use the IBM System Storage DS8000 Storage Manager or the IBM System Storage DS command-line interface to define these elements and return to this topic after they are configured.

This task assumes that you have already launched the IBM System Storage DS8000 Storage Manager.

Perform the following steps to configure the IBM DS8000:

1. Click **Real-time manager** → **Manage hardware** → **Host systems**.
2. Select **Create** from the **Select Action** list. The Create Host System wizard begins.
3. Perform the following steps to select a host type:
 - a. Select **IBM SAN Volume Controller (SVC)** from the **Host Type** list.
 - b. Enter a unique name of up to 16 characters for each port in the **Nickname** field. The value that you enter in this field is displayed in other panels when you select defined hosts. This is a required field.
 - c. Optionally, enter a detailed description of up to 256 characters in the **Description** field.
 - d. Click **Next**. The Define host wizard panel is displayed.
4. Perform the following steps in the Define host panel:
 - a. Enter the number of ports that you are defining for the SAN Volume Controller node in the **Quantity** field.

Note: You must add all of the SAN Volume Controller node ports.
 - b. Select **FC Switch fabric (P-P)** from the **Attachment Port Type** list.
 - c. Click **Add**.
 - d. Select **Group ports to share a common set of volumes**.
 - e. Click **Next**. The Define host WWPN panel is displayed.
5. Specify a WWPN for each SAN Volume Controller node port that you are configuring. When you have defined all SAN Volume Controller node port WWPNs, click **Next**.
6. Perform the following steps in the Select storage images panel:
 - a. Select all the available storage units that use the ports that you defined in the previous step.
 - b. Click **Add** to move the selected storage units to the **Select storage images** field.
 - c. Click **Next**. The Specify storage image parameters panel is displayed.
7. Perform the following steps in the Specify storage image parameters panel:
 - a. Select a host attachment identifier from the table.

- b. Click the following specific storage image I/O ports in the **This host attachment can login to** field. The available ports are displayed in the Available storage unit I/O ports table.
- c. Select each port in the Available storage unit I/O ports table.

Note: The **Type** for each port should be **FcSf**. If the listed type is not **FcSf**, click **Configure I/O Ports**. The Configure I/O Ports panel is displayed. Click the port that you want to configure and select **Change to FcSf** from the **Select Action** list.

- d. Click **Apply assignment**.
- e. Click **OK**. The Verification panel is displayed.
8. Verify that the attributes and values that are displayed in the table are correct.
9. Click **Finish** if the values that are displayed in the table are correct. Otherwise, click **Back** to return to the previous panels and change the incorrect values.

Supported firmware levels for the IBM DS8000

The SAN Volume Controller supports the IBM DS8000 series.

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

Supported models of the IBM DS8000

The SAN Volume Controller supports models of the IBM DS8000 series of controllers.

See the following Web site for the latest supported models:

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

User interfaces on the IBM DS8000

Ensure that you are familiar with the user interfaces that support the IBM DS8000.

Web server

You can manage, configure, and monitor the IBM DS8000 through the IBM System Storage DS8000 Storage Manager.

CLI

You can also manage, configure, and monitor the IBM DS8000 through the IBM System Storage DS command-line interface.

Concurrent maintenance for the IBM DS8000

Concurrent maintenance is the capability to perform I/O operations to an IBM DS8000 while simultaneously performing maintenance operations on it.

All IBM DS8000 concurrent maintenance procedures are supported.

Configuring the HDS Lightning series subsystem

This section provides information about configuring the Hitachi Data Systems (HDS) Lightning series subsystem for attachment to a SAN Volume Controller.

The information in this section also applies to the supported models of the Sun StorEdge series and the HP XP series.

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 HP XP models; therefore, the SAN Volume Controller also supports models of the Sun StorEdge and the HP XP.

See the following Web site for the latest supported models:

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

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/storage/support/2145>

Note: Concurrent upgrade of the controller firmware is supported with the SAN Volume Controller.

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.

User interface on HDS Lightning

Ensure that you are familiar with the user interface application that supports the HDS Lightning subsystem.

Service Processor (SVP)

HDS Lightning has a laptop in the controller frame. The laptop runs the Service Processor (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.

Sharing the HDS Lightning 99xxV between a host and the SAN Volume Controller

There are restrictions for sharing an HDS Lightning 99xxV between a host and a SAN Volume Controller.

Sharing ports

The HDS Lightning 99xxV can be shared between a host and a SAN Volume Controller under the following conditions:

- The same host cannot be connected to both a SAN Volume Controller and an HDS Lightning at the same time because the Hitachi HiCommand Dynamic Link Manager (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 that 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 under the following conditions:

- 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 LUN x , the LU must map as LUN x to all the SAN Volume Controller ports in the cluster and must also map as LUN x through all of the controller ports that it is mapped to.
- Command LUNs must not be mapped to the SAN Volume Controller.
- Logical Unit Size 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 connect the SAN Volume Controller to the HDS Lightning.

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.

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, and 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.

Logical Unit Size 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 operations are functionally similar to Metro Mirror. TrueCopy processing 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 processing 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.

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 that are defined. These LUs can be expanded by merging 2 - 36 of these LUs 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.

Logical unit creation and deletion on HDS Lightning

The SAN Volume Controller supports Logical Unit Size 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, do not use LUSE on the disk before you import the data.

Configuring settings for HDS Lightning

The Lightning configuration interface provides functions for configuration.

These options and settings can have the following scope:

- Subsystem
- Port
- Logical unit (LU)

Global settings for HDS Lightning

Global settings apply across an HDS Lightning disk controller system.

Table 25 lists the global settings for HDS Lightning.

Table 25. 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

Controller settings for HDS Lightning

Controller settings are settings that apply across the entire HDS Lightning controller.

Table 26 lists the HDS Lightning controller settings that are supported by the SAN Volume Controller.

Table 26. HDS Lightning controller settings that are supported by the SAN Volume Controller

Option	HDS Lightning default setting	SAN Volume Controller required setting
PCB mode	Standard	Standard

Port settings for HDS Lightning

Port settings are configurable at the port level.

There are no available options 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 27 lists the HDS Lightning port settings that are supported by the SAN Volume Controller.

Table 27. 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

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 28 if the LUN is associated with ports in a switch zone that is accessible to the SAN Volume Controller.

Table 28. 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.

Configuring the HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS

You can attach Hitachi Data Systems (HDS) Thunder, HDS TagmaStore Adaptable Modular Storage (AMS), and HDS TagmaStore Workgroup Modular Storage (WMS) subsystems to a SAN Volume Controller cluster.

Note: In Japan, the HDS Thunder 9200 is referred to as the HDS SANrise 1200. Therefore, the information in this section that refers to the HDS Thunder 9200 also applies to the HDS SANrise 1200.

Supported models of the HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS

You can attach models of the HDS Thunder, HDS TagmaStore Adaptable Modular Storage (AMS), and HDS TagmaStore Workgroup Modular Storage (WMS) subsystems to SAN Volume Controller clusters.

See the following Web site for the latest supported models:

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

Supported firmware levels for HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS

The SAN Volume Controller supports models of the HDS Thunder, HDS TagmaStore Adaptable Modular Storage (AMS), and HDS TagmaStore Workgroup Modular Storage (WMS) subsystems.

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

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

Concurrent maintenance on the HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS

Concurrent maintenance is the capability to perform I/O operations to a subsystem while simultaneously performing maintenance operations on it.

Important: An HDS Field Engineer must perform all maintenance operations.

The SAN Volume Controller supports concurrent hardware maintenance and firmware upgrade operations on these subsystems.

User interface on the HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS

Ensure that you are familiar with the user interface applications that support the Hitachi Data Systems (HDS) Thunder, HDS TagmaStore Adaptable Modular Storage (AMS), and HDS TagmaStore Workgroup Modular Storage (WMS) subsystems.

In-band configuration

Disable the subsystem command LUN when you use the user interface applications.

Storage Navigator Modular GUI

The Storage Navigator Modular (SNM) is the primary user interface application for configuring HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems. Use SNM to upgrade firmware, change settings, and to create and monitor storage.

SNM supports an Ethernet connection to the subsystem. An out-of-band command-line interface is available with SNM that supports the majority of the functions that are provided in SNM.

HiCommand

HiCommand is another configuration user interface that is available for the HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems. You must have access to SNM to use HiCommand to configure settings. HiCommand only allows basic creation of storage and provides some monitoring features.

HiCommand uses Ethernet to connect to the subsystem.

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 you put a controller into diagnostic mode by pressing the reset button on the controller, the user interface allows firmware upgrades and subsystem configuration resets.

Sharing the HDS Thunder, HDS TagmaStore AMS, or HDS TagmaStore WMS between a host and the SAN Volume Controller

You can share the HDS Thunder, HDS TagmaStore Adaptable Modular Storage (AMS), and HDS TagmaStore Workgroup Modular Storage (WMS) subsystems between a host and a SAN Volume Controller cluster with certain restrictions.

The following restrictions apply:

- The same host cannot be connected to both a SAN Volume Controller cluster and an HDS Thunder, HDS TagmaStore AMS, or HDS TagmaStore WMS at the same time because Hitachi Dynamic Link Manager (HDLM) and the subsystem device driver (SDD) cannot coexist.
- For the HDS Thunder 9200, a target port cannot be shared between a host and a SAN Volume Controller cluster. If a target port is used by a SAN Volume Controller cluster, it cannot be present in a switch zone that allows a host to access the port.
- Logical units (LUs) cannot be shared between a host and a SAN Volume Controller cluster. The Thunder 9200 must be set into M-TID M-LUN mode and Mapping Mode must be enabled on Thunder 95xx. No LU can have a LUN

Number that is associated with a port that is zoned for host use while also having a LUN Number that is associated with a port that is zoned for a SAN Volume Controller cluster.

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 must be in each group.
3. Make a note of all of the LUNs that are currently on the array. Each LUN that you allow the SAN Volume Controller cluster to manage must be in one group.
4. Divide the LUNs into groups: You must have one group of LUNs for each group of ports.

5. Perform the followings steps 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.

Quorum disks on HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS

When a SAN Volume Controller cluster initializes, the cluster can choose managed disks (MDisks) that are presented by HDS Thunder, HDS TagmaStore Adaptable Modular Storage (AMS) and HDS TagmaStore Workgroup Modular Storage (WMS), subsystems as quorum disks.

You can use the set quorum disk CLI command or the SAN Volume Controller Console to select quorum disks.

Advanced functions for HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS

Some advanced functions of the HDS Thunder, HDS TagmaStore Adaptable Modular Storage (AMS) and HDS TagmaStore Workgroup Modular Storage (WMS), subsystems are not supported by the SAN Volume Controller clusters.

Advanced copy functions

Advanced copy functions for HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems are not supported for disks that are managed by the SAN Volume Controller clusters because the copy function does not extend to the SAN Volume Controller cache. For example, ShadowImage, TrueCopy, and HiCopy are not supported.

LUN Security

LUN Security enables LUN masking by the worldwide node name (WWNN) of the initiator port. This function is not supported for logical units (LUs) that are used by SAN Volume Controller clusters.

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 cluster and HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems support the partitioning function.

Dynamic array expansion

The HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems allow the last LU that is defined in a RAID group to be expanded. This function is not supported when these storage subsystems are attached to a SAN Volume Controller cluster. Do *not* perform dynamic array expansion on LUs that are in use by a SAN Volume Controller cluster.

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, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems support 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 the ports.

Logical unit creation and deletion on HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems

The HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems Storage Navigator Modular Graphical User Interface (GUI) enables you to create and delete LUNs. You must avoid certain creation and deletion scenarios to prevent data corruption.

Creation and deletion scenarios

For example, the Storage Navigator Modular GUI enables you to create LUN A, delete LUN A, and then create LUN B with the same unique ID as LUN A. If a

SAN Volume Controller cluster is attached, data corruption can occur because the cluster might not realize that LUN B is different than LUN A.

Attention: Before you use the Storage Navigator Modular GUI to delete a LUN, remove the LUN from the managed disk group that contains it.

Adding LUNs dynamically

To prevent the existing LUNs from rejecting I/O operations during the dynamic addition of LUNs, perform the following procedure to add LUNs:

1. Create the new LUNs using the Storage Navigator Modular GUI.
2. Quiesce all I/O operations.
3. Perform either an offline format or an online format of all new LUNs on the controller using the Storage Navigator Modular GUI. Wait for the format to complete.
4. Go into the LUN mapping function of the Storage Navigator Modular GUI. Add mapping for the new LUN to all of the controller ports that are available to the SAN Volume Controller cluster on the fabric.
5. Restart the controller. (Model 9200 only)
6. After the controller has restarted, restart I/O operations.

LUN mapping considerations

If LUN mapping is used as described in the LUN mapping topic, you must restart the controller to pick up the new LUN mapping configuration. For each managed disk group that contains an MDisk that is supported by an LU on the subsystem, all virtual disks in those managed disk groups go offline.

Configuring settings for HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems

The Storage Navigator Modular GUI configuration interface provides functions for configuration.

These options and settings can have the following scope:

- Subsystem
- Port
- Logical unit

Global settings for the HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems

Global settings apply across HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems.

Table 29 lists the global settings for these disk subsystems.

Table 29. HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems global settings supported by the SAN Volume Controller

Option	Default setting	SAN Volume Controller required setting
Start attribute	Dual active mode	Dual active mode
SCSI ID/Port takeover mode	Not applicable	Not applicable
Default controller	Not applicable	Not applicable

Table 29. HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems global settings supported by the SAN Volume Controller (continued)

Option	Default setting	SAN Volume Controller required setting
Data-share mode	Used	Used
Serial number		Same as the subsystem 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	Not applicable	Not applicable
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>
Web title	<Empty>	Any setting supported
Cache mode (Thunder 9200 only)	All off	All off

Table 29. HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems global settings supported by the SAN Volume Controller (continued)

Option	Default setting	SAN Volume Controller required setting
Link separation (Thunder 9200 only)	False	False
ROM Pseudo-response command processing (Thunder 9200 only)	Not applicable	Not applicable
Save data pointer response (Thunder 9200 only)	Not applicable	Not applicable
Controller identifier	False	False
RS232C error information outflow mode	Off	Any
Execute write and verify mode	True	True

Controller settings for HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems

Controller settings apply across the entire HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems. Options are not available within the scope of a single controller.

Port settings for the HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems

Port settings are configurable at the port level.

The settings listed in Table 30 apply to disk controllers that are in a switch zone that contains SAN Volume Controller nodes. If the subsystem is shared between a SAN Volume Controller cluster and another host, you can configure with different settings than shown if both of the following conditions are true:

- 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 cluster

There are no options available with the scope of a single controller.

Table 30. HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystem port settings supported by the SAN Volume Controller

Option	Default setting	SAN Volume Controller required setting
Host connection mode 1	Standard	Standard
VxVM DMP mode (HDS Thunder 9200 only)	False	False
HP connection mode	False	False
Report inquiry page 83H (HDS Thunder 9200 only)	False	True
UA (06/2A00) suppress mode	False	False
HISUP mode	False	False
CCHS mode	False	False

Table 30. HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystem port settings supported by the SAN Volume Controller (continued)

Option	Default setting	SAN Volume Controller required setting
Standard inquiry data expand (HDS Thunder 9200 only)	False	False
Host connection mode 2	False	False
Product ID DF400 mode	False	False
HBA WWN report mode (HDS 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 (HDS 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 (HDS Thunder 950xxV only)	False	False
Topology	Point-to-point	Point-to-point

Logical unit settings for the HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems

Logical unit (LU) settings apply to individual LUs that are configured in the HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems.

You must configure the subsystems LUs as described in Table 31 if the logical unit number (LUN) is associated with ports in a switch zone that is accessible to the SAN Volume Controller cluster.

Table 31. HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems 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 that are accessible by the SAN Volume Controller cluster.

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. Because 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 cluster because this can 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. If the LUN is managed by a SAN Volume Controller cluster, this scenario can cause data corruption because the cluster might not recognize that LUN B is different than LUN A.

Mapping and virtualization settings for HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems

The HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems support different modes of operation. These modes affect LUN mapping or masking and virtualization.

The 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. You must restart the controllers for changes to LUN mapping to take effect.

Attention: The HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS subsystems do not provide an interface that enables a SAN Volume Controller cluster to detect and ensure that the mapping or masking and virtualization options are set properly. Therefore, you must ensure that these options are set as described in this topic.

S-TID M-LUN modes

In S-TID M-LUN mode all LUs are accessible through all ports on the subsystem with the same LUN number on each port. You can use this mode in environments where the subsystem is not being shared between a host and a SAN Volume Controller cluster.

M-TID M-LUN modes

If a subsystem is shared between a host and a SAN Volume Controller cluster, you must use M-TID M-LUN mode. Configure the subsystem so that each LU that is exported to the SAN Volume Controller cluster can be identified by a unique LUN. The LUN must be the same on all ports through which the LU can be accessed.

Example

A SAN Volume Controller cluster can access controller ports x and y. The cluster also sees an LU on port x that has LUN number p. In this situation the following conditions must be met:

- The cluster 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 cannot appear as any other LUN number on port y.
- The LU must not be mapped to any subsystem port that is zoned for use directly by a host in a configuration where the subsystem is shared between a host and a cluster.

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 can be LUN 0 on port 1, LUN 3 on port 2, and not visible at all on ports 3 and 4.

Important: 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 can be LUN 1 and LUN 2 on controller port 1.

Important: The SAN Volume Controller does not support this.

Configuring the HDS TagmaStore USP and NSC subsystems

This section provides information about configuring the Hitachi Data Systems (HDS) TagmaStore Universal Storage Platform (USP) and Network Storage Controller (NSC) subsystems for attachment to a SAN Volume Controller. Models of the HDS USP and NSC are equivalent to HP and Sun models; therefore, the SAN Volume Controller also supports models of the HP StorageWorks XP series and the Sun StorEdge series.

The information in this section also applies to the supported models of the HP XP and the Sun StorEdge series.

Supported models of the HDS USP and NSC

The SAN Volume Controller supports models of the Hitachi Data Systems (HDS) Universal Storage Platform (USP) and Network Storage Controller (NSC) series. Models of the HDS USP and NSC are equivalent to HP and Sun models; therefore, the SAN Volume Controller also supports models of the Sun StorEdge and the HP XP series.

See the following Web site for the latest supported models:

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

Supported firmware levels for HDS USP and NSC

The SAN Volume Controller supports the HDS USP and NSC series of controllers.

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

User interface on the HDS USP and NSC

Ensure that you are familiar with the user interface application that supports the HDS USP and NSC. The HDS USP and NSC is configured, managed, and monitored by a Service Processor (SVP). The SVP is a server that is connected to the HDS USP or NSC through a private local area network (LAN).

Web server

The HDS USP and NSC use the Storage Navigator as the main configuration GUI. The Storage Navigator GUI runs on the SVP and is accessed through a Web browser.

Logical units and target ports on the HDS USP and NSC

Logical units (LUs) that are exported by the HDS USP and NSC report identification descriptors in the vital product data (VPD). The SAN Volume Controller uses the LUN associated binary type-3 IEEE Registered Extended descriptor to identify the LU.

An LU path must be defined before an LU can be accessed by a host. The LU path relates a host group to a target port and to a set of LUs. Host initiator ports are added to the host group by worldwide port name (WWPN).

The HDS USP and NSC do not use LU groups so all LUs are independent. The LU access model is active-active and does not use preferred access ports. Each LU can be accessed from any target port that is mapped to the LU. Each target port has a unique WWPN and worldwide node name (WWNN). The WWPN matches the WWNN on each port.

Special LUs

The HDS USP and NSC can use any logical device (LDEV) as a Command Device. Command Devices are the target for HDS USP or NSC copy service functions. Therefore, do not export Command Devices to a SAN Volume Controller.

Switch zoning limitations for the HDS USP and NSC

There are limitations in switch zoning for the SAN Volume Controller and the HDS USP or NSC.

The SAN Volume Controller can be connected to the HDS USP or NSC with the following restrictions:

- If an LU is mapped to a SAN Volume Controller port as LUN x , the LU must appear as LUN x for all mappings to target ports.
- Only fibre-channel connections can be used to connect a SAN Volume Controller to the HDS USP or NSC subsystem.
- Because the SAN Volume Controller limits the number of worldwide node names (WWNNs) for each storage subsystem and the HDS USP and NSC present a separate WWNN for each port, the number of target ports that the SAN Volume Controller can resolve as one storage subsystem is limited. Perform the following steps to provide connections to more target ports:
 1. Divide the set of target ports into groups of 2 to 4.
 2. Assign a discrete set of LUs to each group.

The SAN Volume Controller can then view each group of target ports and the associated LUs as separate HDS USP or NSC subsystems. You can repeat this process to use all target ports.

Controller splitting

You can split the HDS USP or NSC between other hosts and the SAN Volume Controller under the following conditions:

- A host cannot be simultaneously connected to both an HDS USP or NSC and a SAN Volume Controller.
- Port security must be enabled for target ports that are shared.
- An LU that is mapped to a SAN Volume Controller cannot be simultaneously mapped to another host.

Concurrent maintenance on the HDS USP and NSC

Concurrent maintenance is the capability to perform I/O operations to an HDS USP or NSC while simultaneously performing maintenance operations on it. Concurrent firmware upgrades are supported with the SAN Volume Controller.

Important: An HDS Field Engineer must perform all maintenance procedures.

Quorum disks on HDS USP and NSC

HDS USP and NSC subsystems are not approved hosts for quorum disks. Therefore, configurations that consist of a SAN Volume Controller cluster that is attached to only an HDS USP or NSC is not supported.

Host type for HDS USP and NSC

When the HDS USP and NSC subsystems are attached to a SAN Volume Controller, set the host mode attribute to Windows for each host group.

Advanced functions for HDS USP and NSC

Some advanced functions of the HDS USP and NSC are not supported by the SAN Volume Controller.

Advanced subsystem functions

The following advanced subsystem functions for HDS USP and NSC are not supported for disks that are managed by the SAN Volume Controller:

- TrueCopy
- ShadowImage
- Extended Copy Manager
- Extended Remote Copy
- NanoCopy
- Data migration
- RapidXchange
- Multiplatform Backup Restore
- Priority Access
- HARBOR File-Level Backup/Restore
- HARBOR File Transfer

- FlashAccess

Advanced SAN Volume Controller functions

All advanced SAN Volume Controller functions are supported on logical unit (LU) that are exported by the HDS USP or NSC subsystem.

LU Expansion

The HDS USP and NSC support Logical Unit Expansion (LUSE). LUSE is *not* a concurrent operation. LUSE allows you to create a single LU by concatenating logical devices (LDEVs). Before LUSE can be performed, the LDEVs must be unmounted from hosts and paths must be removed.

Attention:

1. LUSE destroys all data that exists on the LDEV.
2. Do not perform LUSE on any LDEV that is used to export an LU to a SAN Volume Controller.

If data exists on an LDEV and you want to use image mode migration to import the data to a SAN Volume Controller, do not perform LUSE on the disk before you import the data.

LUs that are created using LUSE can be exported to a SAN Volume Controller.

Virtual LVI/LUNs

The HDS USP and NSC support Virtual LVI/LUNs (VLL). VLL is *not* a concurrent operation. VLL allows you to create several LUs from a single LDEV. You can only create new LUs from free space on the LDEV.

Attention: Do not perform VLL on disks that are managed by the SAN Volume Controller.

LUs that are created using VLL can be exported to a SAN Volume Controller.

Configuring HP StorageWorks MA and EMA subsystems

This section provides information about configuring HP StorageWorks Modular Array (MA) and Enterprise Modular Array (EMA) subsystems for attachment to a SAN Volume Controller.

Both the HP MA and EMA use an HSG80 controller.

HP MA and EMA 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 that are 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.
mirrorset	See <i>RAID 1</i> .	mirrorset	A RAID storageset of two or more physical disks that maintains a complete and independent copy of all data on 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.
non-RAID	Disks that are not in a redundant array of independent disks (RAID).	non-RAID	See <i>just a bunch of disks</i> .

IBM term	IBM definition	HP term	HP definition
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 stripe sets.
RAID 1	A form of storage array in which two or more identical copies of data are maintained on separate media. Also known as mirror set.	RAID 1	See <i>mirror set</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>RAID set</i> .
RAID set	See <i>RAID 5</i> .	RAID set	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.
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.
stripe set	See <i>RAID 0</i> .	stripe set	See <i>RAID 0</i> .

Configuring HP MA and EMA subsystems

The HP MA and EMA subsystems provide functions that are compatible with the SAN Volume Controller.

This task assumes that the subsystem is not in use.

Note: When you configure a SAN Volume Controller cluster to work with an HP MA or EMA, you must not exceed the limit of 96 process logins.

Perform the following procedure to enable support of an HP, MA, or EMA subsystem.

1. Verify that the front panel of the SAN Volume Controller is clear of errors.
2. Ensure that the HP StorageWorks Operator Control Panel (OCP) on each subsystem is clear of errors. The Operator Control Panel consists of seven green LEDs 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. Ensure that the subsystem firmware is at a supported level. See the following Web site for the latest firmware support:
<http://www.ibm.com/storage/support/2145>.
 - 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 MIRRORING_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 do not specify Host Based Logging.
8. Issue the **SHOW CONNECTIONS FULL** command to verify that you have enough spare entries for all combinations of SAN Volume Controller ports and HP MA or EMA ports.
9. Connect up to four fibre-channel cables between the fibre-channel switches and the HP MA or EMA subsystem.
10. Ensure that the fibre-channel switches are zoned so that the SAN Volume Controller and the HP MA or EMA subsystem are in a zone.

11. Issue the **SHOW THIS** command and **SHOW OTHER** command to verify that each connected port is running. The following is an example of the output that is displayed: `PORT_1_TOPOLOGY=FABRIC`.
12. Issue the **SHOW CONNECTIONS FULL** command to verify that the new connections have been created for each SAN Volume Controller port and HP MA or EMA port combination.
13. Verify that No rejected hosts is displayed at the end of the **SHOW CONNECTIONS** output.
14. Perform the following steps from the SAN Volume Controller command-line interface (CLI):
 - a. Issue the **svctask detectmdisk** CLI command to discover the controller.
 - b. Issue the **svcinfolcontroller** CLI command to verify that the two HSG80 serial numbers appear under the `ctrl s/n`.
 - c. Issue the **svcinfolmdisk** CLI command to verify that the additional MDisks that correspond to the UNITS shown in the HP MA or EMA subsystem.

You can now use the SAN Volume Controller CLI commands to create an MDisk group. You can also create and map VDIs from these MDisk groups. Check the front panel of the SAN Volume Controller to ensure that there are no errors. After the host has reloaded the fibre-channel driver, you can perform I/O to the VDIs. See the *IBM System Storage SAN Volume Controller: Host Attachment Guide* for detailed information.

Partitioning LUNs on HP MA and EMA subsystems

A SAN Volume Controller cluster that has software level 1.2.0 or later supports up to four fibre-channel connections per HP MA or EMA subsystem. The support for partitioned LUNs is restricted to a single fibre-channel connection.

Attention: Before you make any changes to your subsystem, back up important application data.

You can use the HP StorageWorks command **SHOW UNITS** to display all LUNs that are partitioned. Table 32 provides an example of the information that is provided by the **SHOW UNITS** command.

Table 32. 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 MA or EMA subsystem that is or will be connected to the SAN Volume Controller cluster.

Ensure that all SAN Volume Controller nodes have software level 1.2.0 or later installed. If software 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 an HP MA or EMA subsystem with a SAN Volume Controller cluster that has installed level 1.1.1, and that uses a single fibre-channel attachment or zone. If partitioned units are present on the HP MA or EMA subsystem, the following options are 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 upgrade the SAN Volume Controller cluster to software level 1.2.0. or later.
2. Migrate the data that resides on the partitioned units by performing one of the following actions:
 - Use the **svctask migratevdisk** SAN Volume Controller 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 the HP MA or EMA subsystem that corresponds with 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 units that are not partitioned so you can copy of all the data on MDisks that correspond to the partitioned units. Issue the **svctask rmmdisk** CLI command to delete the MDisks.
3. Rezone to utilize the extra ports on the HP MA or EMA subsystem.

Option 2: Retain partitioned units

Perform a concurrent code load to upgrade the SAN Volume Controller cluster to software 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 MA or EMA subsystem 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 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 MA or EMA subsystem that you want to connect to a SAN Volume Controller cluster that already has installed software level 1.2.0 or later.

You must initially zone in a single fibre-channel connection to one of the HP MA or EMA subsystems, 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 software level is already installed.

Supported models of HP MA and EMA subsystems

The SAN Volume Controller supports models of the HP MA and EMA subsystems.

See the following Web site for the latest supported models:

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

Attention: The SAN Volume Controller only supports configurations in which the HSG80 cache is enabled in writeback mode. Running with only a single controller results in a single point of data loss.

Supported firmware levels for HP MA and EMA subsystems

The HP MA and EMA subsystems must use a firmware level that is supported by the SAN Volume Controller.

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

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

Note: Concurrent upgrade of the subsystem firmware is not supported with the SAN Volume Controller.

Concurrent maintenance on the HP MA and EMA

Concurrent maintenance is the capability to perform I/O operations to an HP MA or EMA subsystem while simultaneously performing maintenance operations on it.

Note: HP MA and EMA maintenance documentation uses the phrase *rolling upgrade* in place of *concurrent maintenance*. See this documentation because in some instances you must reduce the level of I/O before you can perform the maintenance procedure.

The HP MA and EMA subsystems allow concurrent replacement of the following components:

- Drive
- EMU
- Blower
- Dual power supply (one unit can be removed and replaced. The fan speed increases when only one power supply unit is present.)

The following component is hot-pluggable, but concurrent maintenance of SAN Volume Controller I/O is not supported.

- Controller

The HP MA and EMA subsystems do not allow 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 MA and EMA

The Command Console configuration and service utility is the configuration interface for the HP MA and EMA subsystems.

The configuration and service utility can connect to the subsystem in the following ways:

- RS232 interface
- In band over fibre channel
- Over TCP/IP to a proxy agent, which then communicates with the subsystem in band over fibre channel.

In band

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 can therefore also access LUs that are visible to SAN Volume Controller nodes 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.

Sharing the HP MA or EMA between a host and a SAN Volume Controller

There are restrictions for sharing HP MA and EMA subsystems between a host and a SAN Volume Controller.

An HP MA or EMA can be shared between a host and a SAN Volume Controller under the following conditions:

- A host must not be connected to both a SAN Volume Controller and an HP MA or EMA subsystem at the same time.
- Target ports cannot be shared between a host and a SAN Volume Controller. Specifically, if an HSG80 port is in use by a SAN Volume Controller, it cannot be present in a switch zone that enables a host to access the port.
- LUs and RAID arrays cannot 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.

Switch zoning limitations for HP MA and EMA

There are limitations in switch zoning for the SAN Volume Controller and the HP MA and EMA subsystems.

Attention: The HP MA and EMA subsystems are supported with a single HSG80 controller or dual HSG80 controllers. Because the SAN Volume Controller only supports configurations in which the HSG80 cache is enabled in write-back mode, running with a single HSG80 controller results in a single point of data loss.

Switch zoning

For SAN Volume Controller clusters that have installed software version 1.1.1, a single fibre-channel port that is attached to the subsystem can be present in a switch zone that contains SAN Volume Controller fibre-channel ports, whether the HP MA or EMA subsystem uses one or two HSG80 controllers. This guarantees that the nodes in the cluster can access at most one port on the HSG80 controller.

For SAN Volume Controller clusters that have software version 1.2.0 or later installed, switches can be zoned so that HSG80 controller ports are in the switch zone that contains all of the ports for each SAN Volume Controller node.

Connecting to the SAN

Multiple ports from an HSG80 controller must be physically connected to the fibre-channel SAN to enable servicing of the HP MA or EMA subsystem. However, switch zoning must be used as described in this topic.

Note: If the HP Command Console is not able to access a fibre-channel port on each of the HSG80 controllers in a two-controller subsystem, there is a risk of an undetected single point of failure.

Quorum disks on HP MA and EMA subsystems

Managed disks (MDisks) that are presented by the HP MA or EMA are chosen by the SAN Volume Controller as quorum disks.

The SAN Volume Controller uses a logical unit (LU) that is presented by an HSG80 controller as a quorum disk. The quorum disk is used even if the connection is by a single port, although this is not recommended. If you are connecting the HP MA or EMA subsystem with a single fibre-channel port, ensure that you have another subsystem on which to put your quorum disk. You can use the **svctask setquorum** command-line interface (CLI) command to move quorum disks to another subsystem.

SAN Volume Controller clusters that are attached only to the HSG80 controllers are supported.

Advanced functions for HP MA and EMA

Some advanced functions of the HP MA and EMA are not supported by the SAN Volume Controller.

Advanced copy functions

Advanced copy functions for HP MA and EMA subsystems (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

HP MA and EMA support 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 nontransportable disk or storage set can be divided into a maximum of eight partitions.

The following restrictions apply to partitioning:

- Partitioned containers are fully supported if the HSG80 controller is connected to the SAN by a single port.
- Partitioned containers are not configured by the SAN Volume Controller if the HSG80 controller 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)

HP MA and EMA subsystems do not provide dynamic array expansion.

Write protection of LUNs

Write protection of LUNs is not supported for use with the SAN Volume Controller.

SAN Volume Controller advanced functions

Virtual disks (VDisks) that are created from managed disks (MDisks) that are presented by an HSG80 controller can be used in SAN Volume Controller FlashCopy mappings, SAN Volume Controller Metro Mirror relationships, and SAN Volume Controller Global Mirror relationships.

LU creation and deletion on the HP MA and EMA

Ensure you are familiar with the HSG80 controller container types for logical unit (LU) configuration.

Table 33 lists the valid container types.

Table 33. HSG80 controller container types for LU configuration

Container	Number of Members	Maximum Size
JBOD - non transportable Attention: A JBOD provides no redundancy at the physical disk drive level. A single disk failure can result in the loss of an entire managed disk group and its associated virtual disks.	1	disk size minus metadata
Mirrorset	2 - 6	smallest member
RAIDset	3 - 14	1.024 terabytes
Stripeset	2 - 24	1.024 terabytes
Striped Mirrorset	2 - 48	1.024 terabytes

Note: LUs can be created and deleted on an HSG80 controller while I/O operations are performed to other LUs. You do not need to restart the HP MA or EMA subsystem.

Configuring settings for the HP MA and EMA

The HP StorageWorks 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)
- Controller
- Port
- Logical unit
- Connection

Global settings for HP MA and EMA

Global settings apply across HP MA and EMA subsystems.

The following table lists the global settings for HP MA and EMA subsystems:

Table 34. HP MA and EMA global settings supported by the SAN Volume Controller

Option	HSG80 controller default setting	SAN Volume Controller required setting
DRIVE_ERROR_THRESHOLD	800	Default
FAILEDSET	Not defined	n/a

Controller settings for HP MA and EMA

Controller settings apply across one HSG80 controller.

Table 35 on page 319 describes the options that can be set by HSG80 controller command-line interface (CLI) commands for each HSG80 controller.

Table 35. HSG80 controller settings that are supported by the SAN Volume Controller

Option	HSG80 controller 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

Port settings for the HP MA and EMA

Port settings are configurable at the port level.

Restriction: Only one port per HSG80 pair can be used with the SAN Volume Controller.

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 36 lists the HSG80 controller port settings that the SAN Volume Controller supports:

Table 36. HSG80 controller port settings supported by the SAN Volume Controller

Option	HSG80 default setting	SAN Volume Controller required setting
PORT_1/2-AL-PA	71 or 72	Not applicable
PORT_1/2_TOPOLOGY	Not defined	FABRIC

Note: The HP MA and EMA subsystems support LUN masking that is configured with the **SET *unit number* ENABLE_ACCESS_PATH** command. When used with a SAN Volume Controller, the access path must be set to all ("**SET *unit number* ENABLE_ACCESS_PATH=ALL**") and all LUN masking must be handled exclusively by the SAN Volume Controller. You can use the **SHOW CONNECTIONS FULL** command to check access rights.

LU settings for HP MA and EMA

Logical unit (LU) settings are configurable at the LU level.

Table 37 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 37. HSG80 controller LU settings supported by the SAN Volume Controller

Option	HSG80 controller default setting	SAN Volume Controller required setting
TRANSFER_RATE_REQUESTED	20MHZ	Not applicable
TRANSPORTABLE/ NOTTRANSPORTABLE	NOTTRANSPORTABLE	NOTTRANSPORTABLE
ENABLE_ACCESS_PATH	ENABLE_ACCESS_PATH=ALL	ENABLE_ACCESS_PATH=ALL
DISABLE_ACCESS_PATH (See Note.)	NO DEFAULT	NO DEFAULT
IDENTIFIER/ NOIDENTIFIER	NOIDENTIFIER	Not applicable
MAX_READ_CACHE_SIZE	32	Not applicable
MAX_WRITE_CACHE_SIZE	32	64 or higher
MAX_CACHED_TRANSFER_SIZE	32	Not applicable
PREFERRED_PATH/ NOPREFERRED_PATH	NOPREFERRED_PATH is set	Not applicable
READ_CACHE/ NOREAD_CACHE	READ_CACHE	Not applicable
READAHEAD_CACHE/ NOREADAHEAD_CACHE	READAHEAD_CACHE	Not applicable
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 must always be overridden by using ENABLE_ACCESS_PATH=ALL on all connections to the SAN Volume Controller nodes.		

Connection settings for the HP MA and EMA

The HP MA and EMA subsystems provide options that are configurable at the connection level.

Table 38 on page 321 lists the default and required HSG80 controller connection settings:

Table 38. HSG80 connection default and required settings

Option	HSG80 controller default setting	HSG80 controller required setting
OPERATING_SYSTEM	Not defined	WINNT
RESERVATION_STYLE	CONNECTION_BASED	Not applicable
UNIT_OFFSET	0	0 or 199

Mapping and virtualization settings for HP MA and EMA

There are LUN mapping or masking and virtualization restrictions for HP MA and EMA subsystems that are in a SAN Volume Controller environment.

The HP StorageWorks configuration interface requires that you assign a unit number to each logical unit (LU) when it is defined. By default, the LUN is the unit number. It is possible for gaps to exist in the LUN range if the unit numbers that are used in the configuration commands are not contiguous. By default, each LUN is visible on all controller ports on both controllers.

LUN masking

The HP MA and EMA subsystems support the concept of connection names. A maximum of 96 connection names that contain the following parameters are supported:

- HOST_ID
- ADAPTER_ID
- CONTROLLER
- PORT
- REJECTED_HOST

Note: The SAN Volume Controller ports must not be in the REJECTED_HOSTS list. This list can be seen with the **SHOW CONNECTIONS FULL** command.

You cannot use LUN masking to restrict the initiator ports or the target ports that the SAN Volume Controller uses to access 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 that the SAN Volume Controller uses, but the preferred method for this is to use SAN zoning.

LU virtualization

The HP MA and EMA subsystems also provide 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 controller target ports and SAN Volume Controller initiator ports is not supported.

Configuring the HP StorageWorks EVA subsystem

This section provides information about configuring the HP StorageWorks Enterprise Virtual Array (EVA) subsystem for attachment to a SAN Volume Controller.

Supported models of the HP EVA

The SAN Volume Controller supports models of the HP EVA.

See the following Web site for the latest supported models:

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

Supported firmware levels for HP EVA

The SAN Volume Controller supports HP EVA.

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

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

Concurrent maintenance on the HP EVA

Concurrent maintenance is the capability to perform I/O operations to an HP EVA while simultaneously performing maintenance operations on it.

Important: All maintenance operations must be performed by an HP Field Engineer.

The SAN Volume Controller and HP EVA support concurrent hardware maintenance and firmware upgrade.

User interface on HP EVA

Ensure that you are familiar with the user interface that supports the HP EVA subsystem.

Storage Management Appliance

HP 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.

Sharing the HP EVA controller between a host and the SAN Volume Controller

The HP 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 HP EVA subsystem at the same time.
- LUs and RAID arrays must not be shared between a host and a SAN Volume Controller.

Switch zoning limitations for the HP 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.

Quorum disks on HP EVA

The SAN Volume Controller chooses managed disks (MDisks) that are presented by the HP EVA controllers as quorum disks.

Advanced functions for HP EVA

Advanced copy functions for HP 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 HP 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 VDIs 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 can be used for multiple VDIs. 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)

Logical unit creation and deletion on the HP EVA

EVA VDIs are created and deleted using the Command View EVA utility.

VDIs are formatted during creation. The time it takes to format the VDIs 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 HP 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 HP 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 HP 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.

Configuration interface for the HP EVA

The HP 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.

Configuring settings for the HP EVA

The HP 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

Global settings for the HP EVA

Global settings apply across an HP EVA subsystem.

Table 39 lists the subsystem options that you can access using the Command View EVA.

Table 39. HP EVA global settings supported by the SAN Volume Controller

Option	HP EVA default setting	SAN Volume Controller required setting
Console LUN ID	0	Any
Disk replacement delay	1	Any

Logical unit settings for the HP EVA

Logical unit (LU) settings are configurable at the LU level.

Table 40 on page 325 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 40. HP EVA LU settings supported by the SAN Volume Controller

Option	HP 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 HP EVA

Host settings are configurable for the HP EVA.

Table 41 lists the host options that can be accessed using the Command View EVA.

Table 41. HP EVA host settings supported by the SAN Volume Controller

Option	HP EVA Default Setting	SAN Volume Controller Required Setting
OS type	-	Windows
Direct eventing	Disabled	Disabled

Configuring HP MSA subsystems

This section provides information about configuring HP StorageWorks Modular Smart Array (MSA) subsystems for attachment to a SAN Volume Controller.

Supported models of the HP MSA subsystem

The SAN Volume Controller supports models of the HP MSA series of subsystems.

See the following Web site for the latest supported models: <http://www.ibm.com/storage/support/2145>

Supported firmware levels for the HP MSA

The HP MSA subsystem must use a firmware level that is supported by the SAN Volume Controller.

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

User interfaces on the HP MSA

Ensure that you are familiar with the user interface applications that are used by the HP MSA.

You can use the following configuration utilities with an HP MSA in a SAN Volume Controller environment:

- The CLI through an out-of-band configuration that is accessed through a host that is connected to the serial port of the HP MSA.

- The GUI through an in-band configuration that uses the HP ACU (Array Configuration Utility)

Notes:

1. If the HP ACU is installed in a configuration that HP does not support, some of its functionality might not be available.
2. If you use an in-band configuration, you must ensure that LUs that are used by the SAN Volume Controller cannot be accessed by a direct-attached host.

Logical unit creation, deletion and migration for HP MSA

You can create or delete logical units for the HP MSA. You must first read the storage configuration guidelines that are specified in the HP MSA documentation that is provided for this subsystem.

Creating arrays

For the HP MSA, an array is a collection of physical disks. Use the storage configuration guidelines for the SAN Volume Controller to create arrays on the HP MSA.

Creating logical drives

The following types of RAID arrays are supported:

- RAID 1+0
- RAID 1
- RAID 5
- RAID 6 (ADG)

RAID 0 is not supported because it does not provide failure protection.

All stripe sizes are supported; however, use a consistent stripe size for the HP MSA.

Use the following settings for logical drives:

- Set Max Boot to disabled
- Set Array Accelerator to enabled.

Presenting logical units to hosts

Set the Selective Storage Presentation (SSP), also known as ACL to enabled.

Use the following host profile settings:

```
Mode 0 = Peripheral Device LUN Addressing
Mode 1 = Asymmetric Failover
Mode 2 = Logical volumes connect as available on Backup Controller
Mode 3 = Product ID of 'MSA1000 Volume'
Mode 4 = Normal bad block handling
Mode 5 = Logout all initiators on TPRLO
Mode 6 = Fault management events not reported through Unit Attention
Mode 7 = Send FCP response info with SCSI status
Mode 8 = Do not send Unit Attention on failover
Mode 9 = SCSI inquiry revision field contains the actual version
Mode 10 = SCSI inquiry vendor field contains Compaq
Mode 11 = Power On Reset Unit Attention generated on FC Login or Logout
Mode 12 = Enforce Force Unit Access on Write
```

You can use the built-in Linux profile or Default profile to set the host profile settings. If you use the Default profile, you must issue the following Serial port CLI command to change the host profile settings:

```
change mode Default mode number
```

where *mode number* is the numeric value for the mode that you want change.

See the documentation that is provided for your HP MSA for additional information.

Important: You must use the Serial port CLI or the SSP to recheck the connection objects after the configuration is complete.

Migrating logical units from the HP MSA to the SAN Volume Controller

You can use the standard migration procedure with the following restrictions:

- You cannot share the HP MSA between a host and the SAN Volume Controller. You must migrate all hosts at the same time.
- The subsystem device driver (SDD) and securepath cannot coexist because they have different qllogic driver requirements.
- The qllogic driver that is supplied by HP must be removed and the driver that is supported by IBM must be installed.

Sharing the HP MSA between a host and the SAN Volume Controller

You must configure your environment so that only the SAN Volume Controller can access all logical units on the HP MSA. You can zone other hosts to communicate with the HP MSA for in-band configuration, but nothing else.

Concurrent maintenance on the HP MSA

Concurrent maintenance is the capability to perform I/O operations to an HP MSA while simultaneously performing maintenance operations on it.

You can perform nondisruptive maintenance procedures concurrently on the following components:

- HP MSA controller
- HP MSA controller cache
- Cache battery pack
- Variable speed blower

- Power supply
- Disk drive
- SFP transceiver

Quorum disks on the HP MSA

The SAN Volume Controller cannot use logical units (LUs) that are exported by the HP MSA as quorum disks.

Advanced functions for the HP MSA

The SAN Volume Controller Copy Service functions and RAID migration utilities are not supported for logical units (LUs) that are presented by the HP MSA.

Global settings for the HP MSA

Global settings apply across an HP MSA subsystem.

The following table lists the global settings for an HP MSA subsystem:

Option	Required setting
Expand Priority	All supported Note: Performance impact of high priority
Rebuild Priority	All supported Note: Performance impact of high priority
Array Accelerator	On Note: Set on all logical drives that are used by the SAN Volume Controller.
Read-Write cache ratio	All supported
Name of controller	Not important

Configuring NEC iStorage subsystems

This section provides information about configuring NEC iStorage subsystems for attachment to a SAN Volume Controller.

Supported firmware levels for the NEC iStorage

The NEC iStorage subsystem must use a firmware level that is supported by the SAN Volume Controller.

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

Logical unit creation and deletion for NEC iStorage

You can create or delete logical units for the NEC iStorage. See the storage configuration guidelines that are specified in the NEC iStorage documentation that is provided for this subsystem.

Platform type for NEC iStorage

You must set all logical units that the SAN Volume Controller accesses to platform type AX (AIX).

Access control methods for NEC iStorage

You can use access control to restrict access from hosts and SAN Volume Controller clusters. You do not need to use access control to allow a SAN Volume Controller cluster to use all of the defined logical units on the subsystem.

The following table lists the access control methods that are available:

Method	Description
Port Mode	Allows access to logical units that you want to define on a per storage controller port basis. SAN Volume Controller visibility (through switch zoning, physical cabling, etc.) must allow the SAN Volume Controller cluster to have the same access from all nodes and the accessible controller ports have been assigned the same set of logical units with the same logical unit number. This method of access control is not recommended for SAN Volume Controller connection.
WWN Mode	Allows access to logical units using the WWPN of each of the ports of an accessing host device. All WWPNs of all the SAN Volume Controller nodes in the same cluster must be added to the list of linked paths in the controller configuration. This becomes the list of host (SAN Volume Controller) ports for an LD Set or group of logical units. This method of access control allows sharing because different logical units can be accessed by other hosts.

Setting cache allocations for NEC iStorage

Cache allocations can be set manually; however, changes to the default settings can adversely effect performance and cause you to lose access to the subsystem.

Snapshot Volume and Link Volume for NEC iStorage

You cannot use Copy Services logical volumes with logical units that are assigned to the SAN Volume Controller.

Configuring NetApp FAS subsystems

This section provides information about configuring the Network Appliance (NetApp) Fibre-attached Storage (FAS) subsystems for attachment to a SAN Volume Controller. Models of the NetApp FAS subsystem are equivalent to the IBM System Storage N5000 series and the IBM System Storage N7000 series; therefore, the SAN Volume Controller also supports models of the IBM N5000 series and the IBM N7000 series.

The information in this section also applies to the supported models of the IBM N5000 series and the IBM N7000 series.

Supported models of the NetApp FAS subsystem

The SAN Volume Controller supports models of the NetApp FAS200, FAS900, FAS3000 and FAS6000 series of subsystems.

See the following Web site for the latest supported models:

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

Supported firmware levels for the NetApp FAS

The NetApp FAS must use a firmware level that is supported by the SAN Volume Controller.

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

User interfaces on the NetApp FAS

Ensure that you are familiar with the user interface applications that support the NetApp FAS.

See the documentation that is provided with your NetApp FAS subsystem for more information about the Web server and CLI.

Web server

You can manage, configure, and monitor the NetApp FAS through the FileView GUI.

CLI

You can access the command-line interface through a direct connection to the filer serial console port or by using the filer IP address to establish a telnet session.

Logical units and target ports on the NetApp FAS

For the NetApp FAS subsystems, a logical unit (LU) is a subdirectory in an internal file system.

LU's that are exported by the NetApp FAS report identification descriptors in the vital product data (VPD). The SAN Volume Controller uses the LUN-associated binary type-3 IEEE Registered Extended descriptor to identify the LU. For a NetApp LUN that is mapped to the SAN Volume Controller, set the LUN Protocol Type to Linux.

The NetApp FAS does not use LU groups so that all LUs are independent. The LU access model is active-active. Each LU has a preferred filer, but can be accessed from either filer. The preferred filer contains the preferred access ports for the LU. The SAN Volume Controller detects and uses this preference.

The NetApp FAS reports a different worldwide port name (WWPN) for each port and a single worldwide node name (WWNN).

Creating logical units on the NetApp FAS

To create logical units, you must identify a volume from which to create the logical unit and specify the amount of space to use.

Perform the following steps to create logical units:

1. Log on to the NetApp FAS.
2. Go to **Filer View** and authenticate.
3. Click **Volumes** and identify a volume to use to create an LU. A list of volumes is displayed.
4. Identify a volume that has sufficient free space for the LUN size you want to use.

5. Click **LUNs** on the left panel.
6. Click **Add** in the list.
7. Enter the following:
 - a. In the **Path** field, enter `/vol/volx/lun_name` where *volx* is the name of the volume identified above and *lun_name* is a generic name.
 - b. In the **LUN Type** field, enter Image.
 - c. Leave the **Description** field blank.
 - d. In the **Size** field, enter a LUN Size.
 - e. In the **Units** field, enter the LUN Size in units.
 - f. Select the **Space Reserved** box.

Note: If the Space Reserved box is not ticked and the file system is full, the LUN goes offline. The managed disk group also goes offline and you cannot access the virtual disks.

- g. Click **Add**.

Note: To check the LUN settings, go to the Manage LUNs section and click the LUN you want to view. Ensure that the Space Reserved setting is set.

Deleting logical units on the NetApp FAS

You can delete logical units.

Perform the following steps to delete logical units:

1. Log on to the NetApp FAS.
2. Go to **Filer View** and authenticate.
3. Click **LUNs** on the left panel.
4. Click **Manage**. A list of LUNs is displayed.
5. Click the LUN that you want to delete.
6. Click **Delete**.
7. Confirm the LUN that you want to delete.

Creating host objects for the NetApp FAS

You can create host objects.

Perform the following steps to create host objects:

1. Log on to the NetApp FAS.
2. Go to **Filer View** and authenticate.
3. Click **LUNs** on the left panel.
4. Click **Initiator Groups**.
5. Click **Add** in the list.
6. Enter the following:
 - a. In the **Group Name** field, enter the name of the initiator group or host.
 - b. In the **Type** list, select FCP.
 - c. In the **Operating System** field, select Linux.
 - d. In the **Initiators** field, enter the list of WWPNs of all the ports of the nodes in the cluster that are associated with the host.

Note: Delete the WWPNs that are displayed in the list and manually enter the list of SAN Volume Controller node ports. You must enter the ports of all nodes in the SAN Volume Controller cluster.

7. Click **Add**.

Presenting LUNs to hosts for NetApp FAS

You can present LUNs to hosts.

Perform the following steps to present LUNs to hosts:

1. Log on to the NetApp FAS.
2. Go to **Filer View** and authenticate.
3. Click **LUNs** on the left panel.
4. Click **Manage**. A list of LUNs is displayed.
5. Click the LUN that you want to map.
6. Click **Map LUN**.
7. Click **Add Groups to Map**.
8. Select the name of the host or initiator group from the list and click **Add**.

Notes:

- a. You can leave the LUN ID section blank. A LUN ID is assigned based on the information the controllers are currently presenting.
 - b. If you are re-mapping the LUN from one host to another, you can also select the **Unmap** box.
9. Click **Apply**.

Switch zoning limitations for the NetApp FAS

There are limitations in switch zoning for the SAN Volume Controller and NetApp FAS.

Fabric zoning

The SAN Volume Controller switch zone must include at least one target port from each filer to avoid a single point of failure.

Target port sharing

Target ports can be shared between the SAN Volume Controller and other hosts. However, you must define separate initiator groups (igroups) for the SAN Volume Controller initiator ports and the host ports.

Host splitting

A single host cannot be connected to both the SAN Volume Controller and the NetApp FAS to avoid the possibility of an interaction between multipathing drivers.

Controller splitting

You can connect other hosts directly to both the NetApp FAS and the SAN Volume Controller under the following conditions:

- Target ports are dedicated to each host or are in a different igroup than the SAN Volume Controller

- LUNs that are in the SAN Volume Controller igroup are not included in any other igroup

Concurrent maintenance on the NetApp FAS

Concurrent maintenance is the capability to perform I/O operations to a NetApp FAS while simultaneously performing maintenance operations on it.

The SAN Volume Controller supports concurrent maintenance on the NetApp FAS.

Quorum disks on the NetApp FAS

The SAN Volume Controller can use logical units (LUs) that are exported by the NetApp FAS as quorum disks.

Advanced functions for the NetApp FAS

The SAN Volume Controller copy and migration functions are supported for logical units (LUs) that are presented by the NetApp FAS.

Chapter 12. Maintaining the master console software

The topics in this section help you maintain the master console software on your system.

You can perform any of the following activities to maintain your master console software:

- Upgrade all of some of the master console components, including the SAN Volume Controller Console, using the master console installation program.
- Upgrade only the SAN Volume Controller Console component by using a downloaded installation wizard.
- Uninstall individual master console software components.

Upgrading the master console software

The topics in this section guide you through the upgrade process for the master console software using the master console installation program.

Prerequisites for upgrading the master console

This topic provides an overview of the prerequisites for upgrading the master console.

Before you upgrade the master console, you must meet the following prerequisites:

- Ensure that your system meets the master console hardware and software requirements provided in the *IBM System Storage SAN Volume Controller: Planning Guide*.
- You must be logged into the master console system using a user ID with administrative privileges.
- If you are upgrading the master console software from version 3.2 or an earlier version, you might be required to uninstall some components that were previously included with the master console. Table 42 provides a list of master console components that are not supported by version 4.2 and the prerequisite actions to take.

Table 42. Unsupported components and actions to take prior to upgrading

Component	Action
IBM Tivoli® Storage Area Network Manager (Tivoli SAN Manager) Agent	You must manually uninstall the Tivoli SAN Manager Agent. The “Uninstalling Tivoli SAN Manager Agent” on page 353 topic provides instructions for uninstalling this component.
IBM Tivoli SAN Manager	You must manually uninstall the Tivoli SAN Manager. The “Uninstalling Tivoli SAN Manager” on page 353 topic provides instructions for uninstalling this component.

Table 42. Unsupported components and actions to take prior to upgrading (continued)

Component	Action
DS4000 Storage Manager Client (FAStT Storage Manager Client)	Unless you currently use the DS4000 Storage Manager Client, you must uninstall it. "Uninstalling the DS4000 Storage Manager Client (FAStT Storage Manager Client)" on page 354 provides instructions for uninstalling this component.
IBM Connection Manager	The IBM Connection Manager is automatically uninstalled during the master console upgrade process. You do not need to take any action.
IBM Director	The upgrade process works best if you uninstall the IBM Director before you upgrade the master console software, but this is not mandatory.

Upgrading using the master console installation wizard

You can use the master console installation wizard to upgrade all the master console components to version 4.2.

Before you begin the upgrade process, ensure that you have performed the following actions:

- Logged in using a user ID with administrative privileges.
- Uninstalled any components that are not supported by the new version. For more information, see "Prerequisites for upgrading the master console" on page 335.

Perform the following steps to upgrade the master console software:

1. Insert the master console software installation CD in the CD drive.
2. Click **Start** → **Run** to open the Run window.
3. Enter *drive*:\setup.exe, where *drive* is the letter of the drive in which you inserted the CD. Click **OK**.

The following message is briefly displayed:

InstallShield(R) is preparing the InstallShield Wizard, which will guide you through the rest of the process. Please wait

Preparing Java (tm) Virtual Machine

You are then prompted to select the language to use for the installation wizard.

4. Select the language to use, and then click **OK**.
The Welcome panel is displayed.
5. Read the information on the Welcome panel, and then click **Next**.
The License Agreement panel is displayed.
6. Read the license agreement and perform one of the following actions:
 - Click **I accept the terms in the license agreement**, and then click **Next** to continue with the installation.
 - Click **I do not accept the terms of the license agreement**, and then click **Cancel** to exit the installation.

When you click **Next**, the wizard verifies that your system meets the hardware requirements for the installation.

Note: If your system does not meet hardware requirements, the wizard opens a panel that warns you about a decrease in the performance level if these requirements are not met. Click **OK** to close the warning panel.

A panel that shows you the destination directories for your master console software installation is displayed.

7. Click **Next**.

The installation wizard compares the list of components to be installed with the products that are already installed on your system. If the wizard finds that any master console components are already installed, it compares the versions and uses the following logic to determine which components to install:

- If the component is not installed or the installed version is an earlier version than the required version, the component is installed or upgraded by launching the component-specific installation program.
- If the component is installed at the same version level as the version that is installed by the installation wizard, the wizard does not install the component.
- If the installed component is a later version than the version to be installed, the installation wizard does not install the component, but a warning that indicates that this version is not tested with the master console is displayed. Click **OK** if you see this warning. Decide whether you want to continue the installation or exit the installation so that you can first remove the later version from your system. If you remove the later version, restart the master console installation wizard after the removal.
- If a component is incorrectly installed on the system, you are asked to continue with the installation by reinstalling the component with the component-specific installer. If the reinstallation does not succeed, you must exit the master console installation wizard, manually remove the product from your system, and restart the master console installation wizard.

After the wizard performs the comparison, the Product List panel is displayed. This panel provides the following information:

- Versions of existing master console components
- Required versions
- Actions to be done by the installation wizard or by you

8. From the List of products panel, click **Next** to continue upgrading the products.

The installation wizard launches the necessary component-specific installation programs.

Note: Upgrades to the master console software components are also available at the following Web site: <http://www.ibm.com/storage/support/2145>. Instructions for downloading and installing upgraded software packages are available at this site too.

9. Follow the instructions on the panels for each master console component that must be upgraded. Click **Finish** when all the components have been installed.
10. If a system restart is required, accept the prompt to complete the master console installation process.
11. Review the master console installation log (mclog.txt) to ensure that all components are properly installed. The log file is located in

installation_directory\logs, where *installation_directory* is the directory where the master console was installed. The default installation directory is C:\Program Files\IBM\MasterConsole.

Upgrading using the downloaded SAN Volume Controller Console installation program

You can download the SAN Volume Controller Console and upgrade or reinstall an existing SAN Volume Controller Console installation. The PuTTY installation program is also included in the downloaded files.

Upgrade or reinstallation modes

You can use one of two modes to upgrade or reinstall 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.

Overview of the upgrade or reinstallation process

The following list provides an overview of the upgrade or reinstallation tasks, as well as any configuration tasks that you must perform after you upgrade or reinstall the SAN Volume Controller Console:

1. Upgrade or reinstall the SAN Volume Controller Console either in graphical mode with the help of an installation wizard or in unattended (silent) mode. If errors were generated during the upgrade or reinstallation process, you must remove and reinstall the SAN Volume Controller Console.
2. 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
 - IBM Websphere Application Server V5 - SVC
3. Use a Web browser to access the SAN Volume Controller Console.
4. Identify the clusters that are to be managed by the SAN Volume Controller Console.

Using graphical mode for upgrading SAN Volume Controller Console and PuTTY

You can upgrade the SAN Volume Controller Console and PuTTY in graphical mode. You can also use this process for reinstalling an existing installation.

Before you can upgrade or reinstall the SAN Volume Controller Console and PuTTY in graphical mode, you must ensure that you have performed the following tasks:

- Ensure that your system meets the master console hardware and software requirements provided in the *IBM System Storage SAN Volume Controller: Planning Guide*.
- Download the SAN Volume Controller Console zip file from the following Web site:

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

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.

During the upgrade or reinstallation process, you use the IBM System Storage SAN Volume Controller Console Launchpad application. The Launchpad allows you to select from the following options:

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 SAN Volume Controller Console installation program 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 upgrade 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 System Storage 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 System Storage 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 System Storage 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

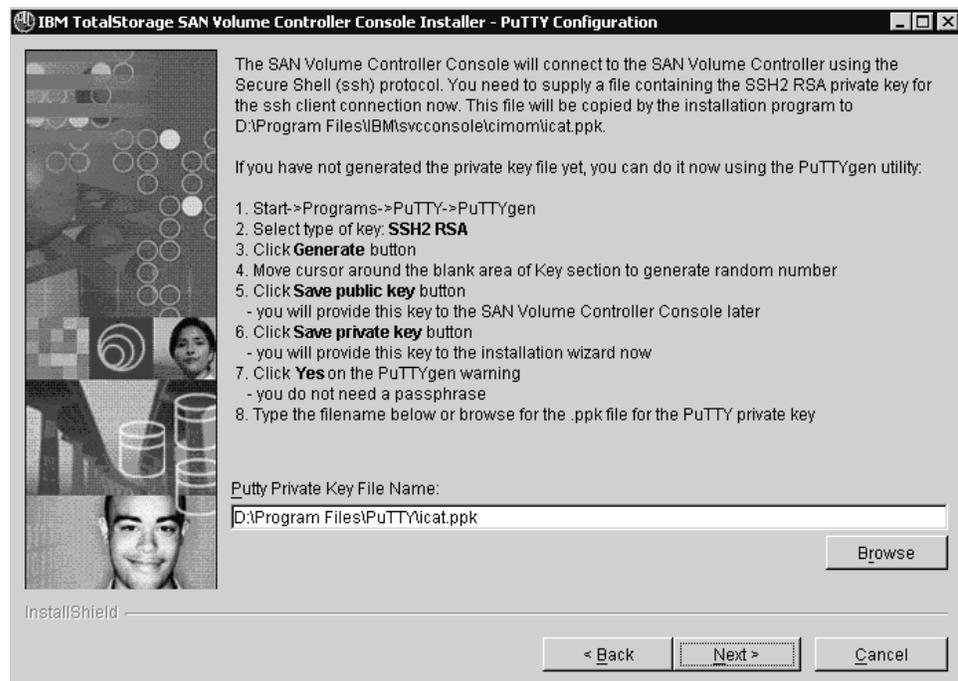
current configuration, the installation program skips the next steps and goes directly to the Installation Confirmation panel. If you do not preserve the current configuration, the Destination Directory panel is displayed.

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.

Notes:

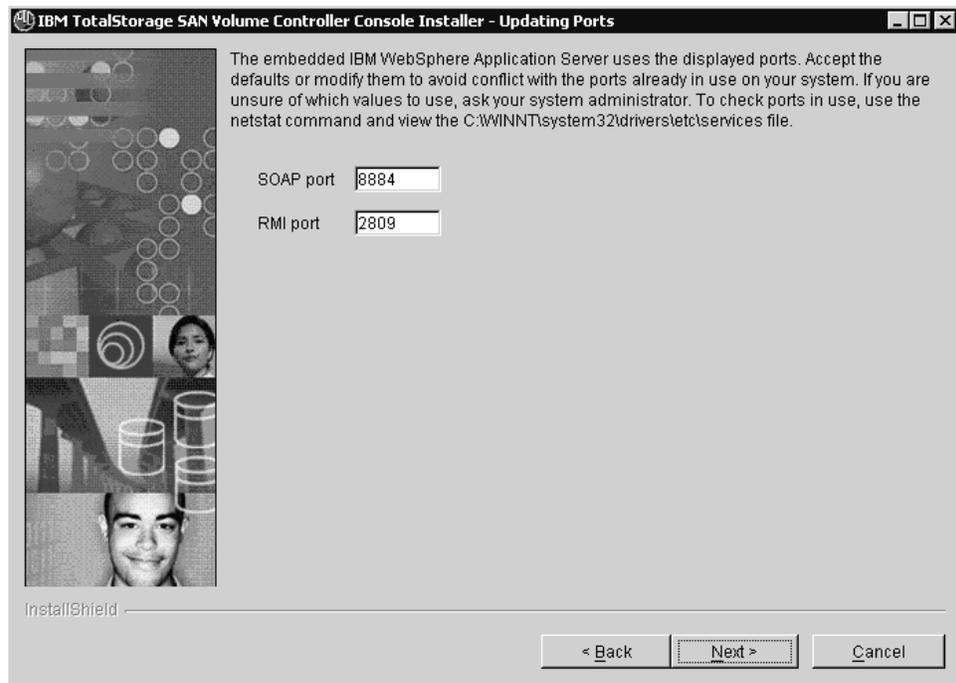
- The directory name, including the drive letter, can be a maximum of 44 characters.
- If the program detects insufficient space for the SAN Volume Controller Console installation in the chosen destination, an error message is displayed. You can free some space on the destination drive and then click **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.



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.



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
 - 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
 - IBM WebSphere Application Server V5 - SVC

Using unattended (silent) mode for upgrading SAN Volume Controller Console

The unattended (silent) mode upgrade option enables you to upgrade or reinstall SAN Volume Controller Console installations without having to interact with the installation wizard. The unattended mode can be helpful when you have to upgrade (or reinstall) multiple SAN Volume Controller Console installations.

Before you can upgrade or reinstall the SAN Volume Controller Console and PuTTY in graphical mode, you must ensure that you have performed the following tasks:

- Ensure that your system meets the master console hardware and software requirements that are provided in the *IBM System Storage SAN Volume Controller: Planning Guide*.
- You must know how to administer commands on a Windows 2000 Server or a Windows Server 2003 operating system.
- Download the SAN Volume Controller Console zip file from the following Web site:
<http://www.ibm.com/storage/support/2145>
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.
- If the secure shell (SSH) client software called PuTTY is not installed on your system, you must install PuTTY so that you can use the PuTTY key generator (PuTTYgen) to generate the SSH keys needed by the installation program. 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.

During the installation process, you 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 System Storage 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 (named `responsefile`), which 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">`

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 344 and customized in step 5 on page 345.

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, 2007 9:36:06 AM), This summary log is an overview of the
sequence of the installation of the IBM System Storage
SAN Volume Controller Console 4.2.0.546
(May 15, 2007 9:38:22 AM), IBM System Storage
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, 2007 9:38:28 AM), Copying Service Location Protocol Files ...
(May 15, 2007 9:38:29 AM), Service Location Protocol successfully installed
(May 15, 2007 9:38:29 AM), Copying CIM Object Manager Files ...
(May 15, 2007 9:39:26 AM), The PuTTY private key successfully copied
into file C:\Program Files\IBM\svconconsole\cimom\icat.ppk
(May 15, 2007 9:39:51 AM), The file setupCmdLine.bat successfully updated.
(May 15, 2007 9:39:51 AM), Compile MOF files started ...
(May 15, 2007 9:40:06 AM), MOF files successfully compiled.
(May 15, 2007 9:40:06 AM), Generate a certificate store started ...
(May 15, 2007 9:40:19 AM), Certificate store called truststore
successfully generated.
(May 15, 2007 9:40:20 AM), IBM CIM Object Manager successfully installed
(May 15, 2007 9:40:20 AM), Installing embedded version of IBM WebSphere
Application Server ...
(May 15, 2007 9:41:42 AM), Websphere Application Server - SVC
successfully installed.
(May 15, 2007 9:43:20 AM), Copying SAN Volume Controller Console Ear Files...
(May 15, 2007 9:46:11 AM), The ICAConsole application successfully installed.
(May 15, 2007 9:47:24 AM), The SVCConsole application successfully installed.
(May 15, 2007 9:48:06 AM), The help application successfully installed.
(May 15, 2007 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, 2007 9:48:27 AM), Command to be executed : net start cimomsrv
(May 15, 2007 9:48:49 AM), Command to be executed : net start
"IBMWAS5Service - SVC"
(May 15, 2007 9:50:15 AM), The following services started successfully:
Service Location Protocol
IBM CIM Object Manager
IBM WebSphere Application Server V5 - SVC
(May 15, 2007 9:50:15 AM), INSTSUCC: IBM
System Storage 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
 - 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
- 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** 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** 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.

Starting your initial SAN Volume Controller Console session

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 Systems Services Representative 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 Systems Services Representative 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 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:
- 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.
 - 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.
 - Select the *administrator* **Access Level** radio button.
 - 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 Console** 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.

Uninstalling master console software

To uninstall the master console software, you must remove the components separately.

Because of product dependencies, you must uninstall the software packages in a specific order. If you have any of the following components and plan to uninstall them, be sure to uninstall them in the following order.

Note: Some of the components that are listed were distributed in previous versions of the master console.

1. IBM Director
2. Tivoli SAN Manager Agent
3. Tivoli SAN Manager
4. DS4000 Storage Manager Client (FAStT Storage Manager Client)
5. Some of the listed components that were distributed in previous versions of the SAN Volume Controller Console.
6. PuTTY
7. Adobe Acrobat Reader
8. Master console

Note: When you remove the master console, you remove some documentation, support utilities, and icons. The documentation that is uninstalled with the master console option is located in `<destination_location>\Documents`, where `<destination_location>` is the location where the master console was installed on the system. The default location is `system_drive\Program Files\IBM\MasterConsole`.

Uninstalling IBM Director

This topic describes how to uninstall IBM Director from the master console using the Add/Remove Programs dialog panel.

This procedure assumes that you have accessed the Add/Remove Programs dialog panel.

Perform the following steps to uninstall IBM Director:

1. In the Add/Remove Programs panel, scroll to **IBM Director**, and click to select it.
2. Click **Change/Remove**.
3. Navigate through the uninstallation wizard, selecting the **Next** button of each window.
4. Wait for the program to be removed, and then click **Finish**.
5. If you are prompted to reboot the system, answer **yes** to reboot the system and complete the removal of the product.

Uninstalling Tivoli SAN Manager Agent

This topic describes how to uninstall Tivoli SAN Manager Agent from your master console hardware using the Add/Remove Programs dialog panel.

This procedure assumes that you have accessed the Add/Remove Programs dialog panel.

Perform the following steps to uninstall the Tivoli SAN Manager Agent:

1. In the Add/Remove Programs panel, scroll to **IBM Tivoli Storage Area Network Manager - Agent**, and click to select it.
2. Click **Change/Remove**.
3. Navigate through the uninstallation wizard, selecting the **Next** button of each window.
4. Wait for the program to be removed, and then click **Finish**.
5. If you are prompted to reboot the system, answer **yes** to reboot the system and complete the removal of the product.

Uninstalling Tivoli SAN Manager

This topic describes how to uninstall Tivoli SAN Manager from your master console hardware using the Add/Remove Programs dialog panel.

This procedure assumes that you have performed the following actions:

- Uninstalled the Tivoli SAN Manager Agent
- Accessed the Add/Remove Programs dialog panel

Perform the following steps to uninstall the Tivoli SAN Manager:

1. In the Add/Remove Programs panel, scroll to **IBM Tivoli Storage Area Network Manager - Manager**, and click to select it.
2. Click **Change/Remove**.
3. Navigate through the uninstallation wizard, selecting the **Next** button of each window.
4. Wait for the program to be removed, and then click **Finish**.

5. If you are prompted to reboot the system, answer **yes** to reboot the system and complete the removal of the product.
6. Remove the directory where the Tivoli SAN Manager and the Tivoli SAN Manager Agent were installed. By default, this directory is C:\Tivoli.

Uninstalling the DS4000 Storage Manager Client (FAStT Storage Manager Client)

This topic describes how to uninstall the DS4000 Storage Manager Client (FAStT Storage Manager Client) from the master console using the Add/Remove Programs dialog panel.

This procedure assumes that you have accessed the Add/Remove Programs dialog panel.

Perform the following steps to uninstall the DS4000 Storage Manager Client (FAStT Storage Manager Client):

1. In the Add/Remove Programs panel, scroll to the product name, and click to select it.
2. Click **Change/Remove**.
3. Navigate through the uninstallation wizard, selecting the **Next** button of each window.
4. Wait for the program to be removed, and then click **Finish**.
5. If you are prompted to reboot the system, answer **yes** to reboot the system and complete the removal of the product.

Uninstalling the SAN Volume Controller Console

This topic describes how to uninstall the SAN Volume Controller Console from your 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 System Storage 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 354. 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\svconconsole.

12. Empty your Windows Recycle Bin to reclaim the disk space that was made available during the removal process.

Uninstalling PuTTY

This topic describes how to uninstall PuTTY from the master console using the Add/Remove Programs dialog panel.

This procedure assumes that you have accessed the Add/Remove Programs dialog panel.

Perform the following steps to uninstall PuTTY:

1. In the Add/Remove Programs panel, scroll to **PuTTY**, and click to select it.
2. Click **Change/Remove**.
3. Navigate through the uninstallation wizard, selecting the **Next** button of each window.
4. Wait for the program to be removed, and then click **Finish**

Uninstalling the master console

This topic describes how to uninstall the master console.

This procedure assumes that you have opened the Microsoft Windows Add or Remove Programs dialog box.

Perform the following steps to uninstall the master console:

1. Find and select **IBM System Storage Master Console for SAN Volume Controller** in the Add or Remove Programs window.
2. Click **Remove** or **Change**.
3. Navigate through the uninstallation wizard, selecting the **Next** button of each window.
4. Wait for the program to be removed, and then click **Finish**
5. If you are prompted to reboot the system, answer **yes** to reboot the system and complete the removal of the product.

To complete the removal process, you can remove the directory where the master console was installed. The default is *system_drive*\Program Files\IBM\Master Console.

Changing the master console host name

You can change the master console host name anytime. When you change the host name, you must also be sure that other master console applications are updated to use the new name.

Perform the following steps to change the host name and to update the name in other master console applications:

1. Right-click on **My Computer** from the desktop.

2. Click **Properties**.
3. Click **Computer Name**.
4. Click **Change**.
5. Type the master console host name in the **Computer name** field.
6. Click **More**.
7. Type the full path information in the **Primary DNS suffix of this computer** field.
8. Click **OK** until you return to the desktop.
9. Click **Yes** to restart the master console system so that the change to the host name is applied.

Chapter 13. Troubleshooting the master console

These topics provide information that can help you troubleshoot and resolve problems with the master console.

In addition to troubleshooting on your own, you can also request an Assist On-site session with an IBM service representative.

Clearing the Microsoft Windows event logs

When you change the master console IP address or host name, you might create entries in the Microsoft Windows event logs.

Clear all three logs to ensure that these log entries do not cause confusion when you try to isolate problems.

The following procedure assumes that your Windows desktop is displayed.

Perform the following steps to clear the event logs:

1. Right-click **My Computer** and select **Manage**.
2. Expand **Event Viewer**.
3. Right-click **Application** and select **Clear All Events**. Click **No** when you are asked if you want to save the log before clearing.
4. Right-click **Security** and select **Clear All Events**. Click **No** when you are asked if you want to save the log before clearing.
5. Right-click **System** and select **Clear All Events**. Click **No** when you are asked if you want to save the log before clearing.
6. Close the computer management window.

Troubleshooting unexpected shutdowns of the SAN Volume Controller Console

If you are working with the SAN Volume Controller Console and you receive a You have signed off message before the SAN Volume Controller Console closes unexpectedly, use these instructions to help you troubleshoot the problem.

You can perform any of the following actions to troubleshoot an unexpected shutdown of the SAN Volume Controller Console:

- Open a new browser window and try to reconnect to the SAN Volume Controller Console. The logoff message is typically caused when an open session times out. This often happens if the browser window was left open from a previous session.
- Increase the allocation of memory for the SAN Volume Controller Console and then try again to reconnect to the application. An unexpected shutdown can occur because of a memory failure on the master console hardware, which causes the SAN Volume Controller Console to run with less memory than is required.
- Determine if the IP address of the master console hardware has changed since the last time that the system was restarted. If it has changed, restart the system to correct the problem.

Troubleshooting Microsoft Windows boot problems

Use this section to help you resolve Microsoft Windows boot problems on the master console system.

Perform the following actions to resolve Windows boot problems:

- If you cannot start the Windows system from the boot drive, try to start the master console system from the second disk drive (the mirrored disk).
- If you continue to have problems starting the system from either the boot disk drive or the second disk drive, you must replace the corrupted disk drive, and then mirror the boot drive again.

Note: After you set up mirroring, the hard disk drive on the system that runs the master console is actually a mirrored pair of hard disks. This strategy protects against loss of access to the master console due to a disk failure. This mirroring can help you start the master console system if the boot disk does not work. Whenever you replace one of the disks on your master console, you must make sure that you mirror the disks again.

Starting the master console hardware from the mirrored disk

During the Microsoft Windows boot process on the master console hardware, if Windows tries to start but fails with an Inaccessible Boot Device message on a blue screen, and another restart attempt does not solve the problem, the Windows boot code on the startup device might be corrupted.

The following instructions require that you use the administrator password for the power-on password when you restart the system. If the system is set up with an administrator password and you use a regular power-on password, you can see only a limited version of the **Configuration/Setup** menu.

Perform the following steps to resolve the problem:

1. Restart the master console system and watch the screen. When the Press F1 for Configuration/Setup message appears, press F1.
The main menu for Configuration/Setup Utility is displayed.
2. Select **Start Options** from the main menu.
3. Select **Start Sequence**.
4. Step down the sequence to the one that contains the hard disk.
5. Use the left and right cursor keys to select the other hard disk. For example, if the hard disk is set to 1, select 0. If the hard disk is set to 0, select 1.
6. Press Esc to exit each menu until the option to save and exit is displayed. Select **Yes** to save the changes and exit the Configuration/Setup Utility.
7. If the master console system starts, proceed with the steps for recovering from a master console disk failure. If the master console system does not start, contact your IBM service representative.

Replace a disk on the master console hardware

If one of the disk drives on the master console system fails, you might need to replace it with a new disk drive. The new drive must be the same capacity or larger than the drive being replaced.

Perform the following steps if one of the mirrored disk drives fails and must be replaced:

1. If you cannot determine which of the two disk drives has failed, restart the system with each disk drive to determine which drive is not functioning.
2. Right-click the **My Computer** icon on your desktop and select **Manage**.
3. Select **Disk Management** from the left navigation panel. The hard drives are displayed in the right panel.
4. If the failing disk drive is displayed, right-click the main volume of the drive and select **Break Mirror**.
5. Shut down the master console hardware and replace the failing disk drive using the procedures that are detailed in the documentation for your replacement hard drive. Ensure that the jumper settings for the new drive are the same as the jumper settings for the drive that is being replaced.

Notes: If the replacement drive has a master boot record (MBR), erase the MBR prior to using the replacement drive. However, if the master console computer fails to start because it cannot find the MBR, change the start sequence in the BIOS to the other hard drive.

6. Restart the computer.
7. Right-click **My Computer** on your desktop and select **Manage**.
8. Select **Disk Management**. The hard drives are displayed in the right panel.
9. If a disk drive is marked **Missing**, right-click the drive and select **Remove Disk**.
10. If a no entry sign is displayed on the new disk drive, right-click that disk drive and select **Write Signature**. This removes the no entry sign.
11. Right-click the new disk drive and select **Upgrade to Dynamic Disk**.
12. Right-click the volume that you want to mirror and select **Add Mirror**. The Add Mirror wizard is started.
13. Use the Add Mirror wizard to configure the second volume.
14. Ignore the window for making changes to the boot.ini file.

The status of both volumes, the existing drive, and the new drive changes to **Regenerating**. After a short period of time, the status shows the percentage of regeneration that has completed. When the regeneration completes, the status is displayed as **Healthy**.

Chapter 14. IBM System Storage 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 Console
- IBM System Storage hardware provider, known as the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software
- Microsoft Volume Shadow Copy Service

The IBM System Storage 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 System Storage 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 System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software must be completed in the correct sequence.

Before you begin, you must have experience with or knowledge of administering a Windows Server 2003 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 System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software.
4. Verify the installation.
5. Create a free pool of volumes and a reserved pool of volumes on the SAN Volume Controller.
6. Optionally, you reconfigure the services to change the configuration that you established during the installation.

System requirements for the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software

Ensure that your system satisfies the following requirements before you install the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software on a Windows Server 2003 operating system.

The following software is required:

- SAN Volume Controller Console software version 2.1.0 or later installed on the master console. You must install the SAN Volume Controller Console *before* you install the IBM System Storage hardware provider.
- SAN Volume Controller nodes with software version 2.1.0 or later installed with the FlashCopy feature enabled.
- IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software version 3.1 or later.
- Windows Server 2003 operating system. The following editions of Windows Server 2003 are supported:
 - Standard Server Edition 32-bit version
 - Enterprise Edition, 32-bit version
 - Standard Server Edition 64-bit version
 - Enterprise Edition, 64-bit version

Installing the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software

This section includes the steps to install the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software on a Windows server.

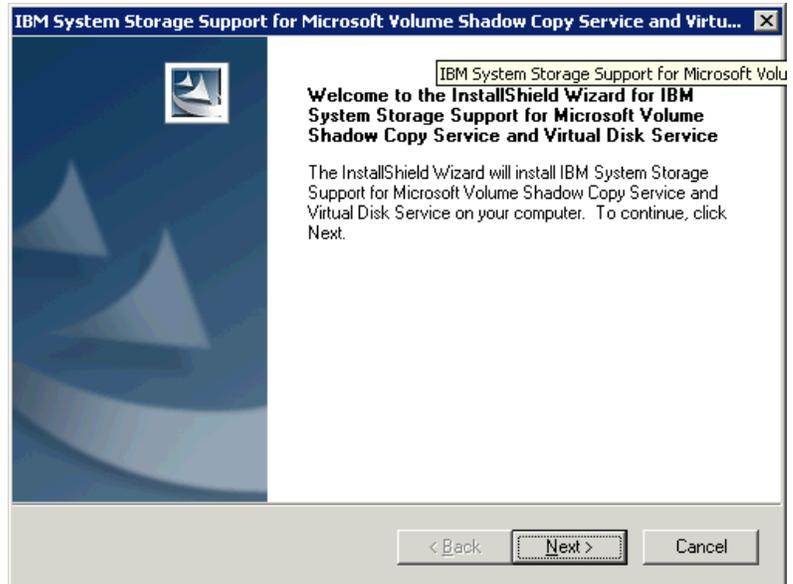
Important: You must install the SAN Volume Controller Console before you install the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software.

You must satisfy all of the prerequisites that are listed in the system requirements section before starting the installation.

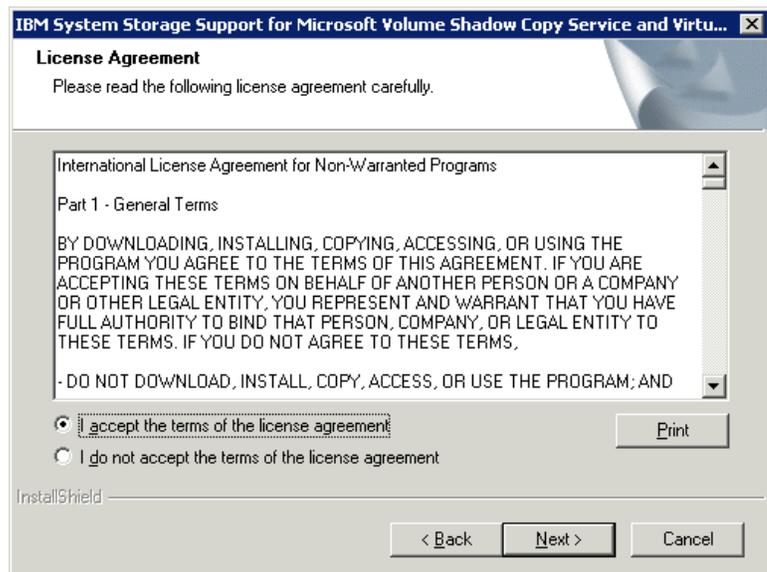
Perform the following steps to install the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software on the Windows server:

1. Log on to Windows as an administrator.

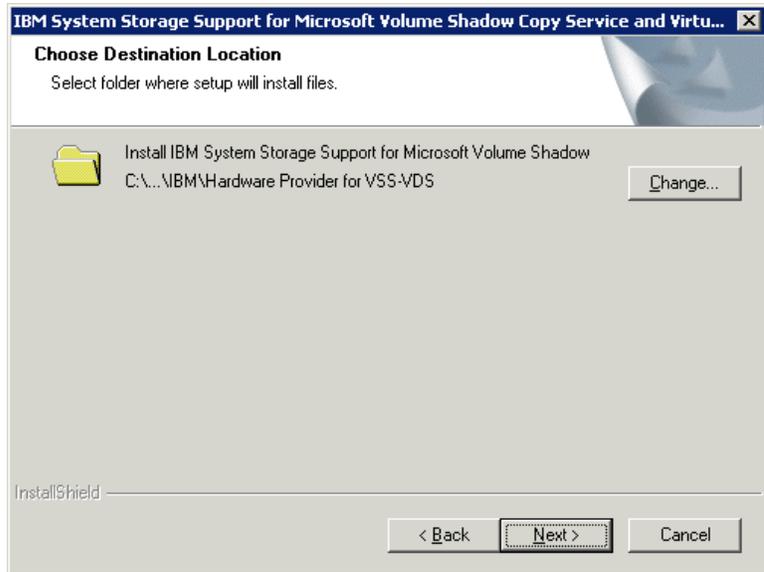
2. Download the IBM VSS Host Installation Package file from the following Web site:
<http://www.ibm.com/storage/support/2145>
3. Double click on the name of the file that you downloaded in step 2 to start the installation process. The Welcome panel is displayed.



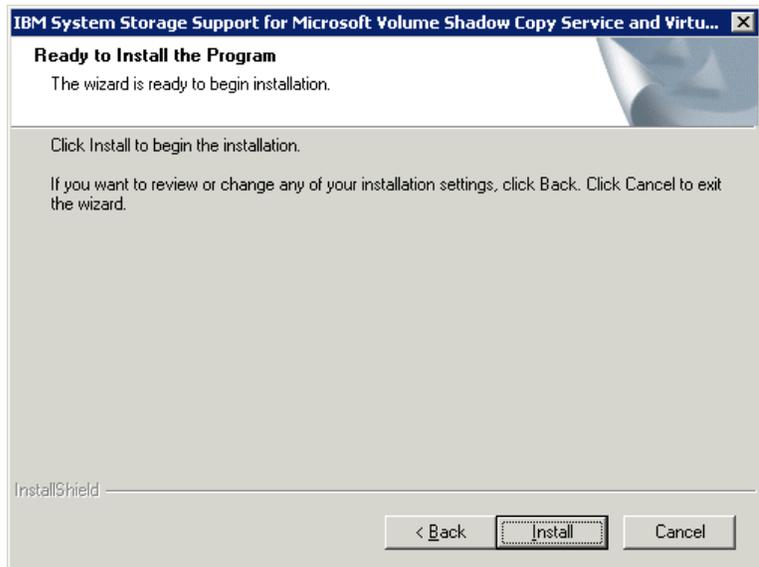
4. Click **Next** to continue. The License Agreement panel is displayed. You can click **Cancel** at any time to exit the installation. To move back to previous screens while using the wizard, click **Back**.



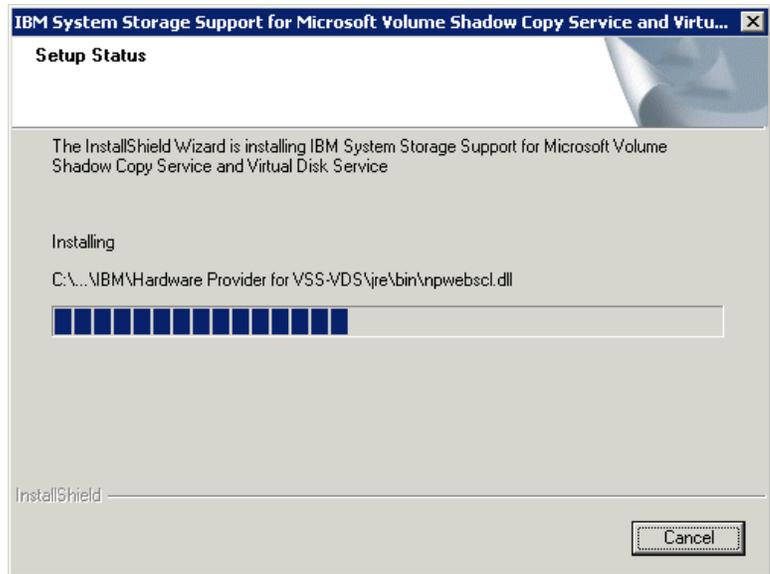
5. Read the license agreement information. Select whether you accept the terms of the license agreement, and click **Next**. If you do not accept, you cannot continue with the installation. The Choose Destination Location panel is displayed.



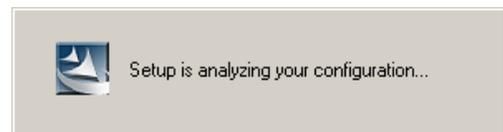
6. Click **Next** to accept the default directory where the setup program will install the files, or click **Change** to select a different directory. Click **Next**. The Ready to Install the Program panel is displayed.



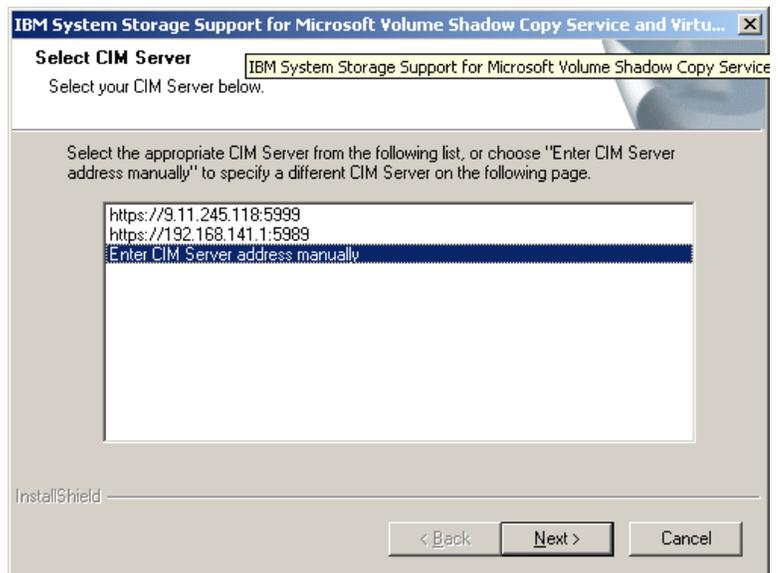
7. Click **Install** to begin the installation. To exit the wizard and end the installation, click **Cancel**. The Setup Status panel is displayed.



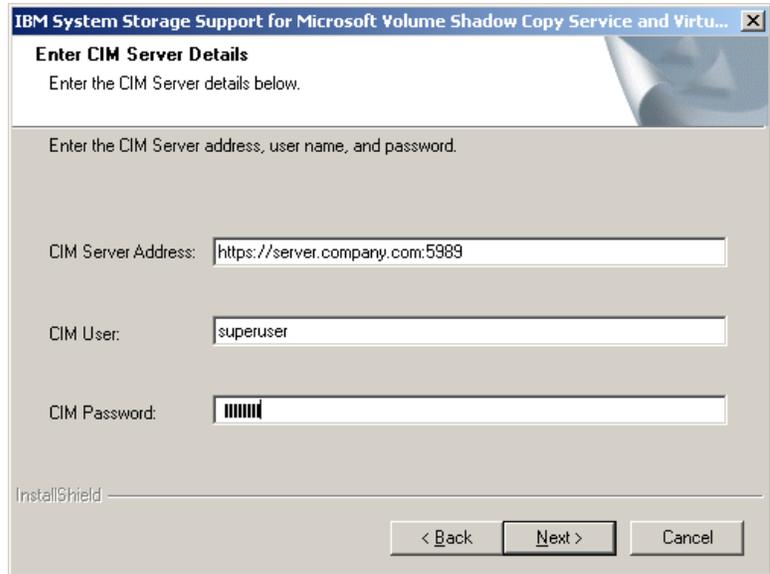
The program setup verifies your configuration.



The Select CIM Server panel is displayed.



8. Select the required CIM server, or select **Enter the CIM Server address manually**, and click **Next**. The Enter CIM Server Details panel is displayed.

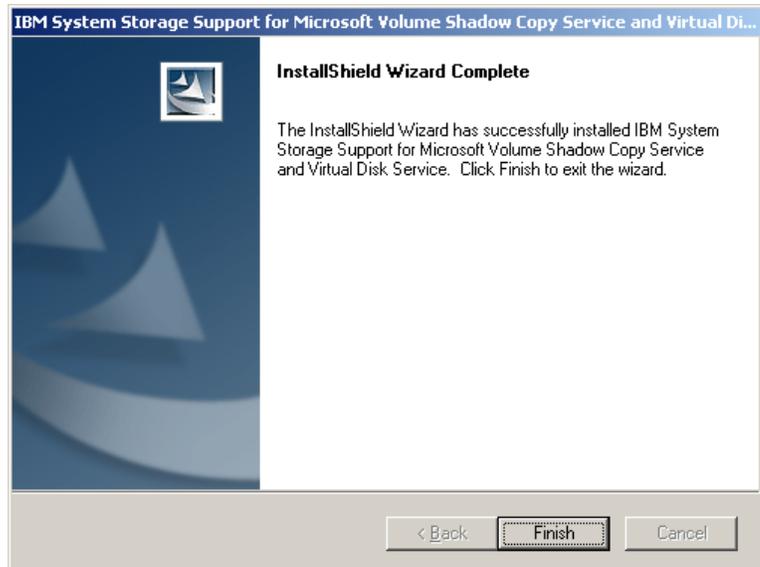


9. Enter the following information in the fields:
 - In the **CIM Server Address** field, type the name of the server where the SAN Volume Controller Console is installed.
 - In the **CIM User** field, type the user name that the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software will use to gain access to the server where the SAN Volume Controller Console is installed. For example, enter the name superuser.
 - In the **CIM Password** field, type the password for the user name that the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software will use to gain access to the SAN Volume Controller Console and click **Next**.

Notes:

- a. If these settings change after installation, you can use the `ibmvcfg.exe` tool to update Microsoft Volume Shadow Copy and Virtual Disk Services software with the new settings.
- b. If you do not have the CIM agent server, port, or user information, contact your CIM agent administrator.

The InstallShield Wizard Complete panel is displayed.



10. Click **Finish**. If necessary, the InstallShield Wizard prompts you to restart the system.

Creating the free and reserved pools of volumes

The IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software 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 System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software 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 `50000000000000001` (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 on page 369 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 System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software with the WWPNs.

Verifying the installation

This task verifies that the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software 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 System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software 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 System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software is listed as a provider.
5. Use the **ibmvcfg listvols** command to test the connection to the master console.

If you are able to successfully perform all of these verification tasks, the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software was successfully installed on the Windows server.

Changing the configuration parameters

You can change the parameters that you defined when you installed the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software. You must use the **ibmvcfg.exe** utility to change the parameters.

Table 43 describes the configuration commands that are provided by the **ibmvcfg.exe** utility.

Table 43. Configuration commands

Command	Description	Example
ibmvcfg showcfg	Lists the current settings.	ibmvcfg showcfg
ibmvcfg set username <username>	Sets the user name to access the SAN Volume Controller Console.	ibmvcfg set username johnny
ibmvcfg set password <password>	Sets the password of the user name that will access the SAN Volume Controller Console.	ibmvcfg set password mypassword

Table 43. Configuration commands (continued)

Command	Description	Example
ibmvcfg set targetSVC <ipaddress>	Specifies the IP address of the SAN Volume Controller on which the VDIs are located when VDIs are moved to and from the free pool with the ibmvcfg add and ibmvcfg rem commands. The IP address is overridden if you use the -s flag with the ibmvcfg add and ibmvcfg rem commands.	set targetSVC 64.157.185.191
set backgroundCopy	Sets the background copy rate for FlashCopy.	set backgroundCopy 80
ibmvcfg set usingSSL	Specifies whether to use Secure Sockets Layer protocol to connect to the SAN Volume Controller Console.	ibmvcfg set usingSSL yes
ibmvcfg set cimomPort <portnum>	Specifies the SAN Volume Controller 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 SAN Volume Controller 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 IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software 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 44 describes the `ibmvfcg.exe` commands for adding or removing volumes from the free pool of volumes.

Table 44. Pool management commands

Command	Description	Example
<code>ibmvfcg listvols</code>	Lists all virtual disks (VDisks), including information about size, location, and VDisk to host mappings.	<code>ibmvfcg listvols</code>
<code>ibmvfcg listvols all</code>	Lists all VDIs, including information about size, location, and VDisk to host mappings.	<code>ibmvfcg listvols all</code>
<code>ibmvfcg listvols free</code>	Lists the volumes that are currently in the free pool.	<code>ibmvfcg listvols free</code>
<code>ibmvfcg listvols unassigned</code>	Lists the volumes that are currently not mapped to any hosts.	<code>ibmvfcg listvols unassigned</code>
<code>ibmvfcg add -s ipaddress</code>	Adds one or more volumes to the free pool of volumes. Use the <code>-s</code> parameter to specify the IP address of the SAN Volume Controller where the VDIs are located. The <code>-s</code> parameter overrides the default IP address that is set with the <code>ibmvfcg set targetSVC</code> command.	<code>ibmvfcg add vdisk12</code> <code>ibmvfcg add 600507</code> <code>68018700035000000</code> <code>0000000BA</code> <code>-s 66.150.210.141</code>
<code>ibmvfcg rem -s ipaddress</code>	Removes one or more volumes from the free pool of volumes. Use the <code>-s</code> parameter to specify the IP address of the SAN Volume Controller where the VDIs are located. The <code>-s</code> parameter overrides the default IP address that is set with the <code>ibmvfcg set targetSVC</code> command.	<code>ibmvfcg rem vdisk12</code> <code>ibmvfcg rem 600507</code> <code>68018700035000000</code> <code>0000000BA</code> <code>-s 66.150.210.141</code>

Error codes

The IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software 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 System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software 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 System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software is installed.

Table 45 lists the errors messages that are reported by the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software.

Table 45. Error messages for the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software

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
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	<code>ibmGetLunInfo</code> failed in <code>AreLunsSupported</code> .	ERR_ARELUNSSUPPORTED_ IBMGETLUNINFO
1800	<code>ibmGetLunInfo</code> failed in <code>FillLunInfo</code> .	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

Table 45. Error messages for the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software (continued)

Code	Message	Symbolic name
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
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 System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software

You must use Windows to uninstall the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software from the Windows server.

Perform the following steps to uninstall the software:

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 System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software** 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 System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software is no longer installed on the Windows server.

Appendix A. Replacing nodes nondisruptively

You can replace SAN Volume Controller 2145-4F2, SAN Volume Controller 2145-8F2, or SAN Volume Controller 2145-8F4 nodes with SAN Volume Controller 2145-8G4 nodes without disrupting your SAN environment. This task does not disrupt your environment because the replacement (new) node uses the same worldwide node name (WWNN) as the node you are replacing.

This task assumes that the following conditions exist:

- The cluster software is at 4.2.0 or higher
- All nodes that are configured in the cluster are present
- All errors in the cluster error log are fixed
- All managed disks (MDisks) are online
- You have a 2145 uninterruptible power supply-1U (2145 UPS-1U) unit for each new SAN Volume Controller 2145-8G4 node.

Perform the following steps to replace nodes:

1. Perform the following steps to record the WWNN of the node that you want to replace:
 - a. Issue the following command from the command-line interface (CLI):

```
svcinfo lsnode -delim : node_name or node_id
```

Where *node_name* or *node_id* is the name or ID of the node for which you want to determine the WWNN.
 - b. Record the WWNN of the node that you want to replace.
2. Issue the following CLI command to delete this node from the cluster and I/O group:

```
svctask rmnode node_name or node_id
```

Where *node_name* or *node_id* is the name or ID of the node that you want to delete.

Notes:

- a. The node is not deleted until the SAN Volume Controller cache is destaged to disk. During this time, the partner node in the I/O group transitions to write through mode.
 - b. You can use the CLI to verify that the deletion process has completed.
3. Issue the following CLI command to ensure that the node is no longer a member of the cluster:

```
svcinfo lsnode node_name or node_id
```

Where *node_name* or *node_id* is the name or ID of the node. The node is not listed in the command output.
 4. Perform the following steps to change the WWNN of the node that you just deleted from the cluster to FFFF:
 - a. From the front panel of the node, press the down button and then use the right and left navigation buttons to display the Node Status menu.
 - b. Press and hold the down button and press the select button. The WWNN of the node is displayed.

- c. Press the down and select buttons to start the WWNN edit mode. The first character of the WWNN is highlighted.
- d. Press the up or down button to increment or decrement the character that is displayed.

Note: The characters wrap F to 0 or 0 to F.

- e. Press the left navigation button to move to the next field or the right navigation button to return to the previous field and repeat step 4d for each field. At the end of this step, the characters that are displayed must be FFFFF.
 - f. Press the select button to retain the characters that you have updated and return to the WWNN Selection menu.
 - g. Press the select button to apply the characters as the new WWNN for the node.
5. Power off and remove the node from the rack.

Tip: Record and mark the order of the fibre-channel cables so that you can use the same order for the replacement node.

6. Install the replacement node in the rack and connect the 2145 UPS-1U cables.

Important: Do not connect the fibre-channel cables during this step.

7. Power-on the node.
8. Perform the following steps to change the WWNN of the replacement node to match the WWNN that you recorded in step 1 on page 375:
 - a. From the front panel of the node, press the down button and then use the right and left navigation buttons to display the Node Status menu.
 - b. Press and hold the down button and press the select button. The WWNN of the node is displayed.
 - c. Press the down and select buttons to start the WWNN edit mode. The first character of the WWNN is highlighted.
 - d. Press the up or down button to increment or decrement the character that is displayed.

Note: The characters wrap F to 0 or 0 to F.

- e. Press the left navigation button to move to the next field or the right navigation button to return to the previous field and repeat step 8d for each field. At the end of this step, the characters that are displayed must be the same as the WWNN you recorded in step 1 on page 375.
 - f. Press the select button to retain the characters that you have updated and return to the WWNN Selection menu.
 - g. Press the select button to apply the characters as the new WWNN for the node.
9. Connect the fibre-channel cables to the node.
 10. Issue the following CLI command to verify that the last five characters of the WWNN are correct:
svcinfolsnodecandidate

Important: If the WWNN is not correct, you must repeat step 8.

11. Add the node to the cluster and I/O group. See the **svctask addnode** CLI command in the *IBM System Storage SAN Volume Controller: Command-Line Interface User's Guide*.

Important:

- a. Both nodes in the I/O group cache data; however, the cache sizes are asymmetric if the remaining partner node in the I/O group is a SAN Volume Controller 2145-4F2 node. The replacement node is limited by the cache size of the partner node in the I/O group. Therefore, the replacement node does not utilize the full size of its cache.
 - b. You do not have to reconfigure the host multipathing device drivers because the replacement node uses the same WWNN as the previous node. The multipathing device drivers should detect the recovery of paths that are available to the replacement node.
 - c. The host multipathing device drivers take approximately 30 minutes to recover the paths.
12. See the documentation that is provided with your multipathing device driver for information on how to query paths to ensure that all paths have been recovered before proceeding to the next step.
 13. Repeat steps 1 on page 375 to 12 for each node that you want to replace.

Note: If you upgrade both nodes in the I/O group to SAN Volume Controller 2145-8G4 nodes, the cache sizes are symmetric and the full 8 GB of cache is utilized.

Appendix B. Replacing nodes disruptively (rezoning the SAN)

You can replace SAN Volume Controller 2145-4F2, SAN Volume Controller 2145-8F2, or SAN Volume Controller 2145-8F4 nodes with SAN Volume Controller 2145-8G4 nodes. This task disrupts your environment because you must rezone your SAN and the host multipathing device drivers must discover new paths. Access to virtual disks (VDisks) is lost during this task.

This task assumes that the following conditions exist:

- The cluster software is at 4.2.0 or higher
- All nodes that are configured in the cluster are present
- All errors in the cluster error log are fixed
- All managed disks (MDisks) are online
- You have a 2145 uninterruptible power supply-1U (2145 UPS-1U) unit for each new SAN Volume Controller 2145-8G4 node.

Perform the following steps to replace nodes:

1. Quiesce all I/O from the hosts that access the I/O group of the node that you are replacing.
2. Delete the node that you want to replace from the cluster and I/O group.

Notes:

- a. The node is not deleted until the SAN Volume Controller cache is destaged to disk. During this time, the partner node in the I/O group transitions to write through mode.
 - b. You can use the command-line interface (CLI) or the SAN Volume Controller Console to verify that the deletion process has completed.
3. Ensure that the node is no longer a member of the cluster.
 4. Power-off the node and remove it from the rack.
 5. Install the replacement (new) node in the rack and connect the uninterruptible power supply (UPS) cables and the fibre-channel cables.
 6. Power-on the node.
 7. Rezone your switch zones to remove the ports of the node that you are replacing from the host and storage zones. Replace these ports with the ports of the replacement node.
 8. Add the replacement node to the cluster and I/O group.

Important: Both nodes in the I/O group cache data; however, the cache sizes are asymmetric. The replacement node is limited by the cache size of the partner node in the I/O group. Therefore, the replacement node does not utilize the full size of its cache.

9. From each host, issue a rescan of the multipathing software to discover the new paths to VDisks.

Notes:

- a. If your system is inactive, you can perform this step after you have replaced all nodes in the cluster.

- b. The host multipathing device drivers take approximately 30 minutes to recover the paths.
10. See the documentation that is provided with your multipathing device driver for information on how to query paths to ensure that all paths have been recovered before proceeding to the next step.
11. Repeat steps 1 on page 379 to 10 for the partner node in the I/O group.

Note: After you have upgraded both nodes in the I/O group, the cache sizes are symmetric and the full 8 GB of cache is utilized.

12. Repeat steps 1 on page 379 to 11 for each node in the cluster that you want to replace.
13. Resume host I/O.

Appendix C. Replacing nodes disruptively (moving VDIs to new I/O group)

You can replace SAN Volume Controller 2145-4F2, SAN Volume Controller 2145-8F2, or SAN Volume Controller 2145-8F4 nodes with SAN Volume Controller 2145-8G4 nodes. This task disrupts your environment because you must move virtual disks (VDIs) from the I/O group of the nodes that you are replacing to a new I/O group.

This task assumes the following:

- The cluster software is at 4.2.0 or higher
- Your cluster contains six or less nodes
- All nodes that are configured in the cluster are present
- All errors in the cluster error log are fixed
- All managed disks (MDisks) are online
- You have a 2145 uninterruptible power supply-1U (2145 UPS-1U) unit for each new SAN Volume Controller 2145-8G4 node.

Perform the following steps to replace nodes:

1. Quiesce all I/O from the hosts that access the I/O groups of the nodes that you are replacing.
2. Zone the ports from the replacement (new) nodes.
3. Add two replacement nodes to the cluster to create a new I/O group.
4. Move all of the VDIs from the I/O group of the nodes you are replacing to the new I/O group.
5. From each host, issue a rescan of the multipathing software to discover the new paths to VDIs. The host multipathing device drivers take approximately 30 minutes to recover the paths.
6. See the documentation that is provided with your multipathing device driver for information on how to query paths to ensure that all paths have been recovered before proceeding to the next step.
7. Delete the nodes that you are replacing from the cluster and remove the ports from the switch zones.
8. Repeat steps 1 to 7 for each node in the cluster that you want to replace.

Appendix D. Error Codes

Error codes provide a unique entry to service procedures. Each error code has an error ID that uniquely identifies the condition that caused the error.

Error IDs are recorded in the error log. When the number of error IDs of a specific type for a specific resource exceeds a predetermined threshold, an SNMP trap is raised and an e-mail is sent. When the SNMP traps are received, the SNMP type is used by the management tools to control how the trap is processed. The SNMP type is used by the Call Home e-mail service to decide the recipients, the title, and the contents of the e-mail. The following SNMP types are possible:

Error This type identifies unexpected conditions that might be the result of a system failure. If configured, this type causes an SNMP trap to be sent to the monitoring application. An e-mail can also be sent to the IBM Support Center and the system administrator.

Warning

This type identifies unexpected conditions that might be experienced during user operations. These conditions can result from device errors or incorrect user actions. If configured, this type causes an SNMP trap to be sent to the monitoring application. An e-mail can also be sent to the system administrator.

Information

This type identifies conditions where a user might want to be notified of the completion of an operation. If configured, this type causes an SNMP trap to be sent to the monitoring application. An e-mail can also be sent to the system administrator.

Table 46 lists the error codes and corresponding error IDs.

Table 46. Error codes

Error ID	SNMP Type	Condition	Error Code
009020	E	An automatic cluster recovery has started. All configuration commands are blocked.	1001
009040	E	The error log is full.	1002
009052	E	The following causes are possible: <ul style="list-style-type: none">• The node is missing• The node is no longer a functional member of the cluster• One or more nodes are not available	1195
009100	W	The software install process has failed.	2010
009101	W	The software upgrade package delivery has failed.	2010
009150	W	Unable to connect to the SMTP (e-mail) server	2600
009151	W	Unable to send mail through the SMTP (e-mail) server	2601
010002	E	The node ran out of base event sources. As a result, the node has stopped and exited the cluster.	2030
010003	E	The number of device logins has reduced.	1630

Table 46. Error codes (continued)

Error ID	SNMP Type	Condition	Error Code
010006	E	The following causes are possible: <ul style="list-style-type: none"> • There is an attempt to access beyond the end of disk • The managed disk does not exist 	2030
010008	E	The block size is invalid, the capacity or LUN identity has changed during the managed disk initialization.	1660
010010	E	The managed disk is excluded because of excessive errors.	1310
010011	E	The remote port is excluded for a managed disk and node.	1220
010012	E	The local port is excluded.	1210
010013	E	The login is excluded.	1230
010017	E	A timeout has occurred as a result of excessive processing time.	1340
010018	E	An error recovery procedure has occurred.	1370
010019	E	A managed disk I/O error has occurred.	1310
010020	E	The managed disk error count threshold has exceeded.	1310
010021	E	There are too many devices presented to the cluster.	1200
010022	E	There are too many managed disks presented to the cluster.	1200
010023	E	The are too many LUNs presented to a node.	1200
010025	W	A disk I/O medium error has occurred.	1320
010026	E	There are no managed disks that can be used as a quorum disk.	1330
010027	E	The quorum disk is not available.	1335
010028	E	The controller configuration is not correct.	1625
010029	E	A login transport fault has occurred.	1360
010030	E	The node or controller reported the following: <ul style="list-style-type: none"> • Sense • Key • Code • Qualifier 	1370
010031	E	The managed disk is offline.	1623
010040	E	The controller subsystem device is only connected to the node through a single initiator port.	1627
010041	E	The controller subsystem device is only connected to the node through a single target port.	1627
010042	E	The controller subsystem device is only connected to the cluster nodes through a single target port.	1627
010043	E	The controller subsystem device is only connected to the cluster nodes through half of the expected target ports.	1627
010044	E	The controller subsystem device has disconnected all target ports to the cluster nodes.	1627
020001	E	There are too many medium errors on the managed disk.	1610
020002	E	A managed disk group is offline.	1620

Table 46. Error codes (continued)

Error ID	SNMP Type	Condition	Error Code
020003	W	There are insufficient virtual extents.	2030
030000	W	The trigger prepare command has failed because of a cache flush failure.	1900
030010	W	The mapping is stopped because of the error that is indicated in the data.	1910
050010	W	A Metro Mirror or Global Mirror relationship has stopped because of a persistent I/O error.	1920
050020	W	A Metro Mirror or Global Mirror relationship has stopped because of an error that is not a persistent I/O error.	1720
072001	E	A system board hardware failure has occurred. This error applies to only the SAN Volume Controller 2145-4F2 model.	1020
072004	E	A CMOS battery failure has occurred. This error applies to the SAN Volume Controller 2145-4F2, the SAN Volume Controller 2145-8F2 and the SAN Volume Controller 2145-8F4 models.	1670
072005	E	A CMOS battery failure has occurred. This error applies to only the SAN Volume Controller 2145-8G4 model.	1670
072101	E	The processor is missing. This error applies to both the SAN Volume Controller 2145-8F2 and the SAN Volume Controller 2145-8F4 models.	1025
072102	E	The processor is missing. This error applies to only the SAN Volume Controller 2145-8G4 model.	1025
073001	E	The fibre-channel adapter card has detected an incorrect number of fibre-channel adapters. This error applies to only the SAN Volume Controller 2145-4F2 model.	1010
073002	E	The fibre-channel adapter has failed. This error applies to only the SAN Volume Controller 2145-4F2 model.	1050
073003	E	The fibre-channel ports are not operational.	1060
073004	E	The fibre-channel adapter has detected a PCI bus error. This error applies to only the SAN Volume Controller 2145-4F2 model.	1012
073005	E	A cluster path failure has occurred.	1550
073006	W	The SAN is not correctly zoned. As a result, more than 512 ports on the SAN have logged into one SAN Volume Controller port.	1800
073101	E	The 2-port fibre-channel adapter card in slot 1 is missing. This error applies to only the SAN Volume Controller 2145-8F2 model.	1014
073102	E	The 2-port fibre-channel adapter in slot 1 has failed. This error applies to only the SAN Volume Controller 2145-8F2 model.	1054
073104	E	The 2-port fibre-channel adapter in slot 1 has detected a PCI bus error. This error applies to only the SAN Volume Controller 2145-8F2 model.	1017
073201	E	The 2-port fibre-channel adapter in slot 2 is missing. This error applies to only the SAN Volume Controller 2145-8F2 model.	1015

Table 46. Error codes (continued)

Error ID	SNMP Type	Condition	Error Code
073202	E	The 2-port fibre-channel adapter in slot 2 has failed. This error applies to only the SAN Volume Controller 2145-8F2 model.	1056
073204	E	The 2-port fibre-channel adapter in slot 2 has detected a PCI bus error. This error applies to only the SAN Volume Controller 2145-8F2 model.	1018
073251	E	The 4-port fibre-channel adapter in slot 1 is missing. This error applies to only the SAN Volume Controller 2145-8G4 model.	1011
073252	E	The 4-port fibre-channel adapter in slot 1 has failed. This error applies to only the SAN Volume Controller 2145-8G4 model.	1055
073258	E	The 4-port fibre-channel adapter in slot 1 has detected a PCI bus error. This error applies to only the SAN Volume Controller 2145-8G4 model.	1013
073301	E	The 4-port fibre-channel adapter in slot 2 is missing. This error applies to only the SAN Volume Controller 2145-8F4 model.	1016
073302	E	The 4-port fibre-channel adapter in slot 2 has failed. This error applies to only the SAN Volume Controller 2145-8F4 model.	1057
073304	E	The 4-port fibre-channel adapter in slot 2 has detected a PCI bus error. This error applies to only the SAN Volume Controller 2145-8F4 model.	1019
073305	E	One or more fibre-channel ports are running at a speed that is lower than the last saved speed. This error applies to both the SAN Volume Controller 2145-8F4 and the SAN Volume Controller 2145-8G4 models.	1065
073310	E	A duplicate fibre-channel frame has been detected, which indicates that there is an issue with the fibre-channel fabric. Other fibre-channel errors might also be generated.	1203
074001	W	Unable to determine the vital product data (VPD) for an FRU. This is probably because a new FRU has been installed and the software does not recognize that FRU. The cluster continues to operate; however, you must upgrade the software to fix this warning.	2040
074002	E	The node warm started after a software error.	2030
075001	E	The flash boot device has failed. This error applies to the SAN Volume Controller 2145-4F2, the SAN Volume Controller 2145-8F2, and the SAN Volume Controller 2145-8F4 models.	1040
075002	E	The flash boot device has recovered. This error applies to the SAN Volume Controller 2145-4F2, the SAN Volume Controller 2145-8F2, and the SAN Volume Controller 2145-8F4 models.	1040
075005	E	A service controller read failure has occurred. This error applies to the SAN Volume Controller 2145-4F2, the SAN Volume Controller 2145-8F2, and the SAN Volume Controller 2145-8F4 models.	1044

Table 46. Error codes (continued)

Error ID	SNMP Type	Condition	Error Code
075011	E	The flash boot device has failed. This errors applies to only the SAN Volume Controller 2145-8G4 model.	1040
075012	E	The flash boot device has recovered. This errors applies to only the SAN Volume Controller 2145-8G4 model.	1040
075015	E	A service controller read failure has occurred. This errors applies to only the SAN Volume Controller 2145-8G4 model.	1044
076001	E	The internal disk for a node has failed.	1030
076002	E	The hard disk is full and cannot capture any more output.	2030
077001	E	The system board service processor shows that fan 1 has failed. This error applies to only the SAN Volume Controller 2145-4F2 model.	1070
077002	E	The system board service processor shows that fan 2 has failed. This error applies to only the SAN Volume Controller 2145-4F2 model.	1070
077003	E	The system board service processor shows that fan 3 has failed. This error applies to only the SAN Volume Controller 2145-4F2 model.	1070
077004	E	The system board service processor shows that fan 4 has failed. This error applies to only the SAN Volume Controller 2145-4F2 model.	1070
077005	E	The system board service processor shows that fan 5 has failed. This error applies to only the SAN Volume Controller 2145-4F2 model.	1071
077011	E	The system board service processor shows that the ambient temperature threshold has exceeded. This error applies to only the SAN Volume Controller 2145-4F2 model.	1075
077012	E	The system board service processor shows that temperature warning threshold has exceeded. This error applies to only the SAN Volume Controller 2145-4F2 model.	1076
077013	E	The system board service processor shows that the soft or hard shutdown temperature threshold has exceeded. This error applies to only the SAN Volume Controller 2145-4F2 model.	1077
077021	E	The system board service processor shows that Voltage 1, (12 volt) is outside the set thresholds. This error applies to only the SAN Volume Controller 2145-4F2 model.	1080
077022	E	The system board service processor shows that Voltage 2, (5 volt) is outside the set thresholds. This error applies to only the SAN Volume Controller 2145-4F2 model.	1080
077023	E	The system board service processor shows that Voltage 3, (3.3 volt) is outside the set thresholds. This error applies to only the SAN Volume Controller 2145-4F2 model.	1080
077024	E	The system board service processor shows that Voltage 4, (2.5 volt) is outside the set thresholds. This error applies to only the SAN Volume Controller 2145-4F2 model.	1081
077025	E	The system board service processor shows that Voltage 5, (1.5 volt) is outside the set thresholds. This error applies to only the SAN Volume Controller 2145-4F2 model.	1081

Table 46. Error codes (continued)

Error ID	SNMP Type	Condition	Error Code
077026	E	The system board service processor shows that Voltage 6, (1.25 volt) is outside the set thresholds. This error applies to only the SAN Volume Controller 2145-4F2 model.	1081
077027	E	The system board service processor shows that Voltage 7, (CPU volts) is outside the set thresholds. This error applies to only the SAN Volume Controller 2145-4F2 model.	1081
077101	E	The service processor shows a fan 40×40×28 failure. This error applies to both the SAN Volume Controller 2145-8F2 and the SAN Volume Controller 2145-8F4 models.	1090
077102	E	The service processor shows a fan 40×40×56 failure. This error applies to both the SAN Volume Controller 2145-8F2 and the SAN Volume Controller 2145-8F4 models.	1091
077105	E	The service processor shows a fan failure. This errors applies to only the SAN Volume Controller 2145-8G4 model.	1089
077111	E	The node ambient temperature threshold has exceeded. This error applies to both the SAN Volume Controller 2145-8F2 and the SAN Volume Controller 2145-8F4 models.	1094
077112	E	The node processor warning temperature threshold has exceeded. This error applies to both the SAN Volume Controller 2145-8F2 and the SAN Volume Controller 2145-8F4 models.	1093
077113	E	The node processor or ambient critical threshold has exceeded. This error applies to both the SAN Volume Controller 2145-8F2 and the SAN Volume Controller 2145-8F4 models.	1092
077121	E	System board - any voltage high. This error applies to both the SAN Volume Controller 2145-8F2 and the SAN Volume Controller 2145-8F4 models.	1100
077124	E	System board - any voltage low. This error applies to both the SAN Volume Controller 2145-8F2 and the SAN Volume Controller 2145-8F4 models.	1105
077128	E	A power management board voltage failure has occurred. This error applies to both the SAN Volume Controller 2145-8F2 and the SAN Volume Controller 2145-8F4 models.	1110
077161	E	The node ambient temperature threshold has exceeded. This errors applies to only the SAN Volume Controller 2145-8G4 model.	1094
077162	E	The node processor warning temperature threshold has exceeded. This errors applies to only the SAN Volume Controller 2145-8G4 model.	1093
077163	E	The node processor or ambient critical threshold has exceeded. This errors applies to only the SAN Volume Controller 2145-8G4 model.	1092
077171	E	System board - any voltage high. This errors applies to only the SAN Volume Controller 2145-8G4 model.	1101
077174	E	System board - any voltage low. This errors applies to only the SAN Volume Controller 2145-8G4 model.	1106

Table 46. Error codes (continued)

Error ID	SNMP Type	Condition	Error Code
077178	E	A power management board voltage failure has occurred. This errors applies to only the SAN Volume Controller 2145-8G4 model.	1110
078001	E	A power domain error has occurred. Both nodes in a pair are powered by the same uninterruptible power supply.	1155
079000	W	Data has not been recovered on virtual disks (VDisks).	1850
079500	W	The limit on the number of cluster secure shell (SSH) sessions has been reached.	2500
081001	E	An Ethernet port failure has occurred.	1400
082001	E	A server error has occurred.	2100
083001	E	An uninterruptible power supply (UPS) communications failure has occurred. The RS232 connection between a node and its UPS is faulty. This error applies to only the 2145 UPS model.	1145
083002	E	The uninterruptible power supply (UPS) output is unexpectedly high. The UPS is probably connected to a non-SAN Volume Controller load. This error applies to only the 2145 UPS model.	1165
083003	E	The uninterruptible power supply battery has reached end of life. This error applies to only the 2145 UPS model.	1190
083004	E	An uninterruptible power supply battery failure has occurred. This error applies to only the 2145 UPS model.	1180
083005	E	An uninterruptible power supply electronics failure has occurred. This error applies to only the 2145 UPS model.	1170
083006	E	Uninterruptible power supply frame fault	1175
083007	E	Uninterruptible power supply frame fault overcurrent. This error applies to only the 2145 UPS model.	1160
083008	E	An uninterruptible power supply failure has occurred. This error applies to only the 2145 UPS model.	1185
083009	E	Uninterruptible power supply AC input power fault. This error applies to only the 2145 UPS model.	1140
083010	E	An uninterruptible power supply configuration error has occurred. This error applies to only the 2145 UPS model.	1150
083011	E	Uninterruptible power supply ambient over temperature. This error applies to only the 2145 UPS model.	1135
083012	E	Uninterruptible power supply over temperature warning. This error applies to only the 2145 UPS model.	3000
083013	E	The cross cable test was bypassed because of an internal uninterruptible power supply software error. This error applies to only the 2145 UPS model.	3010
083101	E	An uninterruptible power supply (UPS) communications failure has occurred. The RS232 connection between a node and its UPS is faulty. This error applies to only the 2145 UPS-1U model.	1146

Table 46. Error codes (continued)

Error ID	SNMP Type	Condition	Error Code
083102	E	The uninterruptible power supply (UPS) output is unexpectedly high. The UPS is probably connected to a non-SAN Volume Controller load. This error applies to only the 2145 UPS-1U model.	1166
083103	E	The uninterruptible power supply battery has reached end of life. This error applies to only the 2145 UPS-1U model.	1191
083104	E	An uninterruptible power supply battery failure has occurred. This error applies to only the 2145 UPS-1U model.	1181
083105	E	An uninterruptible power supply electronics failure has occurred. This error applies to only the 2145 UPS-1U model.	1171
083107	E	Uninterruptible power supply overcurrent. This error applies to only the 2145 UPS-1U model.	1161
083108	E	An uninterruptible power supply failure has occurred. This error applies to only the 2145 UPS-1U model.	1186
083109	E	Uninterruptible power supply AC input power fault. This error applies to only the 2145 UPS-1U model.	1141
083110	E	An uninterruptible power supply configuration error has occurred. This error applies to only the 2145 UPS-1U model.	1151
083111	E	Uninterruptible power supply ambient over temperature. This error applies to only the 2145 UPS-1U model.	1136
083112	E	Uninterruptible power supply over temperature warning. This error applies to only the 2145 UPS-1U model.	3001
083113	E	An uninterruptible power supply software error has occurred. This error applies to only the 2145 UPS-1U model.	3011

Appendix E. 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 is raised.

Configuration event codes are generated when configuration parameters are set. Configuration event codes are recorded in a separate log and do not raise SNMP traps. 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, if configured, an SNMP trap is raised and an e-mail is sent.

Information event codes can be either SNMP trap type I (information) or type W (warning). You can use the SNMP trap type that is included in the e-mail to determine if the information event was caused by an expected or unexpected condition. An information event report of type (W) might require user attention. Table 47 provides a list of information event codes, the SNMP type, and the meaning of the event code.

Table 47. Information event codes

Event code	SNMP 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	The managed disk is not on the preferred path.
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 47. Information event codes (continued)

Event code	SNMP Type	Description
982004	W	Migration suspended due to insufficient virtual extents or too many media errors on the source managed disk.
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 or Global Mirror background copy is complete.
985002	I	Metro Mirror or Global Mirror is 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. They do not raise SNMP traps or send e-mails. Their error fixed flags are ignored. Table 48 provides a list of the configuration event codes and their meanings.

Table 48. Configuration event codes

Event code	Description
990101	Modify cluster (attributes in the <code>svctask chcluster</code> command)

Table 48. Configuration event codes (continued)

Event code	Description
990105	Delete node from cluster (attributes in the svctask rmnode command)
990106	Create host (attributes in the svctask mkhost command)
990112	Cluster configuration dumped to file (attributes from the svcluster -x dumpconfig 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)

Table 48. Configuration event codes (continued)

Event code	Description
990187	Prepare a FlashCopy mapping (attributes in the svctask prestartfcmap command)
990188	Prepare a FlashCopy consistency group (attributes in the svctask prestartfcconsistgrp command)
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 mkfcconsistgrp command)
990205	Modify FlashCopy consistency group (attributes in the svctask chfcconsistgrp command)
990206	Delete FlashCopy consistency group (attributes in the svctask rmfcconsistgrp 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 a Metro Mirror or Global Mirror or Global Mirror relationship (attributes in the svctask mkrrelationship command)
990226	Modify a Metro Mirror or Global Mirror relationship (attributes in the svctask chrrelationship command)
990227	Delete a Metro Mirror or Global Mirror relationship (attributes in the svctask rmrrelationship command)

Table 48. Configuration event codes (continued)

Event code	Description
990229	Start a Metro Mirror or Global Mirror relationship (attributes in the svctask startcrrelationship command)
990230	Stop a Metro Mirror or Global Mirror relationship (attributes in the svctask stopcrrelationship command)
990231	Switch a Metro Mirror or Global Mirror relationship (attributes in the svctask switchcrrelationship command)
990232	Start a Metro Mirror or Global Mirror consistency group (attributes in the svctask startcrconsistgrp command)
990233	Stop a Metro Mirror or Global Mirror consistency group (attributes in the svctask stopcrconsistgrp command)
990234	Switch a Metro Mirror or Global Mirror consistency group (attributes in the svctask switchcrconsistgrp 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 a Metro Mirror or Global Mirror consistency group (attributes in the svctask mkrconsistgrp command)
990241	Modify a Metro Mirror or Global Mirror consistency group (attributes in svctask chrconsistgrp)
990242	Delete a Metro Mirror or Global Mirror consistency group (attributes in the svctask rmrconsistgrp 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	The software installation has started.
990415	The software installation has completed.
990420	The software installation has failed.

Table 48. Configuration event codes (continued)

Event code	Description
990423	The software installation has stalled.
990425	The software installation has stopped.
990430	The Planar Serial Number has changed.
990501	The featurization has changed. See the feature log for details.
990510	The configuration limits have been changed.
991024	I/O tracing has finished and the managed disk has been triggered.

Appendix F. SCSI error reporting

SAN Volume Controller nodes can notify their hosts of errors for SCSI commands that are issued.

SCSI status

Some errors are part of the SCSI architecture and are handled by the host application or device drivers without reporting an error. Some errors, such as read and write I/O errors and errors that are associated with the loss of nodes or loss of access to backend devices, cause application I/O to fail. To help troubleshoot these errors, SCSI commands are returned with the Check Condition status and a 32-bit event identifier is included with the sense information. The identifier relates to a specific error in the SAN Volume Controller cluster error log.

If the host application or device driver captures and stores this error information, you can relate the application failure to the error log.

Table 49 describes the SCSI status and codes that are returned by the SAN Volume Controller nodes.

Table 49. SCSI status

Status	Code	Description
Good	00h	The command was successful.
Check condition	02h	The command failed and sense data is available.
Condition met	04h	N/A
Busy	08h	An Auto-Contingent Allegiance condition exists and the command specified NACA=0.
Intermediate	10h	N/A
Intermediate - condition met	14h	N/A
Reservation conflict	18h	Returned as specified in SPC2 and SAM2 where a reserve or persistent reserve condition exists.
Task set full	28h	The initiator has at least one task queued for that LUN on this port.
ACA active	30h	This is reported as specified in SAM-2.
Task aborted	40h	This is returned if TAS is set in the control mode page 0Ch. The SAN Volume Controller node has a default setting of TAS=0, which is cannot be changed; therefore, the SAN Volume Controller node does not report this status.

SCSI Sense

SAN Volume Controller nodes notify the hosts of errors on SCSI commands. Table 50 on page 398 defines the SCSI sense keys, codes and qualifiers that are returned by the SAN Volume Controller nodes.

Table 50. SCSI sense keys codes and qualifiers

Key	Code	Qualifier	Definition	Description
2h	04h	01h	Not Ready. The logical unit is in the process of becoming ready.	The node lost sight of the cluster and cannot perform I/O operations. The additional sense does not have additional information.
2h	04h	0Ch	Not Ready. The target port is in the state of unavailable.	The following conditions are possible: <ul style="list-style-type: none"> The node lost sight of the cluster and cannot perform I/O operations. The additional sense does not have additional information. The node is in contact with the cluster but cannot perform I/O operations to the specified logical unit because of either a loss of connectivity to the backend controller or some algorithmic problem. This sense is returned for offline virtual disks (VDisks).
3h	00h	00h	Medium error	This is only returned for read or write I/Os. The I/O suffered an error at a specific LBA within its scope. The location of the error is reported within the sense data. The additional sense also includes a reason code that relates the error to the corresponding error log entry. For example, a RAID controller error or a migrated medium error.
4h	08h	00h	Hardware error. A command to logical unit communication failure has occurred.	The I/O suffered an error that is associated with an I/O error that is returned by a RAID controller. The additional sense includes a reason code that points to the sense data that is returned by the controller. This is only returned for I/O type commands. This error is also returned from FlashCopy target VDisks in the prepared and preparing state.
5h	25h	00h	Illegal request. The logical unit is not supported.	The logical unit does not exist or is not mapped to the sender of the command.

Reason codes

The reason code appears in bytes 20-23 of the sense data. The reason code provides the SAN Volume Controller node specific log entry. The field is a 32-bit unsigned

number that is presented with the most significant byte first. Table 51 lists the reason codes and their definitions.

If the reason code is not listed in Table 51, the code refers to a specific error in the SAN Volume Controller cluster error log that corresponds to the sequence number of the relevant error log entry.

Table 51. Reason codes

Reason code (decimal)	Description
40	The resource is part of a stopped FlashCopy mapping.
50	The resource is part of a Metro Mirror or Global Mirror relationship and the secondary LUN is the offline.
51	The resource is part of a Metro Mirror or Global Mirror and the secondary LUN is read only.
60	The node is offline.
71	The resource is not bound to any domain.
72	The resource is bound to a domain that has been recreated.
73	Running on a node that has been contracted out for some reason that is not attributable to any path going offline.

Appendix G. Object types

You can use the object code to determine the object type.

Table 52 lists the object codes and corresponding object types.

Table 52. 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

Accessibility

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

Features

These are the major accessibility features in the SAN Volume Controller Console :

- You can use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. The following screen readers have been tested: WebKing v5.5 and Window-Eyes v5.5.
- You can operate all features using the keyboard instead of the mouse.
- You can change the initial delay and repeat rate of the up and down buttons to two seconds when you use the front panel of the SAN Volume Controller to set or change an IP address. This feature is documented in the applicable sections of the SAN Volume Controller publications.

Navigating by keyboard

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

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

Accessing the publications

You can view the publications for the SAN Volume Controller in Adobe Portable Document Format (PDF) using the Adobe Acrobat Reader. The PDFs are provided at the following Web site:

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

Related reference

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

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.*

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

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

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

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATIONS "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

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

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

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

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

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

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

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

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

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

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

Related reference

"Trademarks" on page 407

Trademarks

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

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

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

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

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

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

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.

asymmetric virtualization

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

auxiliary virtual disk

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

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, up to four pairs of nodes that provide a single configuration and service interface.

concurrent maintenance

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

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

configuration node

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

consistency group

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

consistent copy

In a Metro or Global Mirror relationship, a copy of a secondary virtual

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

container

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

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

copied

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

copying

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

counterpart SAN

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

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 *IBM System Storage SAN Volume Controller: Service Guide*.

disconnected

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

discovery

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

disk controller

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

disk 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*.

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

FlashCopy mapping

A relationship between two virtual disks.

FlashCopy relationship

See *FlashCopy mapping*.

FlashCopy service

In SAN Volume Controller, a copy service that duplicates the contents of a

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

G

Global Mirror

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

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

IBM TotalStorage Enterprise Storage Server (ESS)

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

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

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

illegal configuration

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

image mode

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

image VDisk

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

inconsistent

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

initiator

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

input/output (I/O)

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

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).

managed space mode

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

mapping

See *FlashCopy mapping*.

master virtual disk

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

MDisk

See *managed disk*.

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

paused

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

pend To cause to wait for an event.

point-in-time copy

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

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

primary virtual disk

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

PuTTY

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

Q

queue depth

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

quorum disk

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

R

RAID See *redundant array of independent disks*.

RAID 0

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

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

RAID 1

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

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

HP definition: See *mirrorset*.

RAID 5

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

IBM definition: See the SNIA definition.

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

RAID 10

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

redundant array of independent disks (RAID)

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

redundant SAN

A storage area network (SAN) configuration in which any one single component might fail, but connectivity between the devices within the SAN is maintained, possibly with degraded performance. This configuration is normally achieved by splitting the SAN into two, independent, counterpart SANs. See also *counterpart SAN*.

rejected

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

relationship

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

S

SAN See *storage area network*.

SAN Volume Controller fibre-channel port fan in

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

SATA See *Serial Advanced Technology Attachment*.

SCSI See *Small Computer Systems Interface*.

secondary virtual disk

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

sequential VDisk

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

Serial Advanced Technology Attachment (SATA)

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

Serial ATA

See *Serial Advanced Technology Attachment*.

Simple Mail Transfer Protocol (SMTP)

An Internet application protocol for transferring mail among users of the Internet. SMTP specifies the mail exchange sequences and message format. It assumes that the Transmission Control Protocol (TCP) is the underlying protocol.

Simple Network Management Protocol (SNMP)

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

Small Computer System Interface (SCSI)

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

SMTP See *Simple Mail Transfer Protocol*.

SNMP

See *Simple Network Management Protocol*.

stand-alone relationship

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

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

stopped

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

storage area network (SAN)

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

stripeset

See *RAID 0*.

subsystem device driver (SDD)

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

superuser authority

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

T**trigger**

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

U**unconfigured mode**

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

uninterruptible power supply (UPS)

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

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.

worldwide port name (WWPN)

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

WWNN

See *worldwide node name*.

WWPN

See *worldwide port name*.

Index

Numerics

- 2145 uninterruptible power supply-1U
 - configuration 13
 - operation 14

A

- about this guide xv
- Access Logix 257
- accessibility
 - keyboard 403
 - repeat rate of up and down buttons 403
 - shortcut keys 403
- Add/Remove Programs dialog panel 352
- adding
 - managed disks 180
 - managed disks (MDisks) 126, 129
 - nodes 112
 - storage controllers
 - using the 249
 - using the CLI (command-line interface) 250
- analyzing error logs 157
- AOS remote service 30
- application programming interface 6
- Assist On-site 30
- audience xv

B

- backup cluster configuration files
 - creating 219
- backup configuration files
 - creating 217
 - deleting 223
 - using the CLI 224
 - restoring 221
- book
 - about this xv
- boot drive mirroring 92
- Brocade
 - core-edge fabrics 75
 - switch ports 75
- browsers
 - / see also Web browsers 96

C

- Call Home 30
- changing
 - cluster password 111
 - passwords 212
- checking
 - status of node ports 115
- CLI (command-line interface)
 - configuring PuTTY 99
 - getting started 165

- CLI (command-line interface) (*continued*)
 - issuing commands from a PuTTY SSH client system 168
 - preparing SSH client systems 166, 167
 - upgrading software 225
 - using to set cluster features 171
- cluster
 - split configurations 14
- clusters
 - adding managed disks (MDisks) 126
 - adding nodes 112
 - backing up configuration file 8, 217
 - backing up configuration file using the CLI 219
 - Call Home e-mail 30
 - changing fabric speed 124
 - changing password 111
 - creating
 - from the front panel 101
 - deleting nodes 122, 208
 - error logs 213
 - gateway address
 - changing 210
 - Global Mirror partnerships
 - deleting 154
 - including managed disks (MDisks) 126
 - IP address
 - changing 209
 - IP failover 9
 - logs 213
 - maintaining 155
 - Metro Mirror partnerships
 - deleting 154
 - overview 7
 - properties 112, 171
 - recovering nodes 197
 - removing nodes 122, 208
 - renaming 124
 - resetting the SSH fingerprint 162
 - restoring backup configuration files 221
 - setting
 - features 171
 - time 170
 - setting date 110
 - setting time 110
 - shutting down 124, 214
 - subnet mask
 - changing 210
 - viewing feature logs 213
 - viewing properties 112
- codes
 - configuration events 392
 - events 391
 - information events 391
- command-line interface (CLI)
 - configuring 98
 - configuring PuTTY 99
 - getting started 165

- command-line interface (CLI) (*continued*)
 - issuing commands from a PuTTY SSH client system 168
 - preparing SSH clients 166, 167
 - upgrading software 225
 - using to set cluster features 171
 - using to set cluster time 170
- commands
 - ibmvmcfc add 372
 - ibmvmcfc listvols 372
 - ibmvmcfc rem 372
 - ibmvmcfc set cimomHost 370
 - ibmvmcfc set cimomPort 370
 - ibmvmcfc set namespace 370
 - ibmvmcfc set password 370
 - ibmvmcfc set trustpassword 370
 - ibmvmcfc set truststore 370
 - ibmvmcfc set username 370
 - ibmvmcfc set usingSSL 370
 - ibmvmcfc set vssFreeInitiator 370
 - ibmvmcfc set vssReservedInitiator 370
 - ibmvmcfc showcfg 370
 - svconfig backup 219
 - svconfig restore 221
 - svctask detectmdisk 246
- Common Information Model (CIM) 6
- communications
 - determining between hosts and virtual disks 189
- configuration
 - event codes 392
 - maximum sizes 28
 - mesh 65
 - node failover 9
 - rules 65
- configuration requirements 86
- configuring
 - clusters 106
 - command-line interface (CLI) 98
 - disk controllers 235, 236, 237, 238
 - DS4000 series Storage Manager 280
 - Enterprise Storage Server 241, 276
 - error notification settings 155
 - FAStT Storage Manager 241
 - FAStT Storage Server 241
 - host name 95, 356
 - IBM DS4000 Storage Server 279
 - IBM DS6000 287
 - IBM DS8000 289
 - master console 95, 96, 98
 - nodes 73
 - PuTTY 99
 - SAN 68
 - SAN Volume Controller 73, 98
 - settings
 - error notification 155
 - switches 74
 - Web browsers 96
- consistency group, Mirror 56
- consistency groups, FlashCopy 43
 - creating 147

consistency groups, FlashCopy
(continued)

- deleting 148
- filtering 147
- modifying 148
- starting 147
- stopping 148

console

- master 29
- SAN Volume Controller
 - master console 5, 29
 - portfolio 104
 - starting 105
 - task bar 104
 - user interface 6
 - work area 105

controllers

- adding
 - using the 249
 - using the CLI (command-line interface) 250
- advanced functions

- DS4000 series 283
- EMC CLARiiON 263
- EMC Symmetrix 269
- EMC Symmetrix DMX 269
- Fujitsu ETERNUS 276
- HDS Lightning 292
- HDS NSC 307
- HDS TagmaStore AMS 299
- HDS TagmaStore WMS 299
- HDS Thunder 299
- HDS USP 307
- HP EVA 323
- HP MSA 328
- HP StorageWorks EMA 316, 317
- HP StorageWorks MA 316, 317
- HP XP 307
- IBM Enterprise Storage Server 278
- IBM N5000 335
- NetApp FAS 335
- Sun StorEdge 307

concurrent maintenance

- DS4000 series 282
- EMC CLARiiON 261
- EMC Symmetrix 267
- EMC Symmetrix DMX 267
- Enterprise Storage Server 277
- Fujitsu ETERNUS 275
- HDS Lightning 291
- HDS NSC 307
- HDS TagmaStore AMS 296
- HDS TagmaStore WMS 296
- HDS Thunder 296
- HDS USP 307
- HP MSA 327
- HP StorageWorks EMA 314
- HP StorageWorks MA 314
- HP XP 307
- IBM DS6000 288
- IBM DS8000 290
- IBM N5000 333
- NetApp FAS 333
- Sun StorEdge 307

configuration

- Bull FDA 256

controllers (continued)

configuration (continued)

- EMC CLARiiON 257, 258, 260, 264
- EMC Symmetrix 266, 270
- EMC Symmetrix DMX 266, 270
- Enterprise Storage Server 276
- Fujitsu ETERNUS 272
- HDS Lightning 291
- HDS NSC 305
- HDS SANrise 1200 296
- HDS TagmaStore AMS 296
- HDS TagmaStore WMS 296
- HDS Thunder 296
- HDS USP 305
- HP EVA 322
- HP MSA 325
- HP StorageWorks 324
- HP StorageWorks EMA 308, 311, 312, 318
- HP StorageWorks MA 308, 311, 312, 318
- HP XP 291, 305
- IBM DS4000 279
- IBM DS6000 287
- IBM DS8000 289
- IBM N5000 329
- IBM N7000 329
- NEC iStorage 328
- NetApp FAS 329
- StorageTek D 279
- StorageTek Flexline 279
- Sun StorEdge 291, 305

controller settings

- EMC CLARiiON 265

firmware

- Bull FDA 256
- DS4000 series 282
- EMC CLARiiON 261
- EMC Symmetrix 267
- EMC Symmetrix DMX 267
- Fujitsu ETERNUS 272
- HDS Lightning 291
- HDS NSC 305
- HDS TagmaStore AMS 296
- HDS TagmaStore WMS 296
- HDS Thunder 296
- HDS USP 305
- HP EVA 322
- HP MSA 325
- HP StorageWorks EMA 314
- HP StorageWorks MA 314
- HP XP 305
- IBM DS6000 288
- IBM DS8000 290
- IBM Enterprise Storage Server 277
- IBM N5000 330
- NEC iStorage 328
- NetApp FAS 330
- Sun StorEdge 305

global settings

- EMC CLARiiON 264
- EMC Symmetrix 270
- EMC Symmetrix DMX 270
- HDS TagmaStore AMS 300
- HDS TagmaStore WMS 300

controllers (continued)

global settings (continued)

- HDS Thunder 300
- HP EVA 324
- IBM DS4000 series 285
- Lightning 294

host settings

- HP EVA 325

host type

- HDS NSC 307
- HDS USP 307
- HP XP 307
- Sun StorEdge 307

interface

- DS4000 series 284
- HP StorageWorks 324
- HP StorageWorks EMA 315
- HP StorageWorks MA 315

logical unit creation and deletion

- EMC CLARiiON 264
- EMC Symmetrix 269
- HDS TagmaStore AMS 299
- HDS TagmaStore WMS 299
- HDS Thunder 299
- HP EVA 323
- HP StorageWorks EMA 318
- HP StorageWorks MA 318
- IBM DS4000 series 284
- IBM Enterprise Storage Server 278

logical unit presentation

- HP EVA 324

logical units

- Bull FDA 256
- HDS NSC 306
- HDS USP 306
- HP MSA 326
- HP XP 306
- IBM N5000 330
- NEC iStorage 328
- NetApp FAS3000 330
- Sun StorEdge 306

LU configuration

- HDS Lightning 293

LU settings

- EMC CLARiiON 265
- EMC Symmetrix 271
- EMC Symmetrix DMX 271
- HDS TagmaStore AMS 303
- HDS TagmaStore WMS 303
- HDS Thunder 303
- HP EVA 324
- HP StorageWorks EMA 320
- HP StorageWorks MA 320
- IBM DS4000 series 286
- Lightning 295

mapping settings

- EMC Symmetrix 271
- EMC Symmetrix DMX 271

models

- EMC CLARiiON 261
- EMC Symmetrix 266
- EMC Symmetrix DMX 266
- Fujitsu ETERNUS 272
- HDS Lightning 291
- HDS NSC 305
- HDS TagmaStore AMS 296

- controllers (*continued*)
 - models (*continued*)
 - HDS TagmaStore WMS 296
 - HDS Thunder 296
 - HDS USP 305
 - HP EVA 322
 - HP MSA 325
 - HP StorageWorks EMA 314
 - HP StorageWorks MA 314
 - HP XP 291, 305
 - IBM DS4000 281
 - IBM DS6000 288
 - IBM DS8000 290
 - IBM Enterprise Storage Server 277
 - IBM N5000 329
 - IBM N7000 329
 - NetApp FAS 329
 - StorageTek 281
 - Sun StorEdge 291, 305
 - port selection 246
 - port settings
 - EMC CLARiiON 265
 - EMC Symmetrix 270
 - EMC Symmetrix DMX 270
 - HDS Lightning 295
 - HDS TagmaStore AMS 302
 - HDS TagmaStore WMS 302
 - HDS Thunder 302
 - HP StorageWorks EMA 319
 - HP StorageWorks MA 319
 - quorum disks
 - DS4000 series 283
 - EMC CLARiiON 263
 - EMC Symmetrix 269
 - HDS Lightning 292
 - HDS NSC 307
 - HDS TagmaStore AMS 298
 - HDS TagmaStore WMS 298
 - HDS Thunder 298
 - HDS USP 307
 - HP EVA 323
 - HP MSA 328
 - HP StorageWorks EMA 316
 - HP StorageWorks MA 316
 - HP XP 307
 - IBM Enterprise Storage Server 278
 - IBM N5000 333
 - NetApp FAS 333
 - Sun StorEdge 307
 - registering
 - EMC CLARiiON 258
 - removing 251
 - using the CLI (command-line interface) 253
 - settings
 - DS4000 series 285
 - HDS TagmaStore AMS 300, 302
 - HDS TagmaStore WMS 300, 302
 - HDS Thunder 300, 302
 - HP StorageWorks EMA 318
 - HP StorageWorks MA 318, 321
 - HP StorageWorks MA EMA 321
 - IBM DS4000 series 285
 - Lightning 294, 295

- controllers (*continued*)
 - sharing
 - DS4000 series 282
 - EMC CLARiiON 262
 - EMC Symmetrix 268
 - EMC Symmetrix DMX 268
 - HDS Lightning 292
 - HDS TagmaStore AMS 297
 - HDS TagmaStore WMS 297
 - HDS Thunder 297
 - HP EVA 322
 - HP StorageWorks EMA 315
 - HP StorageWorks MA 315
 - IBM Enterprise Storage Server 278
 - switch zoning
 - EMC CLARiiON 262
 - EMC Symmetrix 268
 - EMC Symmetrix DMX 268
 - HDS NSC 306
 - HDS USP 306
 - HP EVA 323
 - HP StorageWorks EMA 316
 - HP StorageWorks MA 316
 - HP XP 306
 - IBM Enterprise Storage Server 278
 - IBM N5000 332
 - NetApp FAS 332
 - Sun StorEdge 306
 - target port groups
 - Enterprise Storage Server 289
 - target ports
 - Bull FDA 256
 - HDS NSC 306
 - HDS USP 306
 - HP MSA 326
 - HP XP 306
 - IBM N5000 330
 - NEC iStorage 328
 - NetApp FAS3000 330
 - Sun StorEdge 306
 - user interface
 - EMC CLARiiON 262
 - EMC Symmetrix 267
 - EMC Symmetrix DMX 267
 - Fujitsu ETERNUS 272
 - HDS Lightning 291
 - HDS NSC 306
 - HDS TagmaStore AMS 297
 - HDS TagmaStore WMS 297
 - HDS Thunder 297
 - HDS USP 306
 - HP EVA 322
 - HP MSA 325
 - HP XP 306
 - IBM DS4000 series 282
 - IBM DS6000 288
 - IBM DS8000 290
 - IBM Enterprise Storage Server 277
 - IBM N5000 330
 - NetApp FAS 330
 - Sun StorEdge 306
 - zoning 78
 - conventions
 - numbering xvii

- Copy Services
 - FlashCopy 35
 - incremental 37
 - mappings 37
 - multiple target 37
 - states 37
 - Global Mirror 49
 - Metro Mirror 49
 - overview 35
 - creating
 - clusters
 - from the front panel 101
 - FlashCopy
 - mappings 144, 185, 187
 - FlashCopy consistency groups 147
 - Global Mirror
 - partnerships 154
 - managed disk (MDisk) groups 128
 - Metro Mirror
 - partnerships 154
 - Mirror
 - consistency groups 152
 - quorum disks 254
 - virtual disk-to-host mappings 183
 - virtual disks (VDisks) 131
 - creating, SSH keys 98

D

- data
 - migrating 238, 239
- data migration
 - DS4000 series 283
- deleting
 - backup configuration files 223
 - using the CLI 224
 - FlashCopy
 - mappings 146
 - Global Mirror
 - consistency groups 153
 - partnerships 154
 - hosts 144
 - Metro Mirror
 - consistency groups 153
 - partnerships 154
 - Mirror
 - relationships 151
 - nodes 122, 208
 - virtual disks 136
- determining
 - communications between hosts and virtual disks 189
- discovering
 - managed disks 125, 177, 255
 - MDisks 125
- disk controller systems
 - renaming 248
- disk controllers
 - configuring 235, 236
 - overview 16
- disk failure 360
- disks
 - migrating 204
 - migrating image mode 207
 - display on front panel
 - Node rescue request 232

- disruptive software upgrade
 - using the CLI (command-line interface) 231
- DS4000 Storage Manager Client (FASfT Storage Manager Client)
 - uninstalling 354

E

- e-mail
 - Call Home 31
 - inventory information 32
- e-mail error notification 212
- emphasis in text xvii
- error codes 383
- error ID 383
- error messages, IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software 372
- errors
 - notification settings 155
- Ethernet
 - link failures 9
- Ethernet port, entering 96
- event
 - notification 30
- events
 - codes 391
 - configuration 392
 - information 391
- examples
 - SAN environments 80
 - SAN Volume Controller cluster in a SAN fabric 66
- expanding
 - logical units 245
 - virtual disks 200
- extents
 - migrating
 - using the CLI (command-line interface) 202

F

- fabric, SAN 66
- features
 - setting
 - using the CLI (command-line interface) 171
 - viewing logs 155
- fibre-channel
 - network, rescanning 125
- fibre-channel switches 74
- filtering
 - FlashCopy
 - consistency groups 147
 - mappings 145
 - Global Mirror
 - relationships 149
 - Metro Mirror
 - relationships 149
 - Mirror
 - consistency groups 152
- FlashCopy 42, 60
 - consistency groups 43

- FlashCopy (*continued*)
 - copy rate 48
 - creating consistency groups 147
 - creating mappings 185, 187
 - definition 412
 - for Volume Shadow Copy service 363
 - incremental 37
 - mapping 411
 - mappings 37, 184
 - memory 183
 - multiple target 37
 - overview 35
 - renaming consistency groups 148
 - states 37
- free pool of volumes 369
- front panel
 - password 172

G

- gateway address
 - changing 210
- generating a SSH key pair 98
- getting started
 - using the CLI (command-line interface) 165
 - using the command-line interface (CLI) 165
 - using the SAN Volume Controller Console 109
- Global Mirror 52, 53, 55, 59
 - bandwidth 57
 - configuration requirements 53
 - consistency groups
 - deleting 153
 - starting 149, 153
 - stopping 150, 153
 - deleting partnerships 154
 - gmlinktolerance feature 61
 - memory 183
 - overview 49
 - partnerships 52
 - creating 154
 - relationships
 - filtering 149
 - starting 149, 153
 - stopping 150, 153
 - requirements 58
 - restarting relationships 60
 - upgrading cluster software 225
 - zoning considerations 82
- Global Mirror performance, monitor
 - monitor
 - monitor performance 60
- governing 12
- guide
 - about this xv
 - who should read xv
- guidelines
 - zoning 78

H

- HBAs (host bus adapters)
 - configuration 72

- HBAs (host bus adapters) (*continued*)
 - node 73
 - replacing 143
- HDS TagmaStore AMS
 - support 296
- HDS TagmaStore WMS
 - support 296
- HDS Thunder
 - support 296
- host bus adapters (HBAs)
 - configuration 72
 - node 73
 - replacing 143
- host name, configuring 95, 356
- host objects
 - creating 182
- hosts
 - creating 140
 - deleting 144
 - determining VDisk names 189
 - filtering 140
 - flushing data 37
 - mapped virtual disks (VDisks) 141
 - mapping virtual disks (VDisks) 183
 - overview 25
 - replacing HBA 143
 - supported 6
 - traffic 55
 - viewing details 140
 - viewing mapped I/O groups 141
 - viewing ports 141
 - zoning 78

I

- I/O governing 12
- I/O groups 73
 - overview 11
 - renaming 123
- IBM Director
 - uninstalling 353
- IBM System Storage hardware provider
 - installation procedure 363
 - system requirements 364
- IBM System Storage Support for Microsoft Volume Shadow Copy Service
 - installing 364
- IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software
 - creating pools of volumes 369
 - error messages 372
 - ibmvfcg.exe 370, 372
 - installation procedure 363
 - overview 363
 - system requirements 364
 - uninstalling 374
 - verifying the installation 370
- IBM TotalStorage Productivity Center 60
 - ibmvfcg.exe 370, 372
- image mode
 - VDisks 137
- image mode VDisks
 - converting to managed mode
 - using 137
 - using CLI (command-line interface) 206

- including
 - managed disks (MDisks) 126
- information
 - event codes 391
- information center xvii
- installing 89
 - IBM System Storage Support for Microsoft Volume Shadow Copy Service 364
 - master console software 89
 - PuTTY 89, 338, 343
 - SAN Volume Controller Console 89
 - software 225
 - verification 349
- installing PuTTY 165
- interfaces
 - application program 6
 - user 6
- interswitch link (ISL)
 - congestion 77
 - maximum hop count 75
 - oversubscription 77
- inventory information 30
- IP addresses
 - changing 209
 - modifying 111
- IP network connection, configuring 96
- issuing
 - CLI commands 168

K

- key
 - replacing SSH keys 161
- keyboard
 - navigating by 403
 - shortcuts 403

L

- language 213
- legal notices 405
- license
 - disabling features 155
 - enabling features 155
 - viewing features 155
- links, physical 78
- listing
 - dump files 156
 - log files 156
- Local Area Connection 96
- logical unit mapping 245
- logical units
 - expanding 245

M

- maintaining
 - passwords 111, 172
 - SSH keys 210
- maintenance procedures
 - clusters 155
- managed disk (MDisk) 18
- managed disk (MDisk) groups
 - adding
 - managed disks 129

- managed disk (MDisk) groups (*continued*)
 - creating 128
 - renaming 130
- managed disk groups
 - creating using the CLI 178
- managed disks (MDisks)
 - adding 126, 180
 - discovering 125, 177, 255
 - displaying groups 128
 - expanding 245
 - including 126
 - rebalancing access 177, 255
 - removing from a managed disk group 129
 - removing from an MDisk group 129
 - renaming 126
 - virtual disks (VDisks)
 - relationships 190
- managed mode virtual disks
 - converting from image mode
 - using the 137
 - using the CLI (command-line interface) 206
- managing
 - tools 29
- mapping events 42
- mappings, FlashCopy
 - copy rate 48
 - creating 144
 - deleting 146
 - events 42
 - filtering 145
 - modifying 146
 - starting 145
 - stopping 146
- master console
 - configuring 95, 96, 98
 - disk failure 360
 - hardware option 29
 - installing software 89
 - overview 29
 - removing 352, 356
 - software components 29
 - software-only option 29
 - troubleshooting 359, 360
 - upgrading software 336
- master console software 89
- maximum configuration 28
- MDisk (managed disk) 18
- MDisk (managed disk) groups
 - deleting 131
 - forced 131
 - renaming 130
- MDisks (managed disks)
 - adding 180
 - VDisk (virtual disks)
 - relationships 190
- measurements xvii
- mesh configuration 65
- metro mirror
 - upgrading cluster software 225
- Metro Mirror 52, 53, 55, 59
 - bandwidth 57
 - consistency groups
 - deleting 153
 - starting 149, 153
 - stopping 150, 153

- Metro Mirror (*continued*)
 - deleting partnerships 154
 - memory 183
 - overview 49
 - partnerships 52
 - creating 154
 - relationships
 - filtering 149
 - starting 149, 153
 - stopping 150, 153
 - upgrading cluster software 225
 - zoning considerations 82
- Microsoft Windows, troubleshooting 360
- migrating
 - data 238, 239
 - extents
 - using the CLI (command-line interface) 202
 - VDisks (virtual disks) 192
 - virtual disks (VDisks) 117, 139
- migration 132, 283
- Mirror
 - consistency groups
 - creating 152
 - filtering 152
 - overview 56
 - relationships
 - deleting 151
- mirroring, boot drive 92
- modifying 245
 - FlashCopy
 - consistency groups 148
 - mappings 146
 - Global Mirror
 - partnerships 154
 - relationships 150
 - Metro Mirror
 - partnerships 154
 - relationships 150
- monitoring
 - software upgrades 227, 233, 234

N

- node
 - failover 9
- node status 9
- nodes
 - adding 112, 172
 - configuration 10, 73
 - deleting 122, 208
 - HBAs 73
 - overview 7
 - removing 122, 208
 - renaming 121
 - replacing 118, 193
 - rescuing 232
 - returning to cluster 197
 - shutting down 125
 - status 115
 - viewing
 - general details 115, 176
 - virtual disks (VDisks) 73
- notifications
 - inventory information 32
 - sending 30

O

- object classes and instances 401
- object codes 401
- object types 401
- operating over long distances 83
- ordering publications xxi
- oversubscription 65
- overview
 - pscp application 226
 - SAN fabric 66
 - SAN Volume Controller 1
 - zoning 80

P

- partnerships, Global Mirror
 - modifying 154
- partnerships, Metro Mirror
 - modifying 154
- passwords
 - changing 212
 - front panel 172
- plink utility
 - running 168
- port speed 74
- preinstalled software
 - recovering from installation failures 234
- preparing
 - SSH client system 166
 - overview 165
 - to issue CLI commands 167
- public secure shell keys 160
- publications
 - accessing 403
 - ordering xxi
- PuTTY 29, 99
 - configuring 99
 - generating a key pair 98
 - installing 89, 165
 - issuing CLI commands from 168
 - running the plink utility 168
 - scp (pscp) 226
 - SSH *See* secure shell 165
 - uninstalling 356
 - upgrading 336
 - upgrading or reinstalling 338, 343

Q

- quorum disks
 - creating 254
 - DS4000 series 283
 - setting 127

R

- rebalancing
 - managed disks (MDisks) access 177, 255
- recovering
 - offline virtual disks (VDisks) 135
 - using CLI 196
 - software automatically 233

- reinstalling
 - SAN Volume Controller Console 338, 343
- related information xvii
- relationships, Global Mirror
 - creating 149
 - modifying 150
 - overview 50
 - starting 149, 153
 - stopping 150, 153
- relationships, Metro Mirror
 - creating 149
 - modifying 150
 - overview 50
 - starting 149, 153
 - stopping 150, 153
- relationships, Mirror
 - deleting 151
 - modifying 152
- remote service 30
- removing
 - master console 356
 - master console software 352
 - nodes 122, 208
 - SAN Volume Controller Console 354
 - storage controllers 251
 - using the CLI (command-line interface) 253
- renaming
 - a Mirror consistency group 152
 - disk controller systems 248
 - I/O groups 123
 - managed disks 126
 - MDisks 126
 - nodes 121
- replacing
 - nodes 118, 193
 - SSH private key 161
 - SSH public key 161
- requirements
 - upgrading SAN Volume Controller Console software 338
 - upgrading the SAN Volume Controller Console software 338
 - Web browsers 96
- rescanning the fibre-channel network 125
- reserved pool of volumes 369
- resetting
 - SSH fingerprint for a cluster 162
- running
 - cluster maintenance procedure 155
 - PuTTY plink utility 168

S

- SAN (storage area network)
 - configuring 68
- SAN fabric
 - configuring 65
 - example 66
 - overview 66
- SAN Volume Controller
 - adding to cluster 172
 - configuring nodes 73
 - Console
 - banner 104

SAN Volume Controller (continued)

- Console (continued)
 - layout 103
 - master console 29
 - portfolio 104
 - post-installation tasks 349
 - starting 105
 - task bar 104
 - user interfaces 6
 - work area 105
- console configuring 98
- copying using PuTTY scp 226
- features 5
- front panel password 172
- minimum requirements 5
- overview 1
- properties 176
- renaming 121
- replacing nodes 193
- shutting down 125
- software upgrade problems 233, 234
- upgrading software 228
- upgrading software
 - automatically 227
 - upgrading software using the CLI 230
- SAN Volume Controller Console
 - backing up configuration file 217
 - banner 104
 - installing 89
 - launching the Web application 109
 - layout 103
 - starting 109
 - storing SSH keys 99
 - troubleshooting 359
 - uninstalling 354
 - upgrading 336
 - upgrading or reinstalling 338, 343
- scanning
 - fibre-channel network 177, 255
 - rebalancing MDisk access 177, 255
- SDD 5
- secure shell
 - adding keys 160
 - client system
 - overview 165
 - creating keys 98
 - PuTTY 99
 - secure shell (SSH)
 - client system
 - issuing CLI commands from 168
 - preparing to issue CLI commands 167
- Secure Shell (SSH)
 - adding keys 210
 - keys
 - replacing key pair 160
 - replacing private key 161
 - listing keys 210
 - managing keys 159
 - overview 29
 - PuTTY 29
 - resetting fingerprint 162
- service
 - actions, UPS 14
 - remote via Assist On-site 30

- setting
 - cluster date 110
 - cluster features
 - using the CLI (command-line interface) 171
 - cluster time 110
 - using the CLI (command-line interface) 170
 - copy direction 151
 - features
 - using the CLI (command-line interface) 171
 - quorum disks 127
 - time
 - using the CLI (command-line interface) 170
- settings
 - error notification 211
 - language 213
- shortcut keys 403
- shrinking
 - VDisks 133
- shutting down
 - clusters 124
 - nodes 125
- SNMP traps 30, 155, 211
- software
 - automatic recovery 233
 - automatic upgrades 227
 - copying using PuTTY scp 226
 - installing 225
 - manual recovery 234
 - multipathing 6
 - option, master console 29
 - recovering automatically 233
 - recovering manually 234
 - upgrading 225, 228
 - upgrading automatically 227
 - upgrading SAN Volume Controller Console 338
 - upgrading using the command-line interface (CLI) 230
- software upgrades
 - recovering 233, 234
- software, upgrading
 - disruptive
 - using the CLI (command-line interface) 231
 - using the CLI (command-line interface) 225
- SSH
 - /see also Secure Shell (SSH) 29
- SSH (secure shell)
 - adding keys 210
 - client system
 - issuing CLI commands from 168
 - preparing to issue CLI commands 167
 - listing keys 210
 - resetting fingerprint 162
- SSH (Secure Shell)
 - keys
 - replacing key pair 160
 - replacing private key 161
 - managing keys 159
- SSH *See* secure shell 98, 160, 166

- SSH keys
 - replacing 161
 - storing 99
- starting
 - FlashCopy
 - consistency groups 147
 - mappings 145
 - Global Mirror
 - consistency groups 149, 153
 - relationships 149, 153
 - Metro Mirror
 - consistency groups 149, 153
 - relationships 149, 153
- status
 - of 2145 UPS-1U 13
 - of node 9
 - of node ports 115
- stopping
 - FlashCopy
 - mappings 146
 - Global Mirror
 - consistency groups 150, 153
 - relationships 150, 153
 - Metro Mirror
 - consistency groups 150, 153
 - relationships 150, 153
 - Remote Copy
 - consistency groups 148
- storage area network (SAN)
 - configuring 68
- storage controllers
 - adding
 - using the 249
 - using the CLI (command-line interface) 250
 - removing 251
 - using the CLI (command-line interface) 253
- storage subsystems
 - servicing 255
- storing SSH keys 99
- strategy
 - software upgrade
 - using the CLI (command-line interface) 225
- subnet mask
 - changing 210
- subsystem device driver (SDD) 5
- support
 - multipathing software 6
 - Web sites *xxi*
- switches
 - Brocade 75
 - Cisco 75
 - configuring 74
 - director class 77
 - fibre-channel 74
 - McData 75
 - mixing 75
 - operating over long distances 83
 - zoning 80
- system requirements, IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software 364

T

- text emphasis *xvii*
- time
 - setting
 - using the CLI (command-line interface) 170
- Tivoli SAN Manager
 - uninstalling 353
- Tivoli SAN Manager Agent
 - uninstalling 353
- trademarks 407
- transceiver 78
- troubleshooting
 - event notification e-mail 30
 - master console 360
 - Microsoft Windows boot problems 360
 - SAN Volume Controller Console 359
 - using Assist On-site (AOS) 30

U

- uninstalling
 - DS4000 Storage Manager Client (FAST Storage Manager Client) 354
 - IBM Director 353
 - master console 356
 - master console software 352
 - PuTTY 356
 - SAN Volume Controller Console 354
 - Tivoli SAN Manager 353
 - Tivoli SAN Manager Agent 353
- uninterruptible power supply
 - 2145 UPS-1U
 - configuration 13
 - operation 14
 - overview 13
 - configuration 13
 - operation 14
- upgrading
 - master console software 336
 - PuTTY 336
 - SAN Volume Controller Console 336, 338, 343
 - software 225, 228
 - software automatically 227
 - software using the command-line interface (CLI) 230
- upgrading software
 - disruptive
 - using the CLI (command-line interface) 231
 - strategy
 - using the CLI (command-line interface) 225
- user interfaces 6
- using
 - 2145 UPS-1U 13

V

- VDisk (virtual disk)
 - expanding 200
- VDisk (virtual disks)
 - determining mappings 189

- VDisks (virtual disks)
 - image mode 137
- VDisks (virtual disks)
 - converting
 - from image mode to managed mode 137, 206
 - creating 131, 181
 - deleting 136
 - determining name of 189
 - expanding 200
 - MDisks (managed disks)
 - relationships 190
 - migrating 117, 139, 192, 206
 - moving offline 198
 - offline 136
 - overview 22
 - recovering from offline 135
 - using CLI 196
 - shrinking 133
- viewing
 - clusters
 - feature logs 155
- virtual disk-to-host mapping
 - description 26
- virtual disks (VDisks) 200
 - converting
 - from image mode to managed mode 137, 206
 - creating 131
 - determining mappings 189
 - determining name of 189
 - image mode 137
 - managed disks (MDisks)
 - relationships 190
 - migrating 117, 132, 139, 192
 - moving offline 198
 - nodes 73
 - offline 136
 - overview 22
 - recovering from offline 135
 - using CLI 196
 - shrinking 133
- virtualization
 - asymmetric 4
 - overview 2
 - symmetric 4

W

- Web browsers
 - configuring 96
 - requirements 96
- Web sites xxi
- who should read this guide xv

Z

- zoning
 - considerations for Global Mirror 82
 - considerations for Metro Mirror 82
 - controllers 78
 - guidelines 78
 - hosts 78
 - overview 80

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

IBM System Storage SAN Volume Controller
Software Installation and Configuration Guide
Version 4.2.1

Publication No. SC23-6628-01

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

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

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

Comments:

Thank you for your support.

Submit your comments using one of these channels:

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

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

Name

Address

Company or Organization

Phone No.

E-mail address



Fold and Tape

Please do not staple

Fold and Tape



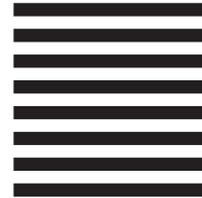
NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

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: 31P1048

Printed in USA

SC23-6628-01



(1P) P/N: 31P1048



Spine information:



IBM System Storage SAN Volume
Controller

SAN Volume Controller Software Installation and
Configuration Guide

Version 4.2.1