

IBM® TotalStorage® Virtualization Family  
SAN Volume Controller™



# Configuration Guide

*Version 1 Release 1*



IBM® TotalStorage® Virtualization Family  
SAN Volume Controller™



# Configuration Guide

*Version 1 Release 1*

**Note**

Before using this information and the product it supports, read the information in "Notices" on page 281.

**First Edition (June 2003)**

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

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

---

# Contents

<b>About this guide . . . . .</b>	<b>vii</b>
Who should use this guide . . . . .	vii
Numbering conventions . . . . .	vii
How to send your comments . . . . .	vii
<b>Part 1. Overview . . . . .</b>	<b>1</b>
<b>Chapter 1. Overview of the Storage Area</b>	<b>3</b>
<b>Network Volume Controller . . . . .</b>	<b>3</b>
Virtualization . . . . .	4
Asymmetric virtualization . . . . .	5
Symmetric virtual storage networks . . . . .	6
<b>Chapter 2. Object overview . . . . .</b>	<b>7</b>
Object descriptions . . . . .	8
Nodes and Clusters . . . . .	9
Nodes . . . . .	9
Configuration node . . . . .	10
Cluster . . . . .	11
I/O groups and Uninterruptible Power Supply . . . . .	12
I/O group . . . . .	13
UPS and power domains . . . . .	14
Controllers and managed disks . . . . .	16
Disk controllers . . . . .	16
Managed disk (MDisk) . . . . .	16
Managed disk groups and virtual disks (VDisks) . . . . .	18
Managed disk (MDisk) group . . . . .	18
Virtual disk (VDisk) . . . . .	21
Hosts and virtual (VDisk) mappings . . . . .	23
Host objects . . . . .	23
<b>Chapter 3. Copy Services . . . . .</b>	<b>27</b>
FlashCopy . . . . .	27
FlashCopy applications . . . . .	27
FlashCopy mappings . . . . .	28
Stand Alone Mappings . . . . .	30
FlashCopy Consistency groups . . . . .	30
FlashCopy consistency group Overview . . . . .	30
Dependent writes . . . . .	30
Operations on consistency groups . . . . .	31
Limits . . . . .	31
FlashCopy indirection layer . . . . .	31
Grains and the FlashCopy bitmap . . . . .	32
Source and target reads . . . . .	32
Writes to the source or target . . . . .	32
Limits . . . . .	32
FlashCopy mapping events . . . . .	33
Background copy . . . . .	34
Host considerations for FlashCopy integrity . . . . .	35
Remote Copy . . . . .	37
Synchronous Remote Copy . . . . .	37
Remote Copy partnerships . . . . .	37
Remote Copy relationships . . . . .	38
Remote Copy consistency groups . . . . .	38
<b>Chapter 4. Configuration rules and requirements . . . . .</b>	<b>41</b>
Configuration rules . . . . .	42
RAID controllers . . . . .	42
Host bus adapters (HBAs) . . . . .	43
Nodes . . . . .	44
Power . . . . .	44
Fibre channel switches . . . . .	44
Configuration requirements . . . . .	46
Maximum configurations . . . . .	47
<b>Part 2. Preparing to configure the SAN Volume Controller . . . . .</b>	<b>49</b>
<b>Chapter 5. Create cluster from the front panel . . . . .</b>	<b>51</b>
<b>Chapter 6. Master console security overview . . . . .</b>	<b>55</b>
Overview of passwords . . . . .	55
<b>Chapter 7. Master console overview . . . . .</b>	<b>57</b>
Configuring the master console . . . . .	57
Secure Shell (SSH) configuration . . . . .	58
Overview of configuring the Secure Shell (SSH) . . . . .	59
Generating an SSH key pair using the SSH client called PuTTY . . . . .	60
Storing keys in the SAN Volume Controller Console software . . . . .	61
Adding SSH keys for hosts other than the master console . . . . .	62
Replacing the SSH key pair . . . . .	62
Replace the client SSH private key known to the SAN Volume Controller software . . . . .	63
Configuring the PuTTY session for the command-line interface . . . . .	64
Configuring the Tivoli SAN Manager (TSanM) . . . . .	64
Starting the TSanM . . . . .	65
Setting up Remote Support . . . . .	66
Enhanced remote support configuration . . . . .	67
Changing the master console hostname . . . . .	67
Overview of IBM Director . . . . .	67
IBM Director . . . . .	68
SAN Volume Controller Call-Home overview using IBM Director . . . . .	69
Configuring IBM Director for the SAN Volume Controller error notification and Call-Home . . . . .	69
Setting up your E-mail notification for the SAN Volume Controller . . . . .	71
Upgrading software on the master console . . . . .	72
Troubleshooting master console problems . . . . .	73

<b>Part 3. SAN Volume Controller Console</b>	<b>75</b>
<b>Chapter 8. Getting started with the SAN Volume Controller Console</b>	<b>77</b>
Accessing the SAN Volume Controller Console	77
SAN Volume Controller Console layout	78
SAN Volume Controller Console banner area	78
SAN Volume Controller task bar	78
SAN Volume Controller Console portfolio	79
SAN Volume Controller Console work area	79
Upgrading the SAN Volume Controller Console software	79
<b>Chapter 9. Overview of creating a cluster using the SAN Volume Controller Console</b>	<b>81</b>
Creating a cluster on the SAN Volume Controller Console	81
Prerequisites for creating a cluster using the SAN Volume Controller Console	81
Creating a cluster using the SAN Volume Controller Console	82
Launching the SAN Volume Controller Application	90
Setting the cluster time using the SAN Volume Controller Console	92
Displaying cluster properties using the SAN Volume Controller Console	93
<b>Chapter 10. Scenario: typical usage for the SAN Volume Controller Console</b>	<b>95</b>
Adding nodes to the cluster	96
Displaying node properties using the SAN Volume Controller Console	100
Create managed disk (MDisk) groups	101
Create virtual disks (VDisks)	103
Creating host objects	104
Create VDisk-to-host mappings	105
Create a FlashCopy consistency group	106
Create a FlashCopy mapping	106
<b>Chapter 11. Advanced function FlashCopy overview</b>	<b>109</b>
Starting FlashCopy mappings	109
Stopping FlashCopy mappings	109
Deleting FlashCopy mappings	109
Starting FlashCopy consistency groups	110
Stopping FlashCopy consistency groups	110
Deleting consistency groups	110
<b>Chapter 12. Advanced functions overview for the SAN Volume Controller Console</b>	<b>111</b>
Guidelines for MDisk group creation	111
Determining a nodes WWPNs using the SAN Volume Controller Console	112
Determining a storage controller name from its SAN Volume Controller name	112
Determining the relationship between VDisks and MDisks using the SAN Volume Controller Console	112
Determining the relationship between MDisks and RAID arrays or LUNs using the SAN Volume Controller Console	113
Increasing the size of your cluster using the SAN Volume Controller Console	113
Adding a node to increase the size of your cluster	114
Migrating a VDisk to a new I/O group	114
Replacing a faulty node in the cluster using the SAN Volume Controller	115
Recovering from offline VDisks after a node or an I/O group failed	116
Replacing an HBA in a host using the SAN Volume Controller Console	117
Adding a new storage controller to a running configuration	118
Removing a storage controller	119
Expanding a VDisk using the SAN Volume Controller Console	120
Shrinking a VDisk using the SAN Volume Controller Console	122
Migrating VDisks between MDisk groups	123
Creating image mode virtual disks	124
Advanced function Remote Copy overview	125
Advanced function cluster overview	126
Deleting nodes from a cluster	126
Setting up the cluster features using the SAN Volume Controller Console	127
Enabling the cluster maintenance procedure using the SAN Volume Controller Console	128
Maintaining passwords using the SAN Volume Controller Console	129
Adding subsequent SSH public keys to the SAN Volume Controller	129
Setting up error notifications using the SAN Volume Controller Console	131
Resetting a cached SAN Volume Controller cluster SSH host fingerprint on the SAN Volume Controller Console overview	132
Resetting an refused SSH key relationship between the SAN Volume Controller Console and the SAN Volume Controller cluster overview	133
Modifying IP addresses using the SAN Volume Controller Console	133
Listing log or dump files using the SAN Volume Controller Console	134
Changing the language settings	134
Viewing the feature log using the SAN Volume Controller Console	135
Analyzing the error log using the SAN Volume Controller Console	135
Shutting down a cluster	136
<b>Part 4. Command-Line Interface</b>	<b>139</b>
<b>Chapter 13. Getting started with the Command-Line Interface</b>	<b>141</b>

Preparing the SSH client system overview . . . . .	142	Migrating a VDisk between I/O groups using the CLI . . . . .	188																																																																				
Preparing the SSH client system to issue CLI commands . . . . .	142	Creating image mode virtual disks using the CLI . . . . .	189																																																																				
Issuing CLI commands from a PuTTY SSH Client system . . . . .	143	Advanced function FlashCopy and Remote Copy overview for CLI . . . . .	190																																																																				
Running the PuTTY and plink utilities . . . . .	144	Advanced function cluster overview using the CLI . . . . .	190																																																																				
Configuring the cluster using the CLI . . . . .	146	Deleting a node from a cluster using the CLI . . . . .	190																																																																				
Setting the cluster time using the CLI . . . . .	147	Enabling the cluster maintenance procedure using the CLI . . . . .	191																																																																				
Reviewing and setting the cluster features using the CLI . . . . .	147	Maintaining passwords using the CLI . . . . .	192																																																																				
Displaying cluster properties using the CLI . . . . .	148	Maintaining SSH keys using the CLI . . . . .	192																																																																				
Modifying passwords using the CLI. . . . .	148	Setting up error notifications using the CLI . . . . .	193																																																																				
<b>Chapter 14. Scenario: typical usage for the command-line interface . . . . .</b>	<b>149</b>	Modifying IP addresses using the CLI . . . . .	193																																																																				
Adding nodes to the cluster using the CLI . . . . .	150	Listing log or dump files using the CLI . . . . .	193																																																																				
Displaying node properties using the CLI . . . . .	154	Changing the language setting using the CLI . . . . .	194																																																																				
Discovering MDisks using the CLI . . . . .	155	Viewing the feature log using the CLI . . . . .	195																																																																				
Create managed disk (MDisk) groups using the CLI. . . . .	156	Analyze the error log using the CLI . . . . .	195																																																																				
Adding MDisks to MDisk groups using the CLI . . . . .	159	Shutting down a cluster using the CLI . . . . .	196																																																																				
Create virtual disks (VDisks) . . . . .	159	<b>Part 5. Software upgrade strategy using the CLI and the SAN Volume Controller Console . . . . .</b>	<b>199</b>																																																																				
Creating host objects using the CLI . . . . .	162	<b>Chapter 16. Disruptive software upgrade. . . . .</b>	<b>201</b>	Create VDisk-to-host mappings using the CLI . . . . .	163	<b>Chapter 17. Upgrading software using the SAN Volume Controller Console . . . . .</b>	<b>203</b>	Create FlashCopy mappings using the CLI . . . . .	164	<b>Chapter 18. Error counts . . . . .</b>	<b>207</b>	Create a FlashCopy consistency group and add mappings using the CLI. . . . .	165	<b>Chapter 19. Automatic upgrade. . . . .</b>	<b>209</b>	Prepare and trigger a FlashCopy Consistency Group using the CLI . . . . .	166	<b>Chapter 20. Automatic recovery from upgrade problems: . . . . .</b>	<b>211</b>	<b>Chapter 15. Advanced functions overview for the CLI . . . . .</b>	<b>169</b>	<b>Chapter 21. Secure copy (scp) overview . . . . .</b>	<b>213</b>	Determining a nodes WWPNs using the CLI . . . . .	169	<b>Chapter 22. Installing the upgrade using the CLI . . . . .</b>	<b>215</b>	Determining a storage controller name from its SAN Volume Controller name . . . . .	170	<b>Chapter 23. Accessible CLI commands during the upgrade process . . . . .</b>	<b>217</b>	Determining the VDisk name from the vpath number on the host . . . . .	170	<b>Chapter 24. Manual recovery from software upgrade problems . . . . .</b>	<b>219</b>	Determining the host that a VDisk is mapped to . . . . .	171	<b>Part 6. Other configurations . . . . .</b>	<b>221</b>	Determining the relationship between VDisks and MDisks using the CLI . . . . .	171	<b>Chapter 25. Configuring disk controllers . . . . .</b>	<b>223</b>	Determining the relationship between MDisks and RAID arrays or LUNs using the CLI. . . . .	173	Configuring a balanced storage subsystem . . . . .	224	Increasing the size of your cluster using the CLI . . . . .	173	Adding a node to increase the size of your cluster using the CLI. . . . .	174	Migrating a VDisk to a new I/O group. . . . .	174	Replacing a faulty node in the cluster using the CLI. . . . .	175	Recovering from offline VDisks after a node or an I/O group failed using the CLI . . . . .	176	Replacing an HBA in a host using the CLI. . . . .	179	Adding a new storage controller to a running configuration using the CLI . . . . .	180	Removing a storage controller using the CLI . . . . .	181	Expanding a VDisk using the CLI . . . . .	182	Shrinking a VDisk using the CLI . . . . .	185	Migrating extents using the CLI . . . . .	185	Migrating VDisks between MDisk groups using the CLI. . . . .	187
<b>Chapter 16. Disruptive software upgrade. . . . .</b>	<b>201</b>																																																																						
Create VDisk-to-host mappings using the CLI . . . . .	163	<b>Chapter 17. Upgrading software using the SAN Volume Controller Console . . . . .</b>	<b>203</b>	Create FlashCopy mappings using the CLI . . . . .	164	<b>Chapter 18. Error counts . . . . .</b>	<b>207</b>	Create a FlashCopy consistency group and add mappings using the CLI. . . . .	165	<b>Chapter 19. Automatic upgrade. . . . .</b>	<b>209</b>	Prepare and trigger a FlashCopy Consistency Group using the CLI . . . . .	166	<b>Chapter 20. Automatic recovery from upgrade problems: . . . . .</b>	<b>211</b>	<b>Chapter 15. Advanced functions overview for the CLI . . . . .</b>	<b>169</b>	<b>Chapter 21. Secure copy (scp) overview . . . . .</b>	<b>213</b>	Determining a nodes WWPNs using the CLI . . . . .	169	<b>Chapter 22. Installing the upgrade using the CLI . . . . .</b>	<b>215</b>	Determining a storage controller name from its SAN Volume Controller name . . . . .	170	<b>Chapter 23. Accessible CLI commands during the upgrade process . . . . .</b>	<b>217</b>	Determining the VDisk name from the vpath number on the host . . . . .	170	<b>Chapter 24. Manual recovery from software upgrade problems . . . . .</b>	<b>219</b>	Determining the host that a VDisk is mapped to . . . . .	171	<b>Part 6. Other configurations . . . . .</b>	<b>221</b>	Determining the relationship between VDisks and MDisks using the CLI . . . . .	171	<b>Chapter 25. Configuring disk controllers . . . . .</b>	<b>223</b>	Determining the relationship between MDisks and RAID arrays or LUNs using the CLI. . . . .	173	Configuring a balanced storage subsystem . . . . .	224	Increasing the size of your cluster using the CLI . . . . .	173	Adding a node to increase the size of your cluster using the CLI. . . . .	174	Migrating a VDisk to a new I/O group. . . . .	174	Replacing a faulty node in the cluster using the CLI. . . . .	175	Recovering from offline VDisks after a node or an I/O group failed using the CLI . . . . .	176	Replacing an HBA in a host using the CLI. . . . .	179	Adding a new storage controller to a running configuration using the CLI . . . . .	180	Removing a storage controller using the CLI . . . . .	181	Expanding a VDisk using the CLI . . . . .	182	Shrinking a VDisk using the CLI . . . . .	185	Migrating extents using the CLI . . . . .	185	Migrating VDisks between MDisk groups using the CLI. . . . .	187				
<b>Chapter 17. Upgrading software using the SAN Volume Controller Console . . . . .</b>	<b>203</b>																																																																						
Create FlashCopy mappings using the CLI . . . . .	164	<b>Chapter 18. Error counts . . . . .</b>	<b>207</b>	Create a FlashCopy consistency group and add mappings using the CLI. . . . .	165	<b>Chapter 19. Automatic upgrade. . . . .</b>	<b>209</b>	Prepare and trigger a FlashCopy Consistency Group using the CLI . . . . .	166	<b>Chapter 20. Automatic recovery from upgrade problems: . . . . .</b>	<b>211</b>	<b>Chapter 15. Advanced functions overview for the CLI . . . . .</b>	<b>169</b>	<b>Chapter 21. Secure copy (scp) overview . . . . .</b>	<b>213</b>	Determining a nodes WWPNs using the CLI . . . . .	169	<b>Chapter 22. Installing the upgrade using the CLI . . . . .</b>	<b>215</b>	Determining a storage controller name from its SAN Volume Controller name . . . . .	170	<b>Chapter 23. Accessible CLI commands during the upgrade process . . . . .</b>	<b>217</b>	Determining the VDisk name from the vpath number on the host . . . . .	170	<b>Chapter 24. Manual recovery from software upgrade problems . . . . .</b>	<b>219</b>	Determining the host that a VDisk is mapped to . . . . .	171	<b>Part 6. Other configurations . . . . .</b>	<b>221</b>	Determining the relationship between VDisks and MDisks using the CLI . . . . .	171	<b>Chapter 25. Configuring disk controllers . . . . .</b>	<b>223</b>	Determining the relationship between MDisks and RAID arrays or LUNs using the CLI. . . . .	173	Configuring a balanced storage subsystem . . . . .	224	Increasing the size of your cluster using the CLI . . . . .	173	Adding a node to increase the size of your cluster using the CLI. . . . .	174	Migrating a VDisk to a new I/O group. . . . .	174	Replacing a faulty node in the cluster using the CLI. . . . .	175	Recovering from offline VDisks after a node or an I/O group failed using the CLI . . . . .	176	Replacing an HBA in a host using the CLI. . . . .	179	Adding a new storage controller to a running configuration using the CLI . . . . .	180	Removing a storage controller using the CLI . . . . .	181	Expanding a VDisk using the CLI . . . . .	182	Shrinking a VDisk using the CLI . . . . .	185	Migrating extents using the CLI . . . . .	185	Migrating VDisks between MDisk groups using the CLI. . . . .	187								
<b>Chapter 18. Error counts . . . . .</b>	<b>207</b>																																																																						
Create a FlashCopy consistency group and add mappings using the CLI. . . . .	165	<b>Chapter 19. Automatic upgrade. . . . .</b>	<b>209</b>	Prepare and trigger a FlashCopy Consistency Group using the CLI . . . . .	166	<b>Chapter 20. Automatic recovery from upgrade problems: . . . . .</b>	<b>211</b>	<b>Chapter 15. Advanced functions overview for the CLI . . . . .</b>	<b>169</b>	<b>Chapter 21. Secure copy (scp) overview . . . . .</b>	<b>213</b>	Determining a nodes WWPNs using the CLI . . . . .	169	<b>Chapter 22. Installing the upgrade using the CLI . . . . .</b>	<b>215</b>	Determining a storage controller name from its SAN Volume Controller name . . . . .	170	<b>Chapter 23. Accessible CLI commands during the upgrade process . . . . .</b>	<b>217</b>	Determining the VDisk name from the vpath number on the host . . . . .	170	<b>Chapter 24. Manual recovery from software upgrade problems . . . . .</b>	<b>219</b>	Determining the host that a VDisk is mapped to . . . . .	171	<b>Part 6. Other configurations . . . . .</b>	<b>221</b>	Determining the relationship between VDisks and MDisks using the CLI . . . . .	171	<b>Chapter 25. Configuring disk controllers . . . . .</b>	<b>223</b>	Determining the relationship between MDisks and RAID arrays or LUNs using the CLI. . . . .	173	Configuring a balanced storage subsystem . . . . .	224	Increasing the size of your cluster using the CLI . . . . .	173	Adding a node to increase the size of your cluster using the CLI. . . . .	174	Migrating a VDisk to a new I/O group. . . . .	174	Replacing a faulty node in the cluster using the CLI. . . . .	175	Recovering from offline VDisks after a node or an I/O group failed using the CLI . . . . .	176	Replacing an HBA in a host using the CLI. . . . .	179	Adding a new storage controller to a running configuration using the CLI . . . . .	180	Removing a storage controller using the CLI . . . . .	181	Expanding a VDisk using the CLI . . . . .	182	Shrinking a VDisk using the CLI . . . . .	185	Migrating extents using the CLI . . . . .	185	Migrating VDisks between MDisk groups using the CLI. . . . .	187												
<b>Chapter 19. Automatic upgrade. . . . .</b>	<b>209</b>																																																																						
Prepare and trigger a FlashCopy Consistency Group using the CLI . . . . .	166	<b>Chapter 20. Automatic recovery from upgrade problems: . . . . .</b>	<b>211</b>	<b>Chapter 15. Advanced functions overview for the CLI . . . . .</b>	<b>169</b>	<b>Chapter 21. Secure copy (scp) overview . . . . .</b>	<b>213</b>	Determining a nodes WWPNs using the CLI . . . . .	169	<b>Chapter 22. Installing the upgrade using the CLI . . . . .</b>	<b>215</b>	Determining a storage controller name from its SAN Volume Controller name . . . . .	170	<b>Chapter 23. Accessible CLI commands during the upgrade process . . . . .</b>	<b>217</b>	Determining the VDisk name from the vpath number on the host . . . . .	170	<b>Chapter 24. Manual recovery from software upgrade problems . . . . .</b>	<b>219</b>	Determining the host that a VDisk is mapped to . . . . .	171	<b>Part 6. Other configurations . . . . .</b>	<b>221</b>	Determining the relationship between VDisks and MDisks using the CLI . . . . .	171	<b>Chapter 25. Configuring disk controllers . . . . .</b>	<b>223</b>	Determining the relationship between MDisks and RAID arrays or LUNs using the CLI. . . . .	173	Configuring a balanced storage subsystem . . . . .	224	Increasing the size of your cluster using the CLI . . . . .	173	Adding a node to increase the size of your cluster using the CLI. . . . .	174	Migrating a VDisk to a new I/O group. . . . .	174	Replacing a faulty node in the cluster using the CLI. . . . .	175	Recovering from offline VDisks after a node or an I/O group failed using the CLI . . . . .	176	Replacing an HBA in a host using the CLI. . . . .	179	Adding a new storage controller to a running configuration using the CLI . . . . .	180	Removing a storage controller using the CLI . . . . .	181	Expanding a VDisk using the CLI . . . . .	182	Shrinking a VDisk using the CLI . . . . .	185	Migrating extents using the CLI . . . . .	185	Migrating VDisks between MDisk groups using the CLI. . . . .	187																
<b>Chapter 20. Automatic recovery from upgrade problems: . . . . .</b>	<b>211</b>																																																																						
<b>Chapter 15. Advanced functions overview for the CLI . . . . .</b>	<b>169</b>	<b>Chapter 21. Secure copy (scp) overview . . . . .</b>	<b>213</b>	Determining a nodes WWPNs using the CLI . . . . .	169	<b>Chapter 22. Installing the upgrade using the CLI . . . . .</b>	<b>215</b>	Determining a storage controller name from its SAN Volume Controller name . . . . .	170	<b>Chapter 23. Accessible CLI commands during the upgrade process . . . . .</b>	<b>217</b>	Determining the VDisk name from the vpath number on the host . . . . .	170	<b>Chapter 24. Manual recovery from software upgrade problems . . . . .</b>	<b>219</b>	Determining the host that a VDisk is mapped to . . . . .	171	<b>Part 6. Other configurations . . . . .</b>	<b>221</b>	Determining the relationship between VDisks and MDisks using the CLI . . . . .	171	<b>Chapter 25. Configuring disk controllers . . . . .</b>	<b>223</b>	Determining the relationship between MDisks and RAID arrays or LUNs using the CLI. . . . .	173	Configuring a balanced storage subsystem . . . . .	224	Increasing the size of your cluster using the CLI . . . . .	173	Adding a node to increase the size of your cluster using the CLI. . . . .	174	Migrating a VDisk to a new I/O group. . . . .	174	Replacing a faulty node in the cluster using the CLI. . . . .	175	Recovering from offline VDisks after a node or an I/O group failed using the CLI . . . . .	176	Replacing an HBA in a host using the CLI. . . . .	179	Adding a new storage controller to a running configuration using the CLI . . . . .	180	Removing a storage controller using the CLI . . . . .	181	Expanding a VDisk using the CLI . . . . .	182	Shrinking a VDisk using the CLI . . . . .	185	Migrating extents using the CLI . . . . .	185	Migrating VDisks between MDisk groups using the CLI. . . . .	187																				
<b>Chapter 21. Secure copy (scp) overview . . . . .</b>	<b>213</b>																																																																						
Determining a nodes WWPNs using the CLI . . . . .	169	<b>Chapter 22. Installing the upgrade using the CLI . . . . .</b>	<b>215</b>	Determining a storage controller name from its SAN Volume Controller name . . . . .	170	<b>Chapter 23. Accessible CLI commands during the upgrade process . . . . .</b>	<b>217</b>	Determining the VDisk name from the vpath number on the host . . . . .	170	<b>Chapter 24. Manual recovery from software upgrade problems . . . . .</b>	<b>219</b>	Determining the host that a VDisk is mapped to . . . . .	171	<b>Part 6. Other configurations . . . . .</b>	<b>221</b>	Determining the relationship between VDisks and MDisks using the CLI . . . . .	171	<b>Chapter 25. Configuring disk controllers . . . . .</b>	<b>223</b>	Determining the relationship between MDisks and RAID arrays or LUNs using the CLI. . . . .	173	Configuring a balanced storage subsystem . . . . .	224	Increasing the size of your cluster using the CLI . . . . .	173	Adding a node to increase the size of your cluster using the CLI. . . . .	174	Migrating a VDisk to a new I/O group. . . . .	174	Replacing a faulty node in the cluster using the CLI. . . . .	175	Recovering from offline VDisks after a node or an I/O group failed using the CLI . . . . .	176	Replacing an HBA in a host using the CLI. . . . .	179	Adding a new storage controller to a running configuration using the CLI . . . . .	180	Removing a storage controller using the CLI . . . . .	181	Expanding a VDisk using the CLI . . . . .	182	Shrinking a VDisk using the CLI . . . . .	185	Migrating extents using the CLI . . . . .	185	Migrating VDisks between MDisk groups using the CLI. . . . .	187																								
<b>Chapter 22. Installing the upgrade using the CLI . . . . .</b>	<b>215</b>																																																																						
Determining a storage controller name from its SAN Volume Controller name . . . . .	170	<b>Chapter 23. Accessible CLI commands during the upgrade process . . . . .</b>	<b>217</b>	Determining the VDisk name from the vpath number on the host . . . . .	170	<b>Chapter 24. Manual recovery from software upgrade problems . . . . .</b>	<b>219</b>	Determining the host that a VDisk is mapped to . . . . .	171	<b>Part 6. Other configurations . . . . .</b>	<b>221</b>	Determining the relationship between VDisks and MDisks using the CLI . . . . .	171	<b>Chapter 25. Configuring disk controllers . . . . .</b>	<b>223</b>	Determining the relationship between MDisks and RAID arrays or LUNs using the CLI. . . . .	173	Configuring a balanced storage subsystem . . . . .	224	Increasing the size of your cluster using the CLI . . . . .	173	Adding a node to increase the size of your cluster using the CLI. . . . .	174	Migrating a VDisk to a new I/O group. . . . .	174	Replacing a faulty node in the cluster using the CLI. . . . .	175	Recovering from offline VDisks after a node or an I/O group failed using the CLI . . . . .	176	Replacing an HBA in a host using the CLI. . . . .	179	Adding a new storage controller to a running configuration using the CLI . . . . .	180	Removing a storage controller using the CLI . . . . .	181	Expanding a VDisk using the CLI . . . . .	182	Shrinking a VDisk using the CLI . . . . .	185	Migrating extents using the CLI . . . . .	185	Migrating VDisks between MDisk groups using the CLI. . . . .	187																												
<b>Chapter 23. Accessible CLI commands during the upgrade process . . . . .</b>	<b>217</b>																																																																						
Determining the VDisk name from the vpath number on the host . . . . .	170	<b>Chapter 24. Manual recovery from software upgrade problems . . . . .</b>	<b>219</b>	Determining the host that a VDisk is mapped to . . . . .	171	<b>Part 6. Other configurations . . . . .</b>	<b>221</b>	Determining the relationship between VDisks and MDisks using the CLI . . . . .	171	<b>Chapter 25. Configuring disk controllers . . . . .</b>	<b>223</b>	Determining the relationship between MDisks and RAID arrays or LUNs using the CLI. . . . .	173	Configuring a balanced storage subsystem . . . . .	224	Increasing the size of your cluster using the CLI . . . . .	173	Adding a node to increase the size of your cluster using the CLI. . . . .	174	Migrating a VDisk to a new I/O group. . . . .	174	Replacing a faulty node in the cluster using the CLI. . . . .	175	Recovering from offline VDisks after a node or an I/O group failed using the CLI . . . . .	176	Replacing an HBA in a host using the CLI. . . . .	179	Adding a new storage controller to a running configuration using the CLI . . . . .	180	Removing a storage controller using the CLI . . . . .	181	Expanding a VDisk using the CLI . . . . .	182	Shrinking a VDisk using the CLI . . . . .	185	Migrating extents using the CLI . . . . .	185	Migrating VDisks between MDisk groups using the CLI. . . . .	187																																
<b>Chapter 24. Manual recovery from software upgrade problems . . . . .</b>	<b>219</b>																																																																						
Determining the host that a VDisk is mapped to . . . . .	171	<b>Part 6. Other configurations . . . . .</b>	<b>221</b>	Determining the relationship between VDisks and MDisks using the CLI . . . . .	171	<b>Chapter 25. Configuring disk controllers . . . . .</b>	<b>223</b>	Determining the relationship between MDisks and RAID arrays or LUNs using the CLI. . . . .	173	Configuring a balanced storage subsystem . . . . .	224	Increasing the size of your cluster using the CLI . . . . .	173	Adding a node to increase the size of your cluster using the CLI. . . . .	174	Migrating a VDisk to a new I/O group. . . . .	174	Replacing a faulty node in the cluster using the CLI. . . . .	175	Recovering from offline VDisks after a node or an I/O group failed using the CLI . . . . .	176	Replacing an HBA in a host using the CLI. . . . .	179	Adding a new storage controller to a running configuration using the CLI . . . . .	180	Removing a storage controller using the CLI . . . . .	181	Expanding a VDisk using the CLI . . . . .	182	Shrinking a VDisk using the CLI . . . . .	185	Migrating extents using the CLI . . . . .	185	Migrating VDisks between MDisk groups using the CLI. . . . .	187																																				
<b>Part 6. Other configurations . . . . .</b>	<b>221</b>																																																																						
Determining the relationship between VDisks and MDisks using the CLI . . . . .	171	<b>Chapter 25. Configuring disk controllers . . . . .</b>	<b>223</b>	Determining the relationship between MDisks and RAID arrays or LUNs using the CLI. . . . .	173	Configuring a balanced storage subsystem . . . . .	224	Increasing the size of your cluster using the CLI . . . . .	173	Adding a node to increase the size of your cluster using the CLI. . . . .	174	Migrating a VDisk to a new I/O group. . . . .	174	Replacing a faulty node in the cluster using the CLI. . . . .	175	Recovering from offline VDisks after a node or an I/O group failed using the CLI . . . . .	176	Replacing an HBA in a host using the CLI. . . . .	179	Adding a new storage controller to a running configuration using the CLI . . . . .	180	Removing a storage controller using the CLI . . . . .	181	Expanding a VDisk using the CLI . . . . .	182	Shrinking a VDisk using the CLI . . . . .	185	Migrating extents using the CLI . . . . .	185	Migrating VDisks between MDisk groups using the CLI. . . . .	187																																								
<b>Chapter 25. Configuring disk controllers . . . . .</b>	<b>223</b>																																																																						
Determining the relationship between MDisks and RAID arrays or LUNs using the CLI. . . . .	173	Configuring a balanced storage subsystem . . . . .	224	Increasing the size of your cluster using the CLI . . . . .	173	Adding a node to increase the size of your cluster using the CLI. . . . .	174	Migrating a VDisk to a new I/O group. . . . .	174	Replacing a faulty node in the cluster using the CLI. . . . .	175	Recovering from offline VDisks after a node or an I/O group failed using the CLI . . . . .	176	Replacing an HBA in a host using the CLI. . . . .	179	Adding a new storage controller to a running configuration using the CLI . . . . .	180	Removing a storage controller using the CLI . . . . .	181	Expanding a VDisk using the CLI . . . . .	182	Shrinking a VDisk using the CLI . . . . .	185	Migrating extents using the CLI . . . . .	185	Migrating VDisks between MDisk groups using the CLI. . . . .	187																																												
Configuring a balanced storage subsystem . . . . .	224																																																																						
Increasing the size of your cluster using the CLI . . . . .	173																																																																						
Adding a node to increase the size of your cluster using the CLI. . . . .	174																																																																						
Migrating a VDisk to a new I/O group. . . . .	174																																																																						
Replacing a faulty node in the cluster using the CLI. . . . .	175																																																																						
Recovering from offline VDisks after a node or an I/O group failed using the CLI . . . . .	176																																																																						
Replacing an HBA in a host using the CLI. . . . .	179																																																																						
Adding a new storage controller to a running configuration using the CLI . . . . .	180																																																																						
Removing a storage controller using the CLI . . . . .	181																																																																						
Expanding a VDisk using the CLI . . . . .	182																																																																						
Shrinking a VDisk using the CLI . . . . .	185																																																																						
Migrating extents using the CLI . . . . .	185																																																																						
Migrating VDisks between MDisk groups using the CLI. . . . .	187																																																																						

Data migration on an existing FAStT installation which contains partitions . . . . .	227
Configuring FAStT disk controllers for the storage server . . . . .	228
Configuring FAStT disk controllers for the storage manager . . . . .	230
Configuring the Enterprise Storage Server (ESS) . . . . .	230
<b>Chapter 26. Overview about zoning a switch . . . . .</b>	<b>233</b>
Zoning a switch . . . . .	233
Zoning considerations for Remote Copy . . . . .	235
Zoning rules and example . . . . .	236
Switch operations over long distances . . . . .	238
<b>Appendix A. Installing the IBM TotalStorage SAN Volume Controller Console for Windows . . . . .</b>	<b>241</b>
Installation overview for the SAN Volume Controller Console . . . . .	241
SAN Volume Controller Console hardware installation requirements . . . . .	242
SAN Volume Controller Console workstation space requirements . . . . .	242
SAN Volume Controller Console software installation requirements . . . . .	243
Installing the SAN Volume Controller Console in graphical mode . . . . .	244
Installing the SAN Volume Controller Console in unattended (silent) mode . . . . .	253
Verifying the Windows services associated with the SAN Volume Controller Console . . . . .	256
Post Installation Tasks - Getting started using the SAN Volume Controller Console . . . . .	257
Removing the SAN Volume Controller Console . . . . .	259
<b>Appendix B. Valid combinations of FlashCopy and Remote Copy functions . . . . .</b>	<b>263</b>
<b>Appendix C. Setting up SNMP traps . . . . .</b>	<b>265</b>
<b>Appendix D. Configuring IBM Director overview . . . . .</b>	<b>267</b>
Setting up an Event Action Plan . . . . .	267
Setting up an e-mail . . . . .	268
Setting up an e-mail user notification . . . . .	269
<b>Appendix E. Object types . . . . .</b>	<b>271</b>
<b>Appendix F. Event codes . . . . .</b>	<b>273</b>
Information event codes . . . . .	273
Configuration event codes . . . . .	274
<b>Appendix G. Accessibility . . . . .</b>	<b>279</b>
<b>Notices . . . . .</b>	<b>281</b>
<b>Trademarks . . . . .</b>	<b>283</b>
<b>Related information . . . . .</b>	<b>285</b>
<b>Ordering IBM publications . . . . .</b>	<b>287</b>
<b>Glossary . . . . .</b>	<b>289</b>
<b>Index . . . . .</b>	<b>295</b>

---

## About this guide

This guide provides information that helps you configure and use the IBM® TotalStorage® Virtualization Family SAN Volume Controller™. This guide describes the configuration tools, both command-line and Web based, that you can use to define, expand, and maintain the storage of the IBM TotalStorage SAN Volume Controller.

**Related topics:**

- “Who should use this guide”
- “Numbering conventions”

---

## Who should use this guide

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

**Related topics:**

- “About this guide”
- “Numbering conventions”

---

## Numbering conventions

This topic describes the numbering conventions used in this guide and product.

In this guide, one gigabyte (GB) is equal to:

1073741824 bytes

1 MB = 1048576 KB and 1 KB = 1048576 bytes.

**Related topics:**

- “About this guide”
- “Who should use this guide”

---

## 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 SAN Volume Controller 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](mailto: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 or fax

Fill out the Readers' Comments form (RCF) at the back of this book. Return it by mail or fax (1-800-426-6209), or give it to an IBM representative. If the RCF has been removed, you can address your comments to:

International Business Machines Corporation  
RCF Processing Department  
M86/050  
5600 Cottle Road  
San Jose, CA 95193-0001  
U.S.A.

**Related topics:**

- “Related information” on page 285

---

## **Part 1. Overview**

This part provides an overview of the SAN Volume Controller. More specifically, this part provides the following information:

- Chapter 1, “Overview of the Storage Area Network Volume Controller”, on page 3
- “Object descriptions” on page 8
- “FlashCopy” on page 27
- “Remote Copy” on page 37
- Chapter 4, “Configuration rules and requirements”, on page 41



---

## Chapter 1. Overview of the Storage Area Network Volume Controller

The IBM TotalStorage Storage Area Network Volume Controller (see Figure 1) is a storage area network controller. It is a rack-mounted unit that can be installed in a standard Electrical Industries Association (EIA) 19-inch rack.

The SAN Volume Controller provides 4 GB cache and 4 Fibre Channel ports (2 FC adapters)

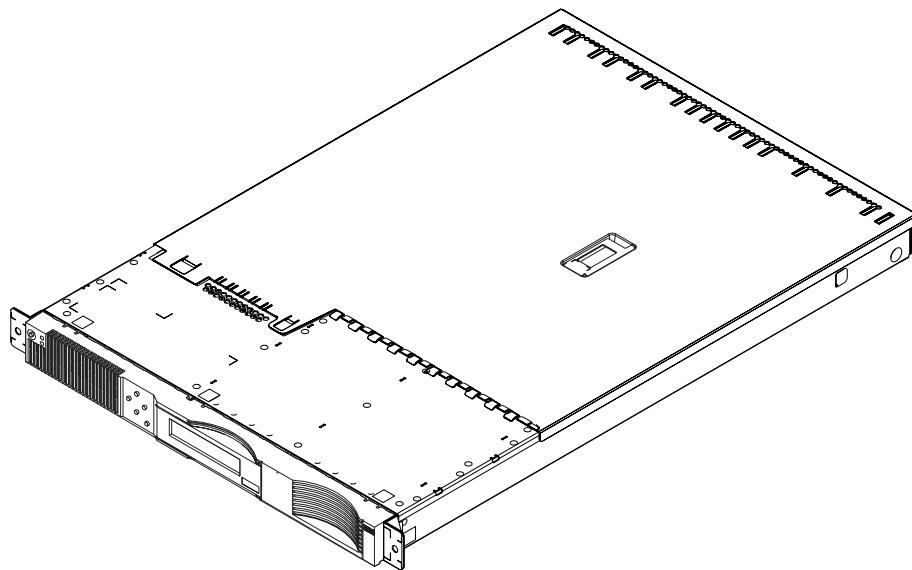


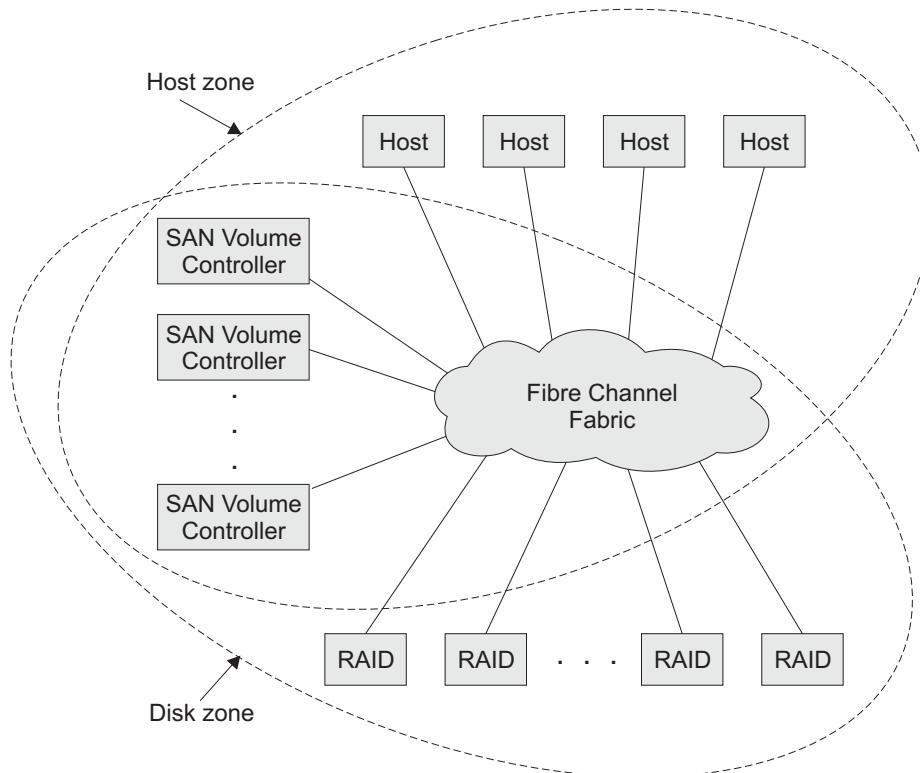
Figure 1. The SAN Volume Controller

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

Each SAN Volume Controller is a *node* that is, it is an end point of a *link*, or it is a junction that is common to two or more links of the SAN. Nodes are grouped into *clusters* of up to four nodes. The cluster is managed as a set, and provides a single point of control for the user for configuration and service activities. For I/O operations, the nodes are grouped into pairs. Each pair has the responsibility to serve I/O for a particular set of VDisks. If one SAN Volume Controller; of a pair fails or is removed, failover occurs to the other SAN Volume Controller. The clusters are attached to the SAN fabric. Also attached to the fabric are RAID controllers and host systems.

The fabric contains two or more distinct zones: a host zone and a disk zone. In the host zone, the host systems can see and address the nodes. In the disk zone, the nodes can see the disk drives. Host systems are not allowed to operate on the disk drives directly; all data transfer occurs through the nodes. Figure 2 on page 4 shows an example of a storage system that is using a SAN Volume Controller.

Several host systems are connected to a SAN fabric. A cluster of SAN Volume controllers are connected to the same fabric and presents virtual disks to the host systems. These virtual disks are created from disks that are presented by the Redundant Array of Independent Disks (RAID) controllers.



*Figure 2. Example of a SAN Volume Controller in a fabric*

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

## Virtualization

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

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

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

#### Types of virtualization:

Virtualization can be performed either asymmetrically or symmetrically:

##### Asymmetric

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

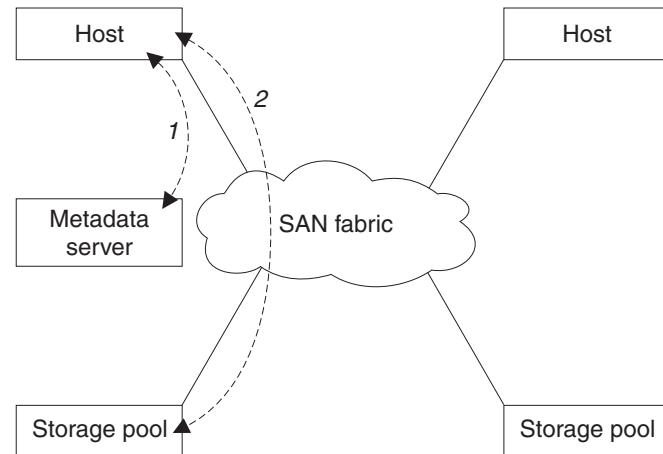
##### Symmetric

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

#### Related topics:

- Chapter 25, “Configuring disk controllers”, on page 223

## Asymmetric virtualization



*Figure 3. Asymmetrical virtualization*

In asymmetric virtual storage networks, the data flow, (2) in the figure above, is separated from the control flow, (1). The meta-data server contains all the mapping and locking tables while the storage devices contain only data.

Because the flow of control is separated from the flow of data, I/O operations can use the full bandwidth of the SAN. A separate network or SAN link is used for control purposes. There are disadvantages, however, to asymmetric virtualization.

For one, data is at risk to increased security exposures and the control network must be protected with a firewall. In addition, meta-data can become very complicated when files are distributed across several devices. Moreover, each host

that accesses the SAN must know how to access and interpret the meta-data. Specific device driver or agent software must therefore be running on each of these hosts. Finally, the meta-data server cannot run advanced functions, such as caching or copy services, because it only knows about the meta-data, and not about the data itself.

---

## Symmetric virtual storage networks

This topic provides overview information about symmetric virtual storage networks.

In symmetric virtual storage networks, see Figure 4, data and control both flow over the same path. Because the separation of the control from the data occurs in the data path, the storage can be pooled under the control of the virtualization engine. The virtualization engine performs the logical to physical mapping.

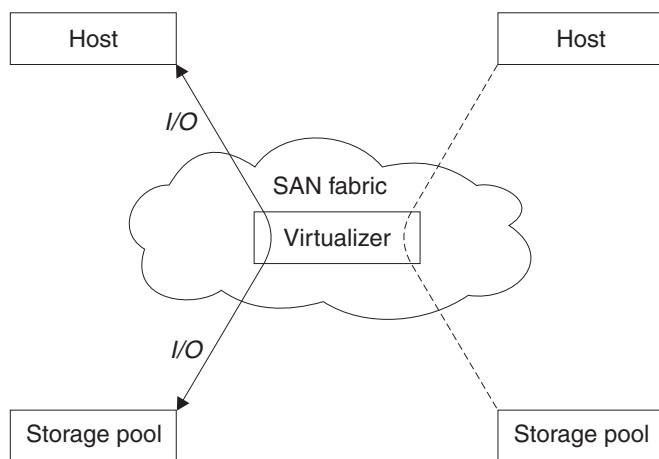


Figure 4. Symmetrical virtualization

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

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

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

---

## Chapter 2. Object overview

This topic provides overview information about object descriptions.

The SAN Volume Controller is based on the following virtualization concepts which are discussed more fully later in this chapter.

A **node** is a single SAN Volume Controller. Nodes are deployed in pairs to make up a cluster. A cluster may have either 1 or 2 node pairs in it. Each pair of nodes is known as an **I/O group**. Each node may be in only one I/O group.

**Virtual disks (Vdisks)** are logical disks that are presented to the SAN by nodes. Virtual disks are also associated with an I/O group. The nodes in the I/O group provide access to the virtual disks in the I/O group. When an application server performs I/O to a virtual disk, it has the choice of accessing the virtual disk via either of the nodes in the I/O group. As each I/O group only has two nodes, the distributed cache the SAN Volume Controller provides is only be 2-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. In such situations, the UPS will maintain power to the nodes while the contents of the distributed cache are dumped to an internal drive.

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

Each managed disk is divided up into a number of **extents** (default size is 16MB) which are numbered, from 0, sequentially from the start to the end of the managed disk.

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

At any one time, a single node in the cluster 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 world wide 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.

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

---

## Object descriptions

This chapter describes the objects in a SAN Volume Controller environment and the relationships between these objects.

The SAN Volume Controller provides block level aggregation and volume management for disk storage within the SAN. In simpler terms, this means that the SAN Volume Controller manages a number of back-end storage controllers and maps the physical storage within those controllers into logical disk images that can be seen by application servers and workstations in the SAN. The SAN is zoned in such a way that the application servers cannot see the back-end physical storage; this prevents any possible conflict between the SAN Volume Controller and the application servers both trying to manage the back-end storage.

The following figure is a representation of the objects described in this chapter and their logical positioning in a virtualized system.

**Note:** Virtual disk to host mappings have not been shown to simplify the figure.

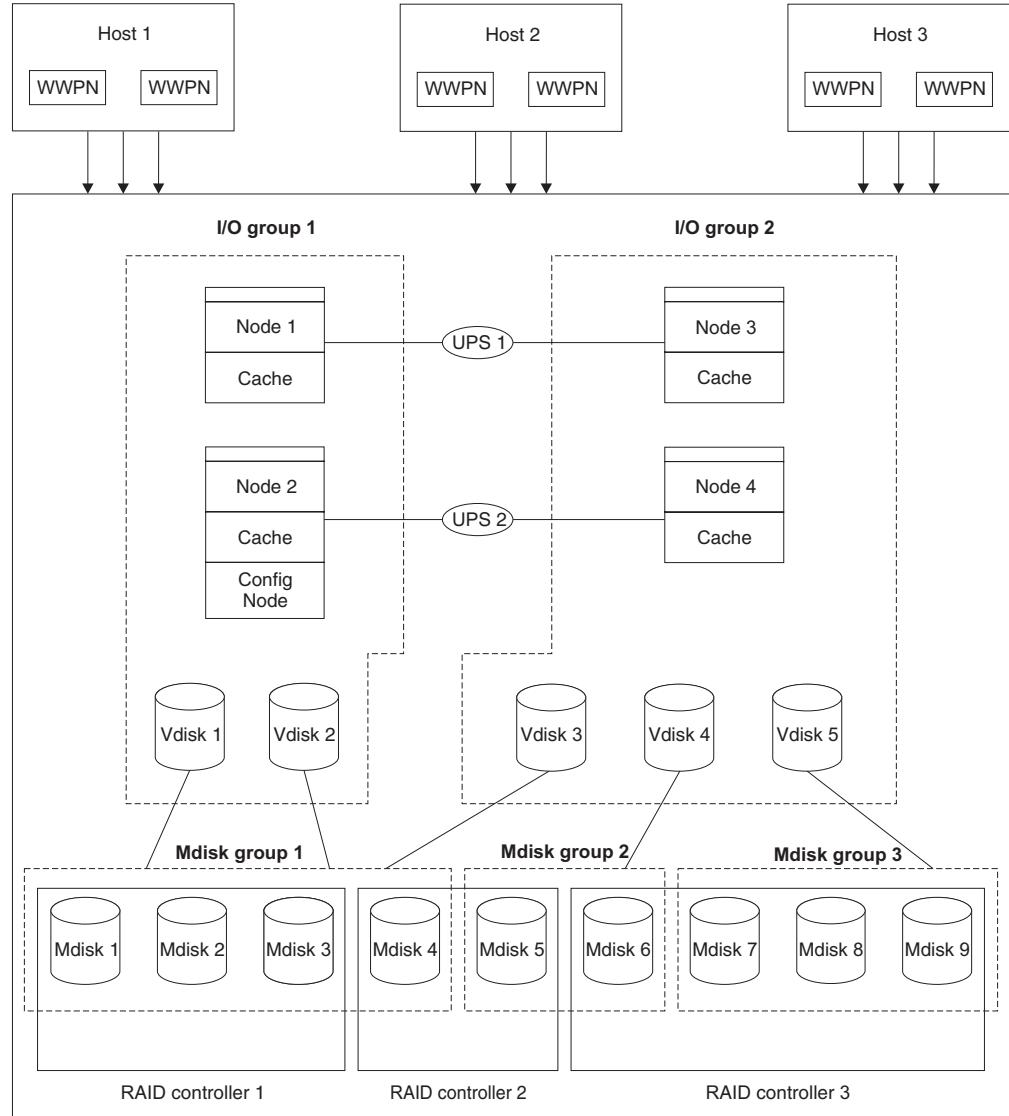


Figure 5. Virtualization

Items such as Virtual Disks, Managed Disks, and Managed Disk Groups are known as objects. It is important that you understand what these objects are, how you use them, and their relationships to one another, so that you can effectively configure and use the SAN Volume Controller.

## Nodes and Clusters

A node is a single SAN Volume Controller. Nodes are deployed in pairs to make up a cluster. A node can be designated as a special node such as configuration node or boss node. Each node in the cluster contains a copy of the cluster state.

### Nodes

A **node** is a single SAN Volume Controller. A node is an end point of a *link*, or a junction that is common to two or more links of the SAN. Nodes are deployed in pairs, for redundancy, to make up a cluster. A cluster may have 1 or 2 node pairs in it. Each pair of nodes is known as an I/O group. Each node may be in *only* one I/O group.

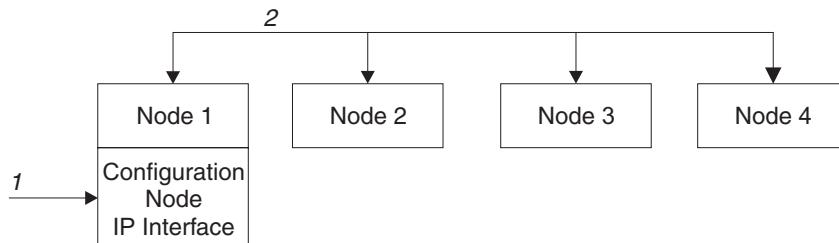
At any one time, a single node in the cluster is used to manage configuration activity. This configuration node manages a cache of the configuration information that describes the cluster configuration and provides a focal point for configuration commands. If the configuration node fails, another node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

The status of a node consists of five settings. The following table describes the different states of a node:

*Table 1. Node status*

Status	Description
<b>Adding</b>	The node has been added to the cluster but has not yet synchronized with the cluster state. (See Note.)
<b>Deleting</b>	The node is in the process of being deleted from the cluster.
<b>Online</b>	The node is operational, assigned to a cluster and has access to the Fibre Channel SAN fabric.
<b>Offline</b>	The node is not-operational. The node has been assigned to a cluster but is not available on the Fibre Channel SAN fabric. Run the Directed Maintenance Procedures to determine the problem (refer to <i>SAN Volume Controller: Service Guide</i> for more information).
<b>Pending</b>	The node is transitioning between states, and in a few seconds will move to one of the other states.

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



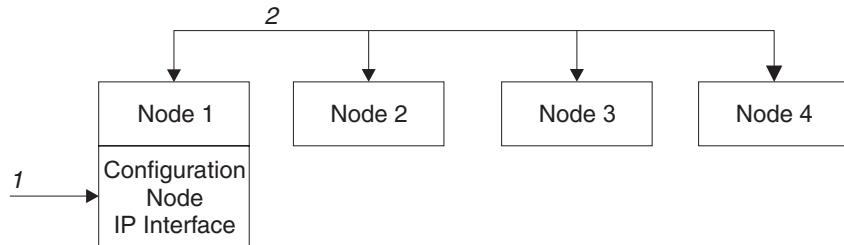
*Figure 6. Configuration node*

## Configuration node

When the cluster is created, the system automatically assigns one node as the *configuration node*. The configuration node binds to the cluster IP address, and provides the configuration interface to the cluster.

If the configuration node fails, the cluster chooses a new configuration node. This action is called configuration node failover. 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.

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



*Figure 7. Configuration node*

## Cluster

A cluster is a group of one or two node pairs. Therefore, you can assign up to four SAN Volume Controller nodes to one cluster. All configuration and service is performed at the cluster level. Some service actions can be performed at node level, but all configuration is replicated across all nodes in the cluster. Because configuration is performed at the cluster level, an IP address is assigned to the cluster instead of to each of the nodes.

### Cluster state and the boss node:

The cluster state holds all configuration and internal cluster data for the cluster. This cluster state information is held in non-volatile memory. If the mainline power fails, the two UPSs maintain the internal power long enough for the cluster state information to be stored on the internal SCSI disk drive of each node. The read and write cache information is also held in non-volatile memory. Similarly if the power fails to a node, the cached data is written to the internal SCSI disk. It is also held in memory and is stored on the internal SCSI disk drives of the nodes in the I/O group that are surfacing that information.

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

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

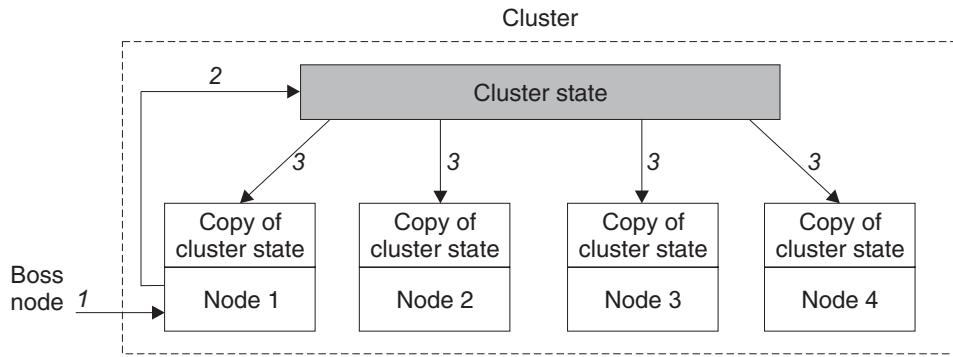


Figure 8. Cluster, nodes, and cluster state.

#### Cluster operation and quorum disks:

The cluster must contain at least half of its nodes to function. That is, when the cluster is formed and becomes stable, only half of the nodes must remain functional for the cluster to continue operating. Note, however, that the cluster can survive more than half the nodes failing as long as the cluster has restabilized between failures. For example, a 4 node cluster that loses 2 nodes would continue to function with the remaining 2 nodes. If, some time later, another node fails, the remaining node would continue to function.

A tie-break situation can occur if exactly half the nodes in a cluster fail at the same time, or if the cluster is divided so that exactly half the nodes in the cluster cannot communicate with the other half. For example, in a cluster of 4 nodes, if any two nodes fail at the same time, or any two cannot communicate with the other two, a tie-break exists and must be resolved.

The cluster automatically chooses three managed disks to be **quorum disks** and assigns them quorum indices of 0, 1, or 2. One of these disks is used to settle a tie-break condition. One of these disks is known to all nodes as the chosen quorum disk at any given time. If this chosen disk becomes offline or unavailable the cluster will select one of the other two disks.

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

You can change the assignment of quorum disks at any time.

## 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 may be in *only* one I/O group.

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

## I/O group

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

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

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

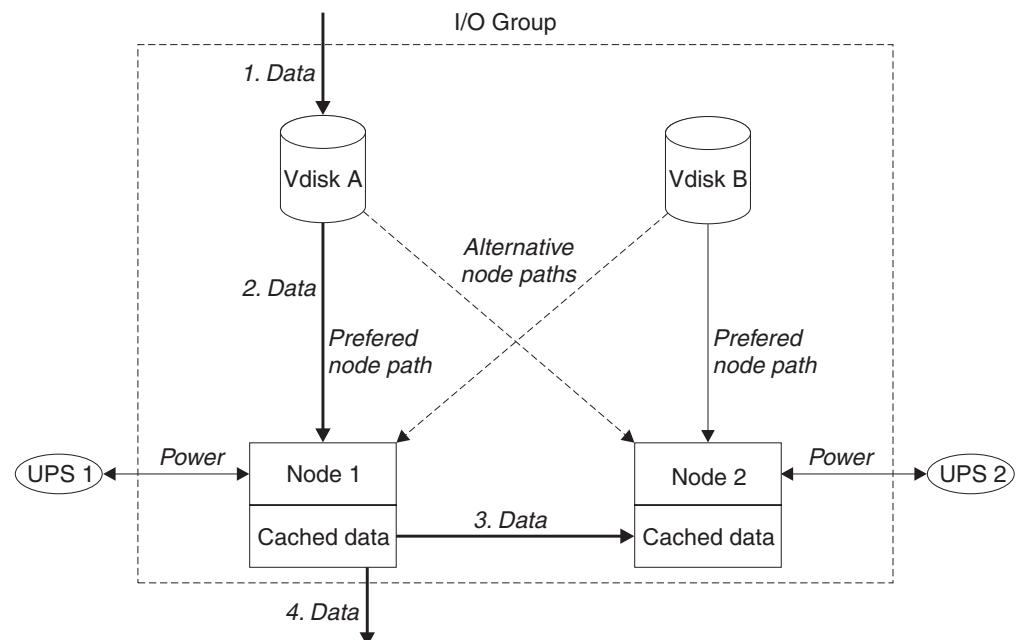


Figure 9. I/O group and uninterruptible power supply

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

If only one node is assigned to an I/O group, or a node has failed in an I/O group, the cache goes into write-through mode. Therefore, any writes for the

virtual disks that are assigned to this I/O group are not cached, it is sent directly to the storage device. If both nodes in an I/O group go offline, the virtual disks that are assigned to the I/O group cannot be accessed.

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

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

## UPS and power domains

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

It is very important that the two nodes in the I/O group are not both connected to the same power domain. Each SAN Volume Controller of an I/O group must be connected to a different uninterruptible power supply. This configuration ensures that the cache and cluster state information is protected against the failure of the uninterruptible power supply or of the mainline power source. To protect against failure of the mainline power source, it is important that the uninterruptible power supply units input power is supplied from different mainline power domains.

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

The figure below shows a cluster of four nodes, with two I/O groups and two uninterruptible power supply units.

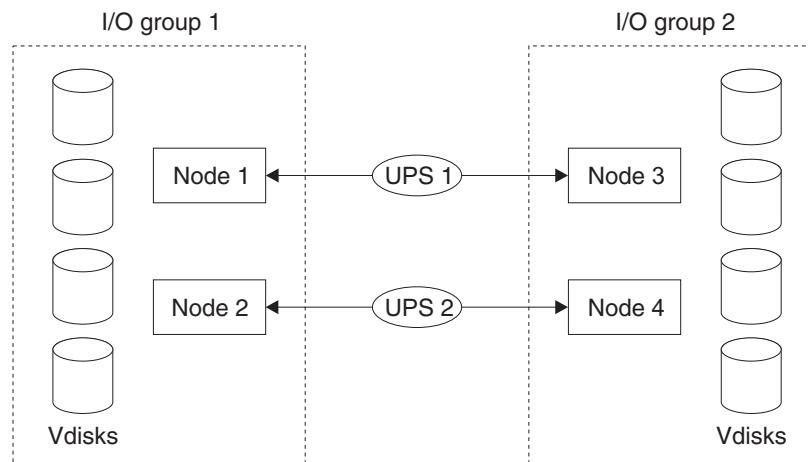
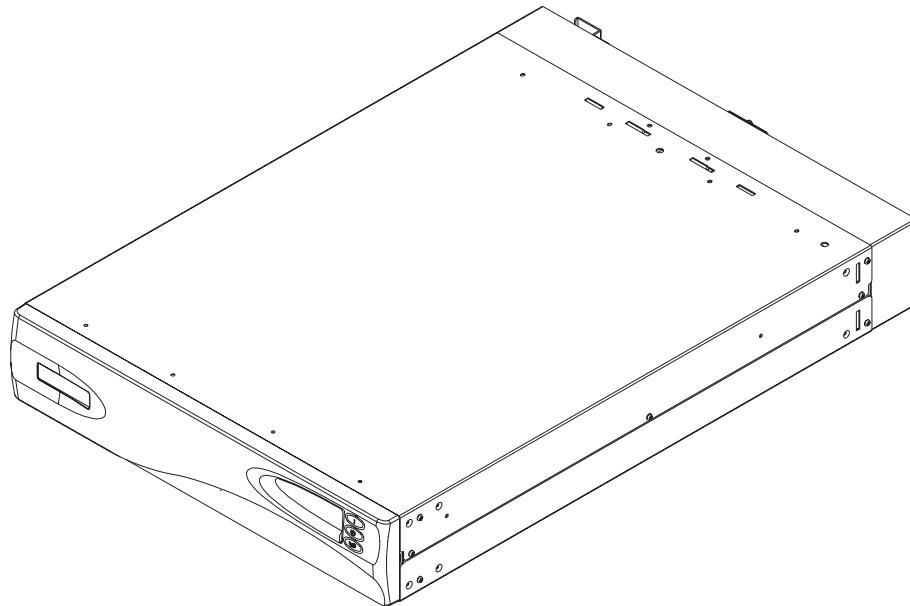


Figure 10. I/O groups and uninterruptible power supply relationship

## Uninterruptible power supply overview

The uninterruptible power supply protects the SAN Volume Controller from power failures, power sags, power surges, and line noise.



*Figure 11. Uninterruptible power supply*

To provide redundancy and concurrent maintenance, the SAN Volume Controller must be installed in pairs. Each SAN Volume Controller of a pair must be connected to a different uninterruptible power supply. Each uninterruptible power supply can support four SAN Volume Controllers. Also, each uninterruptible power supply of a pair must be connected to a separate electrical independent power source (if possible) to reduce the chance of input power failure at both uninterruptible power supply units.

Each uninterruptible power supply includes power (line) cords that will connect the uninterruptible power supply to either a rack power distribution unit (PDU), if one exists, or to an external power source.

The uninterruptible power supply, when connected to the SAN Volume Controller, requires a UL (or equivalent) 250V, 15A circuit breaker.

**Attention:** Under normal circumstances, if power is disconnected from the uninterruptible power supply (UPS), the SAN Volume controller(s) connected to that uninterruptible power supply perform a power down sequence. This operation which saves the configuration and cache data to an internal disk in the SAN Volume Controller typically takes about 3 minutes at which time power is removed from the output of the uninterruptible power supply. In the event of a delay in the completion of the power down sequence, the uninterruptible power supply output power will be removed after 5 minutes from the time power was disconnected to the uninterruptible power supply. As this operation is controlled by the SAN Volume controller, a UPS not connected to an active SAN Volume Controller will NOT shutoff within the 5 minute required period. In the case of an emergency, you will need to manually shutdown the uninterruptible power supply by pushing the uninterruptible power supply power off button. It is important to note that data integrity could be compromised by pushing the uninterruptible power supply power off button.

---

## Controllers and managed disks

The nodes in a cluster see the storage presented by back-end controllers as a number of disks, known as managed disks. The SAN Volume Controller does not attempt to provide recovery from physical disk failures within the back-end controllers. A managed disk is usually, but not necessarily, a RAID array.

### Disk controllers

The nodes in the cluster are connected to one or more Fibre Channel SAN fabrics. Also attached to the SAN fabric are one or more back-end controllers. These are usually RAID controllers. The controllers are the devices that provide the physical storage that the nodes in the cluster detect as managed disks.

The supported RAID controllers are detected by the cluster and reported by the user interfaces. The cluster can also determine which managed disks each controller is presenting, and can provide a view of managed disks filtered by controller. This allows you to associate the managed disks with the RAID arrays that the controller presents.

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

Back-end controllers present storage to other devices on the SAN. The physical storage associated with a back-end controller is normally configured into RAID arrays which provide recovery from physical disk failures. Some back-end controllers also allow physical storage to be configured as RAID-0 arrays (striping) or as JBODs; however, this does not provide protection against a physical disk failure and with virtualization can lead to the failure of many virtual disk.

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

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

### Managed disk (MDisk)

A managed disk (MDisk) is a logical disk (typically a RAID array or partition thereof) that a storage controller has offered on the SAN fabric that the nodes in the cluster are attached. A managed disk may therefore consist of multiple physical disks that are presented as a single logical disk to the SAN. A managed disk always provides usable blocks of physical storage to the cluster whether it has a one to one correspondence with a physical disk or not.

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

is a property of managed disk groups. When an MDisk is added to an MDisk group, the size of the extents the MDisk will be broken into depends on the attribute of the Mdisk group it has been added to.

Managed disks have associated access modes. The mode determines how the cluster will use the disk. The modes are:

#### Unmanaged

Unmanaged mode is the default mode for all newly discovered logical units of storage (for example, RAID arrays) that a storage controller has offered to the SAN. These logical units are MDisks that have not yet been assigned to a managed disk group.

#### Managed

Managed mode is the standard mode for a managed disk when it has been assigned to a managed disk group. The process of assigning a managed disk to a group automatically changes the mode to managed mode. In managed mode, the extents that are available on the managed disk can be used to create virtual disks.

**Image** Image mode is a special mode for a managed disk. You can use this mode to add a disk to the cluster that has existing data on the disk.

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

The figure shows physical disks and managed disks.

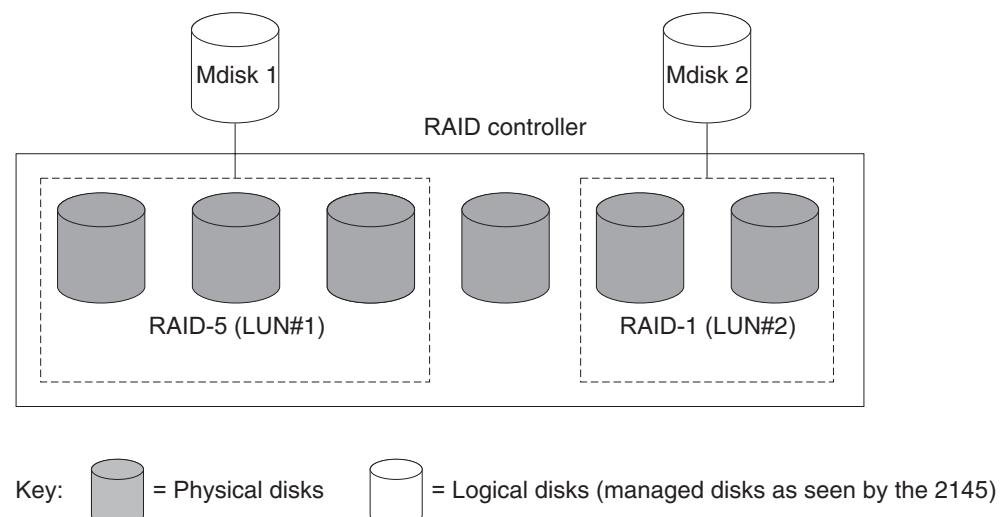


Figure 12. Controllers and MDisks

The status of a managed disk consists of four settings. The following table describes the different states of a managed disk:

Table 2. Managed Disk status

Status	Description
Online	The MDisk can be accessed by all online nodes. That is, all the nodes that are currently working members of the cluster can access this MDisk.

*Table 2. Managed Disk status (continued)*

Degraded	The MDisk cannot be accessed by all the online nodes. That is, one or more (but not all) of the nodes that are currently working members of the cluster cannot access this MDisk. The MDisk may be partially excluded, that is some of the paths to the MDisk (but not all) have been excluded.
Excluded	The MDisk has been excluded from use by the cluster after repeated access errors. Run the Directed Maintenance Procedures to determine the problem. You can reset an MDisk and re-include it in the cluster by running the <code>svctask includemdisk</code> command. See <i>IBM TotalStorage SAN Volume Controller: Command Line Interface Guide</i> for more information.
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.

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

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

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

---

## Managed disk groups and virtual disks (VDisks)

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

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

### Managed disk (MDisk) group

A managed disk (MDisk) group is a group of managed disks. Each group has an extent size that is defined when the group is created. All MDisks that are added to the group are split into extents of that size. Virtual disks (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. 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 these 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 VDisks that are made from the extents that are in the group. If the group is deleted, you cannot recover the mapping that existed between extents that are in the group and the extents that VDisks use. The MDisks that were in the group are returned to unmanaged mode, and can be added to other groups. Because the deletion of a group can cause a loss of data, you must force the delete if VDisks are associated with it.

The status of an MDisk group consists of three settings. The following table describes the different states of a MDisk group:

*Table 3. Managed Disk Group status*

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

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

You should think about the following when creating MDisk groups:

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

#### Extent:

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

You must specify the extent size when you create a new MDisk group. You cannot change the extent size later; it must remain constant throughout the lifetime of the MDisk group. MDisk groups may have different extent sizes however this will place restrictions on the use of data migration. The choice of extent size affects the total amount of storage that can be managed by a SAN Volume Controller Cluster. Table 9 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 virtual disk that is created, using a larger extent size may increase the amount of wasted storage at the end of each virtual disk. Larger extent sizes also reduces the ability of the SAN Volume Controller to distribute sequential I/O workloads across many managed disks and hence may reduce the performance benefits of virtualisation.

*Table 4. Capacities of the cluster given extent size*

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

The figure below shows an MDisk group containing four MDisks.

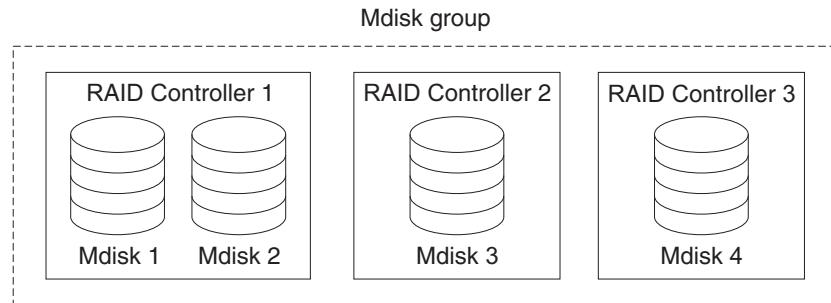


Figure 13. MDisk group

## Virtual disk (VDisk)

VDisks are logical disks that are presented to the SAN by SAN Volume Controller nodes. VDisks, like nodes, are associated with an I/O group. To create a virtual disk, you allocate a set of extents from a managed disk group. The virtualization type determines how the extents are chosen from inside the group. You can create three types of virtual disk:

### Striped

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

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

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

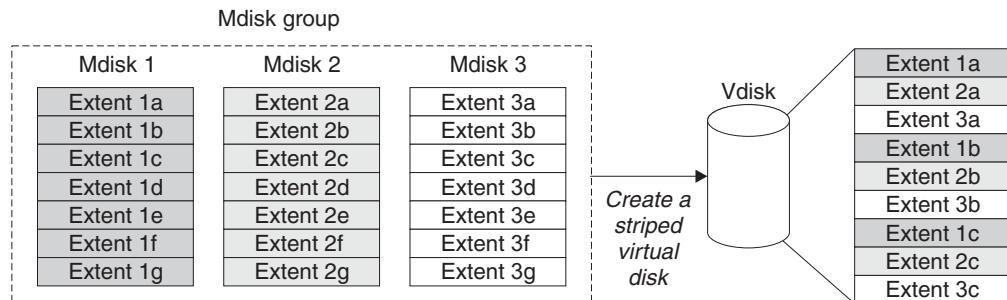


Figure 14. Managed disk groups and VDisks

### Sequential

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

**Image** Image-mode VDisks are special VDisks that have a direct relationship with one managed disk. If you have a RAID array that contains data that you want to merge into the cluster, you can create an image-mode virtual disk.

When you create an image-mode virtual disk, a direct mapping is made between extents that are on the managed disk and extents that are on the virtual disk. The managed disk is not virtualized; meaning, a logical block address (LBA) where  $x$  is arbitrary on the managed disk is the same as LBA  $x$  on the virtual disk.

When you create an image-mode virtual disk, you must assign it to a managed disk group. The extents are managed in the same way as other VDisks. 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. Once you move one or more extents, the virtual disk becomes a real virtualized disk, and the mode of the managed disk changes from image to managed.

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

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

The status of a virtual disk consists of three settings. The following table describes the different states of a virtual disk:

*Table 5. Virtual Disk status*

Status	Description
<b>Online</b>	The virtual disk is online and available if both nodes in the I/O group can access the virtual disk. A single node will only be able to access a VDisk if it can access all the MDisks in the MDisk group associated with the VDisk.
<b>Offline</b>	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.
<b>Degraded</b>	The status of the virtual disk is degraded if one node in the I/O group is online and the other node is either missing or cannot access the virtual disk.

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

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

## Hosts and virtual (VDisk) mappings

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

### Host objects

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

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

When you create a new host object, by typing the **svctask mkhost** command, the configuration interfaces provide a list of unconfigured WWPNs. These WWPNs represent the Fibre Channel ports that the cluster has detected.

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

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.

When you delete a host object, the ports are returned to the unconfigured state, and can be added to a new or existing host. If mappings still exist between the host and the VDisks, you must force the delete operation. This operation deletes the mappings, then the host.

#### **Node Login Counts:**

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

#### **VDisks and host mappings:**

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

presents to the host only those VDisks that are mapped to that host. You must therefore map the VDisks to the hosts that are to access those VDisks.

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

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

The figure below shows two Host servers, and the mappings that exist between the host objects and these VDisks.

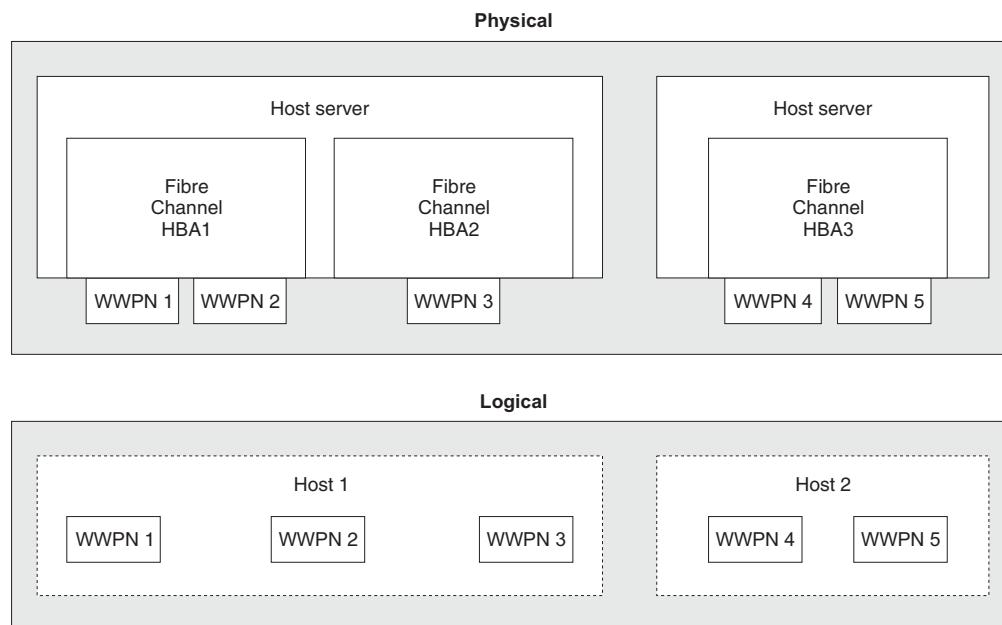


Figure 15. Hosts, WWPNs, and VDisks

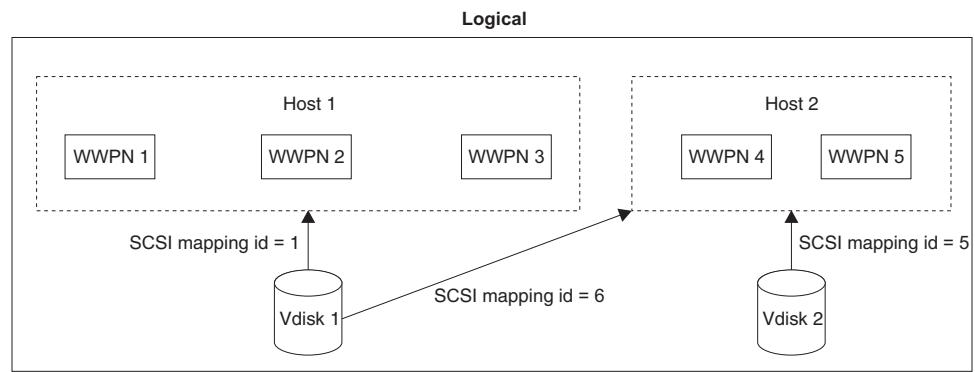


Figure 16. Hosts, WWPNs, VDisks and SCSI mappings



---

## Chapter 3. Copy Services

This topic provides an overview about Copy Services.

The SAN Volume Controller provides Copy Services that enable you to copy a set of virtual disks. You might require such copies for backup purposes, application development activities, or for other uses. There are two methods of making these copies with the SAN Volume Controller. One is called FlashCopy and the other is synchronous Remote Copy. Both methods are described in this section.

---

### FlashCopy

This topic provides an overview about FlashCopy.

**FlashCopy:**

FlashCopy makes a copy of a set of source virtual disks onto a set of target virtual disks. Any data that existed on the target disk is lost and is replaced by the copied data. After the copy operation has been completed, the target virtual disks contain the contents of the source virtual disks as they existed at a single point in time. Although the copy operation takes a finite time to complete, the resulting data on the target is presented in such a way that the copy appears to have occurred immediately. FlashCopy is sometimes described as an instance of a Time-Zero copy ( $T_0$ ) or point-in-time copy technology. Although the FlashCopy operation takes a finite time, this time is several orders of magnitude less than the time which would be required to copy the data using conventional techniques.

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

**Related topics:**

- Appendix B, “Valid combinations of FlashCopy and Remote Copy functions”, on page 263

### FlashCopy applications

This topic provides an overview about FlashCopy applications.

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

Another use of FlashCopy data is in the application testing. It is often very important to test a new version of an application with real business data before the

existing production version of the application is updated or replaced. This testing reduces the risk that the updated application fails because it is not compatible with the actual business data that is in use at the time of the update. Such an application test might require write access to the target data.

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

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

---

## FlashCopy mappings

This topic provides an overview about FlashCopy mappings.

FlashCopy makes an instant copy of a virtual disk at the time it is started. To create a FlashCopy of a virtual disk, you must first create a mapping between the source virtual disk (the disk that is copied) and the target virtual disk (the disk that receives the copy). The source and target must be of equal size. A particular virtual disk can take part in only one mapping; that is, a virtual disk can be the source or target of only one mapping. You cannot, for example, make the target of one mapping the source of another mapping.

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

After a FlashCopy operation has been started, read operations to the source virtual disk operate as normal; that is, they read data from the source virtual disk. If data is written to the source virtual disk, the existing data is copied to the target virtual disk before the new data is written to the source virtual disk. The bitmap is updated to mark that the grain of the source virtual disk has been copied so that later write operations to the same grain do not recopy the data.

It is also possible to read and write to the target virtual disk. Read operations to the target virtual disk use the bitmap to determine whether the grain has been copied or not. If the grain has been copied, data is read from the target virtual disk, otherwise data is read from the source. A write operation to the target virtual disk is similar to a target virtual disk before the new data is written to the target virtual disk. The bitmap is updated to mark that the grain of the source virtual disk has been copied.

The target virtual disk can be accessed after the mapping has been started. Read operations from the target virtual disk are either sent to the source virtual disk (for unchanged data), or read directly from the target virtual disk (for changed data), as described in the bitmap.

When you create a mapping, you specify the background copy rate. This rate determines the priority that is given to the background copy process. If you want to end with a copy of the whole source at the target (so that the mapping can be

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

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

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

The status of flash copy mapping consists of six settings. The following table describes the different states of flash copy mapping:

*Table 6. FlashCopy mapping status*

Status	Description
Idle/Copied	The mapping is idle, or the copy has finished.
Preparing	The mapping is being prepared. The prepare task has been selected.
Prepared	The mapping is prepared. The prepare task has been selected and has completed. The mapping is waiting to be started.
Copying	The mapping is copying. The start task has been selected and the copy is in progress. The <b>Work with Mappings</b> window shows the percentage complete.
Stopped	The mapping has been stopped. The stop task has been selected. The copy can be restarted by issuing the prepare and start tasks again.
Suspended	The mapping has been suspended. The mapping will be suspended if the source VDisk is unavailable, or if the copy bitmap is offline. If the mapping does not return to the Copying state, use the stop task to reset the mapping.

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

## Stand Alone Mappings

This topic provides an overview about stand alone mappings.

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

Stand alone mappings cannot be Started or Prepared.

---

## FlashCopy Consistency groups

This topic provides an overview about consistency groups.

Consistency groups handle the problem in which the using application might have related data that spans multiple virtual disks. FlashCopy must be performed in a way that preserves data integrity across multiple virtual disks. One requirement for preserving the integrity of data being written is to ensure that dependent writes are run in the intended sequence of the application.

## FlashCopy consistency group Overview

This topic provides an overview about FlashCopy consistency groups.

When you copy data from one virtual disk, that data might not include all that you need to enable you to use the copy. For example, the logs for a particular database usually reside on a different disk from the disk that contains the data. In such a condition, you can use a consistency group to ensure that multiple flash copy mappings are all started at the same time.

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. The mappings can have different attributes; for example, different background copy rates.

Mappings that you want to control at an individual level instead of at a consistency group level should not be put into a consistency group. These mappings are known as stand-alone mappings.

## Dependent writes

This topic provides an overview about dependent writes.

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

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

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

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

## Operations on consistency groups

This topic provides an overview about the operations on consistency groups.

You can create, change, and delete consistency groups by using the command line tool that is described in *IBM TotalStorage Virtualization Family SAN Volume Controller Command-Line Interface User's Guide* or you can use the SAN Volume Controller Console that is described in *IBM TotalStorage Virtualization Family SAN Volume Controller Configuration Guide*.

## Limits

This topic provides information about the limits to working with consistency groups.

The SAN Volume Controller supports 128 consistency groups per cluster.

---

## FlashCopy indirection layer

This topic provides an overview about FlashCopy indirection layer.

FlashCopy provides the semantics of a point in time copy by using an indirection layer which intercepts I/Os targeted at both the source and target virtual disks. The act of starting a FlashCopy mapping causes this indirection layer to become active in the I/O path. This occurs as an atomic command across all FlashCopy mappings in the consistency group.

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

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

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

## Grains and the FlashCopy bitmap

This topic provides an overview about grains and the FlashCopy bitmap.

When data is copied from the source virtual disk to the target virtual disk, it is copied in units of address space known as grains. In the SAN Volume Controller, the grain size is 256KB. The FlashCopy bitmap contains one bit for each grain. The bit records whether the associated grain has yet been split by copying the grain from the source to the target.

## Source and target reads

This topic provides an overview about the source and target reads.

### Source reads:

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

### Target reads:

In order for FlashCopy to process a read from the target disk it must consult its bitmap. If the data being read has already been copied to the target then the read is sent to the target disk. If it has not, then the read is sent to the source disk.

Clearly this algorithm requires that, while this read is outstanding, no writes are allowed to execute which would change the data being read from the source. The SAN Volume Controller satisfies this requirement by using a cluster wide locking scheme.

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

## Writes to the source or target

This topic provides an overview about writing to the source or target.

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

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

## Limits

This topic provides an overview about the limits for FlashCopy indirection layer.

There is a per I/O group limit of 16TB on the quantity of source virtual disk address space which may participate in flash copy mappings.

This address space is allocated in units of 8 GB. That is to say, creating a FlashCopy mapping between a pair of virtual disks whose size is less than 8 GB will consume 8 GB of FlashCopy mapping address space.

This limit was calculated to allow for the maximum number of supported virtual disks per SAN Volume Controller cluster (8192) to each be a member of a FlashCopy mapping, provided that these virtual disks are spread evenly across the four I/O groups, with 2048 virtual disks per I/O group and each of the virtual disks is 8 GB or less in size.

The calculation is:

$$8 \text{ GB} * (2048 \text{ virtual disks per I/O group}) = 16 \text{ TB}$$

---

## FlashCopy mapping events

This topic provides an overview about FlashCopy mapping events.

FlashCopy mapping events details the events that modify the state of a FlashCopy mapping. The states are described in “FlashCopy mappings” on page 28.

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

### Prepare

The prepare command is directed to either a consistency group for FlashCopy mappings which are members of a normal consistency group or to the mapping name for FlashCopy mappings which are members of the special consistency group 0. The prepare command places the FlashCopy mapping in the preparing state.

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

### Flush done

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

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

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

During the **start** command:

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

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

#### **Modify**

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

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

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

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

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

#### **Flush failed**

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

#### **Copy complete**

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

#### **Bitmap Online/Offline**

The node has failed.

## **Background copy**

This topic provides an overview about background copy.

A FlashCopy mapping has a property background copy rate. This is expressed as a percentage and can take values between 0 and 100. The background copy rate can be changed when the FlashCopy mapping is in any state.

If a value of 0 is specified, then background copy is disabled. One use of this is for short lived FlashCopy mapping which are to be used for backup purposes only. Since the source data set is not expected to change much during the lifetime of the FlashCopy mapping, it is more efficient in terms of managed disk I/Os not to perform a background copy.

The relationship of the background copy rate value to the attempted number of grains to be split per second is given by the following table.

*Table 7. Background copy*

user percentage	KB/sec	grains/sec
1 - 10	128	0.5
11 - 20	256	1
21 - 30	512	2
41 - 50	2048	8
91 - 100	64 MB	256

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

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

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

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

---

## Host considerations for FlashCopy integrity

This task provides step-by-step instructions to flush data from your host volumes and perform the FlashCopy.

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

before proceeding with the FlashCopy. Host cache flushing is ensured by unmounting the source virtual disks from the source host before starting the FlashCopy.

**Steps:**

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

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

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

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

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

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

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

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

**Related topics:**

- “FlashCopy” on page 27
- “FlashCopy mappings” on page 28

---

## Remote Copy

Remote Copy is a Copy Service available with the SAN Volume Controller. Remote Copy enables you to set up source-target relationships between virtual disks (VDisks). This enables the SAN Volume Controller to copy host data on a particular source VDisk to the target VDisk designated in the relationship. Remote Copy is particularly useful when you have a requirement, such as disaster recovery, in which two copies of a VDisk are separated by some distance. The SAN Volume Controller assumes that the SAN fabric to which it is attached contains hardware that achieves the long-distance requirement.

One VDisk is designated the source (primary) and the other VDisk is the target (secondary). Host applications write data to the primary VDisk, and updates to the primary VDisk are copied to the target VDisk. Normally, host applications do not perform input/output (I/O) operations to the target VDisk.

Remote Copy supports the following features:

- Intracluster copying of a VDisk, in which both VDisks belong to the same cluster
- Intercluster copying of VDisk, in which one VDisk belongs to a cluster and the other VDisk belongs to a different cluster

**Note:** A cluster can communicate with only one other cluster.

- Intercluster and intracluster Remote Copy can be used concurrently within a cluster for different relationships

### Related topics:

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

## Synchronous Remote Copy

In the synchronous mode, Remote Copy provides a *consistent* copy, which means that the target VDisk is always the exact match of the source VDisk. The host application writes data to the source VDisk but does not receive the final status on the write operation until the data is written to the target VDisk. For disaster recovery, this mode is the only practical mode of operation because a consistent copy of the data is maintained. However, synchronous mode is slower than asynchronous mode because of the latency time and bandwidth limitations imposed by the communication link to the secondary site.

### Related topics:

- “Remote Copy”

## Remote Copy partnerships

With Remote Copy, you can copy a VDisk in one cluster to a VDisk in another cluster. The SAN Volume Controller needs to know not only about the relationship between the two VDisks but also about the relationship between the two clusters. A Remote Copy partnership defines the relationship between the two clusters. In a cluster partnership, one cluster is defined as the local cluster and the other cluster as the remote cluster.

### **Background copy management:**

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

### **Related topics:**

- “Remote Copy” on page 37
- “Remote Copy relationships”

## **Remote Copy relationships**

Because Remote Copy copies one VDisk to another VDisk, the SAN Volume Controller needs to know about that relationship. A Remote Copy relationship defines the relationship between two virtual disks: a master VDisk and an auxiliary VDisk. In most cases, the master VDisk contains the production copy of the data and is the VDisk that the application accesses. The auxiliary VDisk contains a backup copy of the data and is used in disaster recovery scenarios.

Additionally, the master VDisk and the auxiliary VDisk can take on specific roles. They can take on the role of a primary VDisk or a secondary VDisk.

### **Primary**

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

### **Secondary**

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

When a relationship is created, the master VDisk is assigned the role of primary VDisk and the auxiliary VDisk is assigned the role of secondary VDisk. Hence, the copying direction is from master to auxiliary.

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 input/output (I/O) group.

### **Related topics:**

- “Remote Copy” on page 37
- “Remote Copy partnerships” on page 37
- “Remote Copy relationships”

## **Remote Copy consistency groups**

Certain uses of Remote Copy require the manipulation of more than one relationship. Remote Copy provides the ability to group relationships so that they are manipulated in unison. To address this requirement, consistency groups have been created.

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

This ensures that the set of secondary VDisks contains usable data. The key property is that these relationships be consistent. Hence, the groups are called consistency groups.

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

#### **Remote Copy states:**

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

##### **Connected**

The two clusters can communicate.

##### **Disconnected**

The two clusters cannot communicate. This state only applies to VDisks that are in different clusters.

Relationships that contain VDisks that operate as secondary VDisks can be described as being consistent, inconsistent, or synchronized.

##### **Consistent**

Contains the same data as the primary VDisk at a point at which an imaginary power failure occurred.

**Note:** Consistent does not mean that the data is up-to-date. A copy could contain data that is frozen at some point in time in the past.

##### **Inconsistent**

Does not contain the same data as the primary VDisk at a point at which an imaginary power failure occurred.

##### **Synchronized**

A secondary VDisk that is consistent and *up-to-date*.

#### **Related topics:**

- “Remote Copy relationships” on page 38
- “Remote Copy” on page 37



---

## Chapter 4. Configuration rules and requirements

This chapter describes the rules and requirements for configuring a SAN Volume Controller. It also provides a list of defined terms that are referenced in the configuration rules.

### Definitions:

Before you read the rules, read these definitions, which can help you to understand the rules.

### Properties:

#### **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, an ISL hop 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, where more than one ISL is in parallel between these switches.

This definition assumes a symmetrical network and a specific workload that is applied equally from all initiators and sent equally to all targets. A symmetrical network means that all initiators are connected at the same level and all the controllers are connected at the same level.

The SAN Volume Controller makes this calculation difficult, because it puts its back-end traffic onto the same network, and this back-end traffic varies by workload. Therefore, the oversubscription that a 100% read hit gives is different from the oversubscription that 100% write-miss gives.

If you have an oversubscription of 1 or less, the network is nonblocking.

#### **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. The way to make a redundant SAN is to split 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 Remote Copy, 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 remote copy, significant distances might exist between the components of the local cluster and those of the remote cluster.

**Local/Remote fabric interconnect**

The SAN components that connect the local fabrics to the remote fabrics. These components might be single-mode optical fibres that are driven by GigaBit Interface Converters (GBICs), or they might be other, more advanced components, such as channel extenders.

**SAN Volume Controller Fibre Channel port fan in**

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

Some controllers recommend that the number of hosts that use each port be limited to prevent excessive queuing at that port. If the port fails, or the path to that port fails, the host might failover to another port, and the fan-in requirements might be exceeded in this degraded mode.

**Invalid configuration**

A configuration that refuses to operate and generates an error code to indicate what caused it to become invalid.

**Unsupported configuration**

A configuration that might operate successfully, but for which IBM does not guarantee to be able to solve problems that might occur.

Usually this type of configuration does not create an error log.

**Valid configuration**

A configuration that is neither invalid nor unsupported.

**Degraded**

A valid configuration that has had a failure, but continues to be neither invalid nor unsupported.

Typically, a repair action is required to restore the degraded configuration to a valid configuration.

---

## Configuration rules

SAN configurations that contain SAN Volume Controller clusters can be set up in various ways. Some configurations, however, do not work, and are known as *invalid*. You can avoid creating invalid configurations if you follow the rules that are given in this section.

A SAN configuration that contains SAN Volume Controllers is valid if it observes *all* of the following rules:

### RAID controllers

All SAN Volume Controller nodes of a SAN Volume Controller cluster must be able to see the same set of back-end storage ports on each back-end controller. Any operation that is in this mode in which two nodes do not see the same set of ports on the same controller is degraded, and the system logs errors that request a repair action. This rule can have important effects on back-end storage such as FASST,

which has exclusion rules that determine to which host bus adapter (HBA) WWNNs a storage partition can be mapped.

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

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

The SAN Volume Controller must be configured to manage only LUNs that are exported by supported RAID controllers. Operation with other RAID controllers is not supported.

The SAN Volume Controller must not share its back-end storage with hosts or controllers. Any operation that is in this mode in which a back-end controller is shared with hosts or controllers is not supported.

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

## Host bus adapters (HBAs)

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

HBAs that are in dissimilar hosts, or dissimilar HBAs that are in the same host must be in separate zones. For example, if you have an AIX host and a Windows NT host, those hosts must be in separate zones. Here, *dissimilar* means that the hosts are running different operating systems, or that they are different hardware platforms. Different levels of the same operating system are, therefore, thought of as similar. This rule helps you to ensure that different SANs can operate with each other. A configuration that breaks this rule is not supported.

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

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

Operation with other HBAs is not supported.

The number of paths from the SAN Volume Controller nodes to a host must not exceed eight. The maximum number of host HBA ports must not exceed four (for example, no more than two two-port HBAs or four one-port HBAs). Because each SAN Volume Controller node in an I/O group presents four images of a virtual disk (VDisk) onto the SAN, and each host SAN attachment has up to four HBA ports, this means that with more simplified zoning the number of paths may be up to 4 SAN Volume Controller ports x 2 nodes per I/O group x 4 host HBA ports, equalling 32. To restrict the number of paths to a host the switches should be

zoned such that each host HBA port is zoned with one SAN Volume Controller port of each node in the cluster. If a host has multiple HBA ports then each port should be zoned to a different set of SAN Volume Controller ports to maximize performance and redundancy.

## Nodes

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

The two nodes of an I/O group must be no more than 100 m (328 ft) apart.

## Power

The uninterruptible power supply must be in the same rack that contains the SAN Volume Controller nodes. The combination power and signal cable for connection between SAN Volume Controller and uninterruptible power supply units is 2m long. The SAN Volume Controller and uninterruptible power supply must communicate through RS-232.

## Fibre channel switches

The SAN must contain only supported switches. The SAN Volume Controller supports specific IBM 2109, McData, and InRange switch models. See the following Web site for the latest models and firmware levels:

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

Operation with other switches is not supported.

Different vendor switches may not be intermixed in the same Counterpart SAN. A redundant SAN, made up of more than one Counterpart SAN may contain different vendor switches, provided the same vendor is used within each Counterpart SAN.

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

Configurations with more than two SANs are not supported.

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

On the Fibre Channel SAN, back-end storage must always and only be connected to SAN switches. Multiple connections are permitted from the redundant controllers of the back-end storage, to improve data bandwidth performance. It is not necessary to have a connection between each redundant controller of the back-end storage and each counterpart SAN. For example, in a FASST configuration in which the FASST contains two redundant controllers, only two controller minihubs are usually used. Controller A of the FASST is therefore

connected to counterpart SAN A, and controller B of the FASST is connected to counterpart SAN B. Any operation that uses direct connection between host and controller is not supported.

The connections between the switches and the SAN Volume Controllers can operate at 1 GB per second or at 2 GB per second. All the ports of SAN Volume Controllers that are in a single cluster, however, must run at one speed. Any operation that runs different speeds on the node-to-switch connections that are in a single cluster is invalid.

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

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

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

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

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

Operation with fibre channel extenders is not supported.

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

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

Because each SAN Volume Controller has four ports, the switches can be zoned so that a particular SAN Volume Controller port is used only for internode communication, for communication to the host, or for communication to back-end storage. Whatever the configuration, each SAN Volume Controller node must remain connected to the full SAN fabric. Zoning must not be used to split the SAN into two parts. In remote copy, additional zones are required that contain only the local nodes and the remote nodes. It is valid for the local hosts to see the remote nodes, or for the remote hosts to see the local nodes. Any zone that contains the local and the remote back-end storage and local nodes or remote nodes, or both, is not valid.

## Configuration requirements

This chapter describes the steps you *must* perform before you configure the SAN Volume Controller.

### Steps:

Perform the following steps:

1. Your IBM customer engineer (CE) must have installed the SAN Volume Controller. See the *IBM TotalStorage SAN Volume Controller: Installation and Hardware Reference*.
2. Install and configure your RAID controllers and create the RAID resources that you intend to virtualize. To prevent loss of data, virtualize only those RAID arrays 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.

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

When in doubt, arrays with between 6 and 8 member disks is recommended.

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

Most back-end storage controllers allow RAID arrays to be divided up into more than one SCSI LU. When configuring new storage for use with the SAN Volume Controller it is not necessary to divide up the array and so it should be presented as one SCSI LU. This will give a one-to-one relationship between MDisks and RAID arrays.

**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 needs. One zone must contain all the RAID controllers and the SAN Volume Controller nodes. For hosts, use switch zoning to ensure that each host FC port is zoned to exactly one FC port of each SAN Volume Controller node in the cluster. The SAN Volume Controller and the master console exists in both zones.

**Note:** The SAN Volume Controller and the master console are defined in each zone.

For more information on zoning your switches, refer to *IBM TotalStorage SAN Volume Controller: Configuration Guide*.

4. In order to have redundant paths to disks exported by the SAN Volume Controller, you must install SDD on all of your hosts that are connected to the

SAN Volume Controller. Otherwise, you will not be able to use the redundancy inherent in the configuration. Install the Subsystem Device Driver (SDD) from the following Web site:

<http://ssddom02.storage.ibm.com/techsup/webnav.nsf/support/sdd>

Be sure to install version 1.4.x.x or higher.

5. Install and configure the SAN Volume Controller master console. The communication between the master console and the SAN Volume Controller runs under a client-server network application called Secure Shell (SSH). Each SAN Volume Controller cluster is equipped with SSH Server software and the master console comes to you equipped with the SSH Client software called PuTTY. You will need to configure the SSH client key pair using PuTTY on the master console. Once you have installed your master console, you can configure and administer the SAN Volume Controller using a graphical interface or a Command-Line Interface.
  - You can configure the SAN Volume Controller using the master console which is a web application providing browser access to the SAN Volume Controller configuration tools.

**Note:** You can also install the master console on another machine which you provide from the master console CD-ROM.

- You can configure the SAN Volume Controller using the Command-Line Interface (CLI) commands.
- You can install an SSH client if you only want to use the CLI commands. If you want to use the CLI from a host other than the master console, you need to download and install an SSH client from the following Web site:

<http://commerce.ssh.com>

#### **Result:**

When you and the IBM customer engineer have completed the initial preparation steps, you must:

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

---

## **Maximum configurations**

This topic provides information about maximum configurations.

The following table displays the maximum configurations supported by a SAN Volume Controller cluster.

**Note:** Not all maximum configurations can be supported simultaneously.

*Table 8. Maximum configurations for the SAN Volume Controller*

Property	Maximum	Notes
nodes	4	Arranged as two pairs.
SAN ports	256	Maximum size of fabric, including all SAN Volume Controller nodes.
Host ports	128	Hosts may be connected through multiple ports. <b>Note:</b> A port is used as a host port or a controller port is not monitored by a SAN Volume Controller.
RAID disk controllers	64	
Controller ports	256	16 ports per controller.
Managed LUNs	4096	Represents an average of 64 per controller.
Virtual Disks (VDisks)	1024	
VDisks per Host ID	512	
Addressability	2.1PB	Maximum extent size is 512MB, arbitrary limit $2^{22}$ extents in a mapping.
LU size	2TB	Defined by 32 bit LBA limit.
SDD	512 SAN Volume Controller vpaths per host	<p>One vpath is created for each VDisk mapped to a host. Although the SAN Volume Controller only permits 512 VDisks to be mapped to a host, the SDD limit can be exceeded by either:</p> <ul style="list-style-type: none"><li>• Creating two (or more) host objects for one physical host and mapping more than 512 VDisks to the host using the multiple host objects.</li><li>• Creating two (or more) clusters and mapping more than 512 VDisks to the host using the multiple clusters.</li></ul> <p><b>Note:</b> Both of these operations are unsupported.</p>

---

## **Part 2. Preparing to configure the SAN Volume Controller**

This part provides a description of tasks that you will need to perform before you begin the configuration of the SAN Volume Controller. Fundamentally, the configuration of the SAN Volume Controller begins by completing a two phase creation (initialization) of the cluster. The first phase is performed from the front panel of the cluster. The completion of the creation of the cluster is performed from the SAN Volume Controller Console which is accessible from a Web server running on the master console. This part provides the following information:

- Chapter 5, “Create cluster from the front panel”, on page 51
- Chapter 7, “Master console overview”, on page 57



---

## Chapter 5. Create cluster from the front panel

The task provides step-by-step instructions you will need to perform to create the cluster from the front panel.

### Prerequisites:

Ensure that the SAN Volume Controller nodes are installed. Prior to configuring the cluster, should you choose to have the customer engineer (CE) initially create the cluster, ensure that you have supplied the following information to your customer engineer:

1. Ensure that you have the correct license. The license will show you whether you are permitted to use flash copy or remote copy. It will also show how much virtualization you are permitted to use.
2. You must supply the following information to enable the customer engineer to start the configuration procedure:
  - Cluster IP address. This address must be unique, otherwise communication problems can occur.
  - Subnet mask
  - Gateway IP address

The customer engineer uses the front panel of the SAN Volume Controller to enter the information that you have supplied. The SAN Volume Controller generates a random password on the display panel that the customer engineer will give to you.

3. Make a note of the password and the IP address. You need it when you connect to the web application to create the cluster.

### Context:

Ensure that you have a new pair of nodes and you want to make a cluster. You also want to gain access to this cluster to start your configuration. The steps are as follows:

1. Choose a node and create a new cluster.
2. Set the IP addresses so you can gain access to the cluster.
3. Configure your cluster.

### Steps:

Perform the following steps to create the cluster:

1. Choose any node that is to become a member of the cluster that you are creating.
2. At the IBM TotalStorage SAN Volume Controller service panel, keep pressing and releasing the up or down navigation button until **Node:** is displayed.
3. Keep pressing and releasing the left or right navigation button until **Create Cluster?** is displayed.
4. Press the **Select** button.

If **IP Address:** is displayed on line 1 of the screen, go to step 5 on page 52.

If Delete Cluster? is displayed in line 1 of the service display screen, this node is already a member of a cluster. Either you have selected the wrong node, or you have already used this node in an existing cluster. The ID of this existing cluster is displayed in line 2 of the service display screen.

- If you selected the wrong node you can exit this procedure now. The procedure cancels automatically after 60 seconds.

**Attention:** When a node is deleted from a cluster, all customer data that is contained in that node is lost.

If you are sure that the existing cluster is not required:

- Press and hold the up button.
- Press and release the select button. The node will be restarted. Once the node has been restarted you must then restart this procedure from step 1 on page 51. IP Address: is displayed on the service display screen.
- Go to step 5.

**Changing the fibre channel port speed** To display the menu that shows the current value of the fibre channel speed setting for the node, press and hold the down button. Then press the select button when the display is showing the status of one of the fibre channel (FC) ports. The setting should be either 1 Gb or 2 Gb. To change the setting, perform the following steps:

- Press the up or down buttons to select the speed.
- Press the select button when the speed you want is displayed.

This action changes the speed of all the fibre channel ports on the node.

- Press the select button.
  - Use the up or down navigation button to change the value of the first field of the IP Address to the value that you have chosen.
  - Use the right navigation button to move to the next field. Use the up or down navigation buttons to change the value of this field.
  - Repeat step 7 for each of the remaining fields of the IP Address.
  - When you have changed the last field of the IP Address, press the select button.
  - Press the right button. Subnet Mask: is displayed.
  - Press the select button.
  - Change the fields for Subnet Mask in the same way that you changed the fields for IP Address.
  - When you have changed the last field of Subnet Mask, press the select button.
  - Press the right navigation button. Gateway: is displayed.
  - Press the select button.
  - Change the fields for Gateway in the same way that you changed the fields for IP Address.
  - When you have changed the last field of Gateway, press the select button.
  - Keep pressing and releasing the right-hand navigation button until Create Now? is displayed.
  - If you are satisfied with your settings, press the select navigation button. 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 the cluster is created successfully, Password: is displayed in line 1 of the service display screen. Line 2 contains a password that you can use to access

the cluster. Make a note of this password now. The password is displayed for only 60 seconds, or until the up, down, left or right navigation button is pressed.

**Attention:** If you do not record the password, you will have to start the cluster configuration procedure again. When the password has been recorded, press the up, down, left, or right navigation button to delete the password from the screen.

20. If the cluster is created successfully:
  - Cluster: is displayed in line 1 of the service display screen,
  - the cluster IP address is displayed on line 2,
  - and you have successfully completed the creating a cluster process.

If the cluster is not created, Create Failed: is displayed in line 1 of the service display screen. Line 2 contains an error code. Refer to the error codes that are given in the *IBM TotalStorage SAN Suite: SAN Volume Controller Service Guide* to find the reason why the cluster was not created.



---

## Chapter 6. Master console security overview

This topic provides information about security.

There are several passwords and IDs that have been set to default values in manufacturing that need to be changed on the SAN Volume Controller system.

**Note:** It is important to change the default passwords to maintain the security of the master console.

---

### Overview of passwords

This topic provides an overview about passwords.

The following passwords must be set:

**Note:** If you forget your Superuser password, you must contact IBM Service.

1. **Windows user ID and password:** Use the Computer Management Administrative tool to change the user ID and passwords. To access this tool, select Start -> Settings -> Control Panel and double-click Administrative Tools -> Computer Management -> Local Users and Groups from the left-hand navigation.

**Notes:**

- a. Any new user ID generated must have administrator privileges, if it is to function with all of the master console software.
  - b. If the Windows password is changed, you also need to make changes to the Tivoli SAN Manager's Host Authorization because it is used to authorize access.
2. **Tivoli SAN Manager (changing the Host Authorization):** You can change this password by performing the following steps:
    - a. Open a command prompt window by selecting Start -> Programs -> Accessories -> Command Prompt.
    - b. Enter CD c:\tivoli\itsanm\manager\bin\w32-ix86.
    - c. Enter

```
srmcp -u -userID -p password ConfigService  
setAuthenticationPw newPassword
```

where *-userID* is your user ID, *password* is your password, and *newPassword* is the new host authentication password.

These internally used IDs and passwords can also be changed if required.

**DB2 user IDs and passwords:**

- **Base User ID = db2admin:** use the Computer Management Administrative tool to change this password. To access this tool, select Start -> Settings -> Control Panel and double-click Administrative Tools -> Computer Management -> Local Users and Groups from the left-hand navigation.

- **Database used by Tivoli SAN Manager = db2user:** use the Computer Management Administrative tool to change this password. If it is changed, then you must perform the following steps:

1. Open a command line window by selecting **Start -> Programs -> Accessories -> Command Prompt**.
2. Enter CD c:\tivoli\itsanm\manager\bin\w32-ix86
3. Enter

```
srmcp -u -userID -p password ConfigService  
setAuthenticationPw newPassword
```

where *-userID* is your user ID, *password* is your password, and *newPassword* is the new host authentication password.

- **SAN Volume Controller user ID and password:** these are set using the SAN Volume Controller Web pages accessed using a Web browser or using the SAN Volume Controller Console function. See Chapter 8, “Getting started with the SAN Volume Controller Console”, on page 77 or Chapter 13, “Getting started with the Command-Line Interface”, on page 141 for more information.

---

## Chapter 7. Master console overview

The SAN Volume Controller provides a master console that is used as a single platform to configure, manage, and service software required to manage the SAN Volume Controller.

The master console allows a system administrator to rapidly integrate the virtualization controller into their environment. The master console monitors the configuration of the whole system and all of the internal components. It offers a standard and central location for all aspects of the operation, including SAN topology rendering, SNMP trap management, Call Home and Remote Service facilities as well as all the configuration and diagnostic utilities for the components.

**Note:** VPN connection is required for Remote Service facilities.

The master console also provides the following functions:

- Browser support for:
  - SAN Volume Controller Console
  - Fibre channel switch
- CLI configuration support using Secure Shell (SSH)
- SAN Topology rendering using Tivoli® SAN Manager
- Remote Service capability through VPN
- IBM Director
  - SNMP Trap management
  - Call Home capability using IBM Director
  - E-mail notification to Customer, for example, System Administrator

---

## Configuring the master console

This topic provides an overview of the steps that you will need to complete to configure the master console.

### Steps:

Perform the following steps to successfully configure the master console:

1. Log onto the master console.
2. Generate an SSH key pair using the SSH client called PuTTY.
3. Configure the PuTTY session (for Command-Line Interface (CLI) access only).
4. Start the SAN Volume Controller Console for the SAN Volume Controller.
5. Store the master console SSH public key file on each SAN Volume Controller cluster.
6. Set up a new zone on the Fibre Channel Switches that includes the master console and all of the 2145 ports
7. Configure the Tivoli SAN Manager (TSanM).
8. Start the TSanM.
9. Configure the Remote Support.
10. Start the IBM Director.

11. Configure the IBM Director.
12. Set up Call-Home and E-mail.
13. Recover disk drives on the master console.
14. Upgrade software on the master console.

**Related topics:**

- Chapter 6, “Master console security overview”, on page 55
- “Generating an SSH key pair using the SSH client called PuTTY” on page 60
- “Configuring the PuTTY session for the command-line interface” on page 64
- “Accessing the SAN Volume Controller Console” on page 77
- “Adding subsequent SSH public keys to the SAN Volume Controller” on page 129
- “Configuring the Tivoli SAN Manager (TSanM)” on page 64
- “Starting the TSanM” on page 65
- “Configuring IBM Director for the SAN Volume Controller error notification and Call-Home” on page 69
- “Upgrading software on the master console” on page 72

---

## Secure Shell (SSH) configuration

This topic describes how to use the secure shell client remotely from your host system.

**Overview:**

SSH is a client-server network application. The SAN Volume Controller cluster acts as the SSH server in this relationship. The secure shell (SSH) client provides a secure environment in which to connect to a remote machine. It uses the principles of public and private keys for authentication.

SSH keys are generated by the SSH software. This includes a public key, which is uploaded and maintained by the cluster and a private key that is kept private to the server that is running the SSH client. These keys authorize specific users to access the administration and service functions on the cluster. Each key is associated with a user-defined ID string that can consist of up to 40 characters. Up to 100 keys can be stored on the cluster. You can also add new IDs and keys or delete unwanted IDs and keys.

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

- the SAN Volume Controller Command-Line Interface (CLI)
- or the system on which the SAN Volume Controller Console is installed.

When an SSH client (A) attempts to connect to an SSH server (B), the key pair is needed to authenticate the connection. The key consists of two halves; the public and private keys. The SSH client public key is put onto the SSH Server (B) by a means outside of the SSH session. When the SSH client (A) tries to connect, the private key on the SSH client (A) is able to authenticate with its public half on the SSH server (B).

In order to use the Command-Line Interface (CLI) or SAN Volume Controller Console system you must have an SSH Client installed on that system, generate the SSH key pair on the client system and store the client's SSH public key on the SAN Volume Controller cluster(s).

The master console has the SSH client software called PuTTY preinstalled. This software provides the Secure Shell (SSH) client function for users logged into the master console who wish to invoke the SAN Volume Controller Command-Line Interface (CLI).

If you wish to run the SAN Volume Controller Command-Line Interface (CLI) from a different system than the master console, you must install an SSH client. For your convenience, the installation program to install the PuTTY software on Windows can be found in the SSH client directory of the SAN Volume Controller Console CD-ROM. You can generate SSH public and private keys using the PuTTY software. You must store the SSH Client public key on all SAN Volume Controller clusters to which the CLI will connect.

The master console also has the SAN Volume Controller Console Web server and Common Information Model (CIM) Object Manager software preinstalled. This software depends on the PuTTY Secure Shell (SSH) client function for the SAN Volume Controller Console to programmatically access the SAN Volume Controller cluster. The master console comes with PuTTY SSH keys preinstalled. You can generate new PuTTY SSH keys unique to your Master Console and copy the private SSH key to the SAN Volume Controller Console directory and store the public SSH key on all clusters to which the SAN Volume Controller will connect.

You can also install the SAN Volume Controller Console on a Windows 2000 server system which you provide. If you intend to install the SAN Volume Controller on a host which you supply, you must install PuTTY first, which is a prerequisite for the SAN Volume Controller Console.

**Related topics:**

- “Generating an SSH key pair using the SSH client called PuTTY” on page 60

---

## Overview of configuring the Secure Shell (SSH)

This topic provides an overview about configuring the SSH client system. The following sections elaborate on each step to configure a PuTTY Secure Shell client system. IBM has preinstalled the PuTTY Secure Shell client software on the master console. You can also install Putty on any Windows 2000 server where you will run the Command-Line Interface (CLI) or where you can choose to install the SAN Volume Controller Console. If you have some other Secure Shell client software to run on another host, follow that software's documentation to perform the tasks equivalent to the following steps.

1. Install SSH client software (not required for master console which has PuTTY preinstalled).
2. Generate SSH keys on the SSH client system.
3. Configure the session, if required, on the SSH client system.
4. If client system is the master console, copy the private key into the SAN Volume Controller Console install directory.
5. Store the SSH client public key on the SAN Volume Controller cluster.

You will perform step 5 to store the SSH client public key on the SAN Volume Controller Cluster when you complete the creation of the SAN Volume Controller

cluster. Once you have defined a cluster to the SAN Volume Controller Console and have therefore enabled SSH communication to the cluster , you can store additional SSH client public keys on the cluster. You can store additional keys through the SAN Volume Controller Console or the Command-Line Interface.

**Related topics:**

- “Generating an SSH key pair using the SSH client called PuTTY”
- “Configuring the PuTTY session for the command-line interface” on page 64
- “Adding subsequent SSH public keys to the SAN Volume Controller” on page 129

---

## Generating an SSH key pair using the SSH client called PuTTY

This task provides step-by-step instructions for generating SSH keys on the PuTTY SSH client system.

**Steps:**

Perform these steps to generate SSH keys on the SSH client system:

1. Start the PuTTY Key Generator to generate public and private keys for SSH client connection to the SSH Server on the San Volume Controller cluster. Select **Start -> Programs -> PuTTY -> PuTTYgen** to open the PuTTY Key Generator Graphical User Interface (GUI) window.
2. Use the PuTTY Key Generator GUI window to generate keys:
  - a. Select the **SSH2 RSA** radio button.
  - b. Leave the number of bits in a generated key value at 1024.
  - c. Click **Generate**.

A message similar to the following is displayed:

Please generate some randomness by moving  
the mouse over the blank area.

in the section of the GUI labeled Key. The *blank area* indicated by the message is the large blank rectangle on the GUI inside the section of the GUI labeled Key. Continue to move the cursor over the blank area until the progress bar reaches the far right. This generates random characters to create a unique key.

- Attention:** Do not enter anything in **Key Passphrase** or **Confirm passphrase** fields.
3. Save the generated SSH keys on your system disk for later use
    - a. Click **Save public key**. You will be prompted for a name and location for the key. Remember the name and location of the SSH public key you save.

**Note:** It is recommended that you use the term pub in naming the public key, for example, pubkey, to easily differentiate the SSH public key from the SSH private key. You will identify the name and location of the SSH public key to the SAN Volume Controller cluster in a later step.

- b. Click **Save Private key**. You will be prompted with a message similar to the following:

Are you sure you want to save this key  
without a passphrase to protect it?  
Yes/No

Click **Yes**. You will be prompted for a name and location for the key. Remember the name and location of the SSH private key you save. You will need to identify the name and location of the SSH private key when you configure the PuTTY session in the following step. You will also need the name and location of the SSH private key if you choose to run the SAN Volume Controller Console installation program on another system other than the master console. The PuTTY key generator will save the private key with an extension of .ppk.

4. Close the PuTTY Key Generator.

**Related topics:**

- “Configuring the PuTTY session for the command-line interface” on page 64

---

## Storing keys in the SAN Volume Controller Console software

This topic provides step-by-step instructions for storing your SSH keys in the SAN Volume Controller Console software.

When the keys that are used to communicate with the SAN Volume Controller are changed, you must store a copy of the new private key in the SAN Volume Controller Console software.

**Steps:**

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 by clicking **Start -> Run**.
2. Type cmd.exe in the Open box. Click **OK**.
3. Type the following command:

```
copy <path><filename> C:\\"Program Files"\IBM \svcconsole\cimom\icat.ppk
```

where *<path><filename>* is the path and file name where you stored the SSH private key when it was generated in the previous procedure.

**Note:** Directory names with embedded spaces must be surrounded by quotation marks

4. Stop and start the IBM CIM Object Manager to make the change take effect. Perform the following:
  - 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**. Wait for Windows to stop the service.
  - e. Select **IBM CIM Object Manager** in the list of services, right click and select **Start**.

---

## Adding SSH keys for hosts other than the master console

This task provides step-by-step instructions for adding additional SSH keys on hosts other than the master console.

### Steps:

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 wish to be able to access 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. Secure copy these public keys from the master console to the cluster.  
Repeat for each hosts public key copied onto the master console in 2.

---

## Replacing the SSH key pair

This topic provides step-by-step instructions for replacing the SSH key pair.

- If you change the SSH keys that will be used by the master console to communicate with the SAN Volume Controller Console you will have to store the client SSH private key in the SAN Volume Controller Console software as described in the section above 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 will not be aware of the existence of the cluster.

The procedure to correct this, is to remove the cluster from the SAN Volume Controller Console and add it back again. To correct these scenarios, perform the following steps:

1. Start the SAN Volume Controller Console by clicking on the desktop icon or by using your Web browser to go to  
  
`http://<IPAddress>:9080/ica`  
  
where <IPAddress> is the IP address of the master console. The Sign on 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 **Remove a cluster** in the selection box.
6. Click **Go**.
7. Click **Clusters** from the portfolio.
8. Click **Add a cluster** in the selection box.
9. Input the IP address of the cluster.
10. Do not check the **Create (Initialize Cluster** box.
11. Click **OK**.

12. Enter the user name and password. When you see the pop-up window, enter the network password and click **OK**.
13. Add the SSH client public key to the SAN Volume Controller cluster:
  - a. Click **Browse...** for the key file to upload and locate the public key or input the key in the **Key (direct input)** field.
  - b. Type an ID in the **ID** field, which uniquely identifies the key to the cluster.
  - c. Select the **administrator** radio button.
  - d. Click **Add Key**.
  - e. Click **Clusters** from the portfolio to check the status of the cluster. If the cluster status remains **SSH Key Refused**, you do not have a good key pair. You can reset the SAN Volume Controller Console private SSH key. However, if you have successfully contacted other clusters, you will break that connectivity.

## Replace the client SSH private key known to the SAN Volume Controller software

This task provides step-by-step instructions to replace the client SSH private key 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 known to the SAN Volume Controller software.

### Steps:

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, stop the IBM CIM Object Manager. Perform the following:
  - 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. Copy the client SSH private key into the appropriate SAN Volume Controller Console directory. Perform the following:
  - a. Open a command prompt window by clicking **Start -> Run**.
  - b. Type **cmd.exe** in the **Open** field.
  - c. Click **OK**.
4. Type the following command:

```
copy <filename> C:\program files\IBM\svcconsole\cimom\icat.ppk
```

where *<filename>* is the path and file name of the client SSH private key.

5. Restart the IBM CIM Object Manager. Select **IBM CIM Object Manager** in the list of services, right click and select **Start**.
6. Log on to the SAN Volume Controller Console.
7. Click **Clusters** in the portfolio.
8. Check the status of the cluster.

---

## Configuring the PuTTY session for the command-line interface

This task provides step-by-step instructions for configuring the PuTTY session for the command-line interface on the SSH client system. This step is only required if you are preparing to run the Command Line Interface (CLI) from the master console.

### Steps:

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

1. Configure the PuTTY client session. Select **Start -> Programs -> PuTTY -> PuTTY** to open the PuTTY Configuration GUI window. The items you select in the large vertical left hand section of the GUI labeled Category will affect the content of the right hand section of the GUI. Complete the following steps:
  - a. Click **Session** in the large vertical left hand section of the GUI labeled Category.
  - b. Select the **SSH** radio button.
  - c. Click **Connection -> SSH** in the Connection tree. This will bring up a different view in the right hand side of the GUI.
  - d. Ensure the radio button labeled **2** is selected.
  - e. Click **Auth** in the SSH tree in the Connection tree in the large vertical left hand section of the GUI labeled Category. This will bring up a different view in the right hand side of the GUI.
  - f. Type the name of the SSH client private key file you specified when you used the PuTTY Key Generator in the **Private key file for authentication** field in the Authentication Parameters. This field is the second section of the right hand side of the GUI. You can click **Browse** to select the file name from the system directory or, alternatively, type the fully qualified file name (for example, C:\Support Utils\PutTY\priv.ppk).
  - g. Click **Session** in the large vertical left hand section of the GUI labeled Category.
  - h. Click **Default Settings -> Save** in the middle right hand section of the GUI labeled Load, save or delete a stored session field.

### Related topics:

- “Configuring the master console” on page 57

---

## Configuring the Tivoli SAN Manager (TSanM)

This task provides information about configuring the TSanM.

### Prerequisites:

When the master console was installed, its name might have been changed from the default name given in manufacture to a name of your choice, if so then it will be necessary to modify some Tivoli SAN Manager configuration files to reflect the new name. Using Windows Notepad, open the following files and make the changes indicated:

1. Open the file called, c:\tivoli\itsanm\manager\bin\w32-ix86\setenv.bat. Find the line in the file *TSNM\_LOCAL\_HOSTNAME=xxxxxx*, replace the *xxxxxx* with the master console’s full DNS name, then save and close the file.

2. Open the file called, c:\tivoli\itsanm\manager\bin\w32-ix86\setDeployEnv.bat. Find the line in the file *NODE=xxxxxx*, replace the *xxxxx* with the master console's Short Name, then save and close the file.
3. Open the file called, c:\tivoli\itsanm\manager\conf\tsnmdbparms.properties. Find the line in the file *tivoli.sanmgmt.jdbc.dbURL=xxxxxxxx*, replace the *xxxxxx* with the master console's full DNS name, then save and close the file.
4. Open the file called, c:\tivoli\itsanm\manager\conf\user.properties. Find the line in the file *SANDomainID=xxxxxx*, replace the *xxxxx* with the master console's full DNS name, then save and close the file.
5. Open the file called, c:\tivoli\itsanm\manager\apps\was\java\jre\lib\orb.properties. Find the line in the file *com.ibm.CORBA.LocalHost=xxxxxxxx*, replace the *xxxxxx* with the master console's Short Name, then save and close the file.
6. Open the file called,  

```
c:\tivoli\itsanm\manager\apps\was\config\cells\DefaultNode\nodes\
DefaultNode\serverindex.xml
```

Find the line in the file *hostName=xxxxxxxx*, replace the *xxxxx* with the master console's Short Name. Find the lines, (there are 8 occurrences), *host=xxxxxxxx*, replace the *xxxxx* with the master console's Short Name, then save and close the file.
7. Open the file called, c:\WINNT\system32\drivers\etc\HOSTS. Find the last line in the file and replace the IP Address with the new address of master console's. Replace the Short Name with the new name of master console's. Replace the full DNS name with the new full DNS name of master console's, then save and close the file.

**Note:** To assist Tivoli SAN Manager present its optimum information, you can install the Tivoli SAN Manager Agent software on each of your host systems.

#### Related topics:

- “Configuring the master console” on page 57

## Starting the TSanM

This task provides step-by-step instructions about how to start the TSanM. Ensure that you configure the Tivoli SAN Manager (TSanM) to meet your requirements.

#### Steps:

Perform the following steps to start the TSanM:

1. Double-click **Tivoli Netview** icon on your desktop, or
2. Select **Start -> Programs -> Tivoli Console**.
3. On the menu bar, select **SAN -> Agent Configuration** and add the IP addresses of your fibre channel switches into the SNMP Agents list.
  - a. In the SNMP Agent section of the Agent Configuration panel, select each entry in turn and click Advanced.
  - b. Type the user ID and password for that particular switch (manufacturing defaults are user ID is Admin and password is passw0rd ), this is to allow TSanM to access the switch to collect zoning information.

- c. To enable this access, configure each Fibre Channel switches to allow SNMP Command Access.

Refer to your Fibre Channel switch documentation for the procedure to setup this access.

You can limit the extent of the discovery to just the components in the SAN Volume Controller. On the menu bar, select **Options** -> **Discovery** by editing the seed file to include the IP addresses of the fibre channel switches and the master console.

4. Verify the installation by running a SAN discovery. From the menu bar, click **SAN** -> **Configure Manager**. This displays the Configure Manager panel. Select **Clear History** -> **OK**. Select **Cancel** on the Configure Manager panel.

Ensure that the TSanM discovers all expected Fibre Channel connections and devices. You can visually check that the TSanM discovers all expected connections and device by displaying the topology map for each fabric and seeing that all the expected devices are presented.

**Related topics:**

- “Configuring the master console” on page 57

---

## Setting up Remote Support

This task provides step-by-step instructions about how to set up Remote Support for the Management Node.

To connect to the IBM Remote Support gateway, the Ethernet Port 2 on the master console must have a globally routable IP address (the machine should not be behind a firewall that does IP masquerading). To test for this, ping the gateway by typing `ping 207.25.252.196` in a Windows command prompt window. If the ping test is unsuccessful, there is no access to the specified IP address (no global access). This issue must be resolved by the network administrator.

**Steps:**

If the ping test is successful, set up the routing table by performing the following steps:

1. Open a command line window by selecting **Start** ->**Programs** ->**Accessories** -> **Command Prompt**.
2. Run the command **route -p add DNS Server IP Address (for Local Area Connection2) MASK 255.255.255.255 IP Address of the master console (for Local Area Connection2) MASK 255.255.255.255**, where the italicized text represents valid IP addresses.
3. Run the command **route -p add SAN Volume Controller Subnet MASK 255.255.255.255 IP Address of the master console (for Local Area Connection2) MASK 255.255.255.255**, where the italicized text represents valid IP addresses. The SAN Volume Controller Subnet item refers to the IP subnet where the SAN Volume Controller equipment (for example, FC switches, FAStTs, and SAN Volume Controller nodes) are located. For example, if all the equipment is in the 192.168.1.xx subnet, then SAN Volume Controller Subnet should be replaced by 192.168.1.0.

**Related topics:**

- “Configuring the master console” on page 57

---

## Enhanced remote support configuration

This topic provides information about enhanced remote support configuration.

Enhanced Remote Support provides full desktop GUI access. For details on how to enable Enhanced Remote Support, go to the Web site below and select the **Enhanced Remote Support** topic:

[www.ibm.com/storage/support/2145](http://www.ibm.com/storage/support/2145)

---

## Changing the master console hostname

This topic provides information about changing the hostname of the master console.

### Context:

If you have changed the hostname of the master console, you must modify some of the IBM WebSphere Application Server files that are used by the SAN Volume Controller Console. If you changed the hostname, it was likely when you initially installed the master console.

### Steps:

Perform the following steps to modify the IBM WebSphere Application Server files:

1. Start Windows notepad on the master console.
2. Open the file named: c:\program files\ibm\svcconsole\console\embeddedWAS\config\cells\DefaultNode\nodes\DefaultNode\serverindex.xml.
3. Find the hostName variable. Change the name in quotes to the new hostname of the master console.

The following example shows the line in question:

```
<serverindex:ServerIndex xmi:id="ServerIndex_1" hostName="old_host_name">
```

4. Find the eight other references to the old hostname in the variables called host. Change all of them to the new hostname.

The following example shows one of the lines:

```
<endPoint xmi:id="EndPoint_1" host="old_host_name" port="2809"/>
```

5. Save and close the file.
6. Using notepad open the file called: c:\program files\ibm\svcconsole\console\embeddedWAS\java\jre\lib\orb.properties.

The last line of this file contains a reference to the old hostname.

7. Change the reference from the old hostname to the new hostname.

The following example shows the line in question:

```
com.ibm.CORBA.LocalHost=old_host_name
```

8. Save and close the file.
9. Close notepad.

---

## Overview of IBM Director

This topic provides overview information about the IBM Director.

IBM Director is a systems-management solution that helps administrators manage single or large groups of IBM and non-IBM devices, Management Nodes, and workstations.

All of the functionality of IBM Director is contained in a IBM Director Console that enables single-click and drag-and-drop commands. IBM Director can manage up to 5,000 clients depending on configuration density. Powerful remote management functions include:

- Sophisticated discovery of network components
- Scheduled asset (hardware and software) inventories with persistent storage of data
- Proactive problem notification and tools for problem resolution
- Hardware system component monitors and thresholds to trigger alerts of impending problems
- Alert management with automated actions, manual intervention, or both
- Process scheduling to automate wide-scale client software maintenance (clean up temp files, restart tasks, backups, and so on) according to any timetable
- Help desk and routine maintenance functions such as remote control and file transfer
- Extensive security and authentication

Administrative tasks are performed at the Console. It is a Java application that serves as the user interface to the Director-managed environment. The console provides comprehensive hardware management using a single click or drag-and-drop operation. You can install the console on a machine at a remote location from the server. In addition, there is no limit to the number of IBM Director Consoles that can connect into the Management Node.

**Related topics:**

- “Configuring the master console” on page 57

## IBM Director

This task provides step-by-step instructions about modify your IBM Director settings.

**Steps:**

Modify the following settings.

1. Change the Discovery settings.
  - a. Start IBM Director by clicking the IBM Director Console icon on the desktop.
  - b. Change the user ID to the unique name that you gave to your master console (for example, MASTERC10\ADMINISTRATOR).
  - c. Type the password passw0rd.
  - d. Select **Options -> Discovery Preferences -> SNMP Discovery**.

**Note:** If the Event Action Wizard starts up when you start IBM Director, cancel it and proceed with this procedure.

2. Add the IP addresses of the master console and the switches, using a mask of all zeros; this ensures that all SNMP devices are listed.
  - a. Add each address.

- b. Click **OK** when you are done.
3. Close IBM Director.

**Related topics:**

- “Configuring the master console” on page 57

---

## SAN Volume Controller Call-Home overview using IBM Director

This topic provides an overview of the SAN Volume Controller Call-Home feature using IBM Director.

**Overview:**

Configuring IBM Director for the SAN Volume Controller Error Notification and Call-Home functions, works as follows:

- The SAN Volume Controller raises an SNMP Trap as the result of a detected error.
- The SAN Volume Controller sends its Traps to a particular machine (for example, the master console) that has IBM Director installed.
- IBM Director collects the trap and sends a specifically formatted e-mail to IBM Retain which interrogates the e-mail and generates a Call-Home request in the IBM Call Management System, which is a user specified location (for example, System Admin). In the master console, most of the setup has already been completed, but you need to input your data to complete the e-mail that is sent when an SNMP Trap occurs. To set up the e-mail for IBM Call-Home, see “Configuring IBM Director for the SAN Volume Controller error notification and Call-Home”.

**Related topics:**

- “Configuring the master console” on page 57

---

## Configuring IBM Director for the SAN Volume Controller error notification and Call-Home

This task provides step-by-step instructions about configuring IBM Director for the SAN Volume Controller error notification and Call-Home.

**Steps:**

SAN Volume Controller Call Home works as follows:

1. The SAN Volume Controller raises an SNMP trap as the result of a detected error.
2. The SAN Volume Controller sends its traps to a particular machine (for example, the master console) that has IBM Director installed.
3. IBM Director collects the traps and sends a specifically formatted e-mail to:
  - IBM Retain, which interrogates the e-mail and generates a Call Home request in the IBM Call Management System
  - A user-specified location (for example, the system administrator)

The Call Home configuration for the 2146 has already been completed during manufacture, but you need to input customer-specific data to complete the Call Home e-mail. To set up this e-mail, perform the following steps:

1. Start IBM Director by clicking the IBM Director Console icon on the desktop.
2. From the **IBM Director Console** menu bar, select **Tasks -> Event Action Plan Builder**.
3. In the **Actions** column, expand the item **Send an Internet (SMTP) E-mail**.
4. Right-click **2145CallHome** and select **Update**. The **Customize Action: 2145CallHome** panel displays.
5. Fill in the following items:

#### **Internet E-mail Address**

- Fill in the IBM Retain e-mail address:
  - a. CALLHOME1@de.ibm.com for customers in the USA and Canada
  - b. CALLHOME0@de.ibm.com for the customers in the rest of the world

#### **Reply to**

- Fill in the e-mail address to which you want any replies to be directed.

#### **SMTP E-mail Server**

- Fill in the address of your e-mail server.

#### **SMTP Port**

- Change this, if required, to your SMTP server port number.

#### **Subject of E-mail Message**

- Fill in 2145 Error Notification.

#### **Body of the E-mail Message**

- Fill in the following information:
  - Contact name
  - Contact phone number
  - Offshift phone number
  - Machine location

6. Click **File ->Save** to save the information.
7. Close the Event Action Plan Builder window.
8. Close the IBM Director Console window.

The following is an example of a completed body of the e-mail message:

```
# Contact name = John Doe
# Contact phone number = 546-247-1522
# Offshift phone number = 546-247-1733
# Machine location = Data Centre 1
# Record Type = 1
&iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.1
&iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.2
&iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.3
&iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.4
&iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.5
&iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.6
&iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.7
&iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.8
&iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.9
&iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.10
```

&iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.11  
&iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.12

For more information, refer to the complete Call Home setup procedure in the IBM TotalStorage Virtualization Family SAN Volume Controller Configuration Guide.

---

## Setting up your E-mail notification for the SAN Volume Controller

This task provides step-by-step instructions about setting up your e-mail notification for the SAN Volume Controller.

### Steps:

Perform the following steps to set up your e-mail notification:

1. Start IBM Director by clicking the IBM Director icon on the desktop.
2. From the IBM Director Console menu bar, select **Tasks ->Event Action Plan Builder**.
3. In the **Actions** column, expand the item **Send an Internet (SMTP) E-mail**.
4. Right-click **2145ErrorNot** and select **Update**. The **Customize Action: 2145ErrorNot** panel displays.
5. Fill in the following items:

#### Internet E-mail Address

- Fill in an e-mail address (for example, the e-mail address of the system administrator)

#### Reply to

- Fill in the e-mail address to which you want any replies to be directed.

#### SMTP E-mail Server

- Fill in the address of your e-mail server.

#### SMTP Port

- Change this, if required, to your SMTP server port number.

#### Subject of E-mail Message

- Fill in 2145 Event Notification.

#### Body of the E-mail Message

Fill in the following information:

- # Machine location = Data Centre 1
- 6. Click **File -> Save** to save the information.
- 7. Close the Event Action Plan Builder window.
- 8. Close the IBM Director Console window.

### Result:

This has completed the IBM Director SAN Volume Controller Call-Home set up.

### Related topics:

- “Configuring the master console” on page 57

## Upgrading software on the master console

This topic provides information about upgrading software on the master console.

### Vendor software:

The following table provides information about upgrading your vendor software.

*Table 9. Upgrading vendor software*

Software	Reasons to upgrade
Microsoft Windows 2000 Server Edition and Service Pack	Requires an upgrade only if new functions are required.
Windows 2000 Security Patches	Requires an upgrade only if a problem is found.
Host Bus Adapter driver	Requires an upgrade only if a problem is found or a new function is required.
PuTTY	Required an upgrade only if a problem is found or a new function is required.
Adobe Acrobat Reader	Requires an upgrade only if a problem is found or a new function is required.

### IBM software:

The following table provides information about upgrading your IBM software.

*Table 10. Caption for a matrix data table*

Software	Reasons to upgrade
IBM Director	Requires an upgrade only if a problem is found or a new function is required.
SAN Volume Controller Console	Requires an upgrade only if a problem is found or a new function is required.
IBM FAST Storage Manager	Requires an upgrade only if a problem is found or a new function is required.
Connection Manager	Requires an upgrade only if a problem is found or a new function is required.
Tivoli SAN Manager	Requires an upgrade only if a problem is found or a new function is required.

All the software packages are provided on CDs with the master console installation instructions for the software packages are located in the individual software installation guides.

For recommended IBM upgrades, see the following Web site:

[www.ibm.com/storage/support/2145](http://www.ibm.com/storage/support/2145)

**Note:** Before upgrading the software for the SAN Volume Controller Console, perform the steps outlined in the procedure, Changing the master console hostname.

### Related topics:

- “Configuring the master console” on page 57
- “Changing the master console hostname” on page 67

---

## Troubleshooting master console problems

If you are presented with a dialogue box containing the words: You have signed off. This window will be closed.

### Investigation steps:

Try the following actions to resolve the problem:

#### The problem could be due to:

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



---

## **Part 3. SAN Volume Controller Console**

This part provides detailed information about the SAN Volume Controller Console. More specifically, it provides information about the following:

- Chapter 8, “Getting started with the SAN Volume Controller Console”, on page 77
- Chapter 9, “Overview of creating a cluster using the SAN Volume Controller Console”, on page 81
- Chapter 10, “Scenario: typical usage for the SAN Volume Controller Console”, on page 95
- Chapter 12, “Advanced functions overview for the SAN Volume Controller Console”, on page 111



---

## Chapter 8. Getting started with the SAN Volume Controller Console

This topic provides information about getting started with the SAN Volume Controller Console.

### Overview:

The SAN Volume Controller is provided with a Console, which is Web browser based. It can be used to create and maintain the configuration of storage associated with the SAN Volume Controller. It also provides user management and access to multiple clusters.

The functions that can be performed with the SAN Volume Controller Console:

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

---

## Accessing the SAN Volume Controller Console

This topic provides information about how to access the SAN Volume Controller Console.

The SAN Volume Controller Console is the centralized web application to manage multiple clusters. You access the SAN Volume Controller Console by pointing a web browser at the following URL on your master console:

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

where `<svcconsoleip>` is the IP address of your master console. Log onto the SAN Volume Controller Console using the superuser userid and password.

Log onto the SAN Volume Controller Console using the superuser userid which is superuser and the superuser password which is `passw0rd`. (Upon first access, you will be required to change the superuser password.)

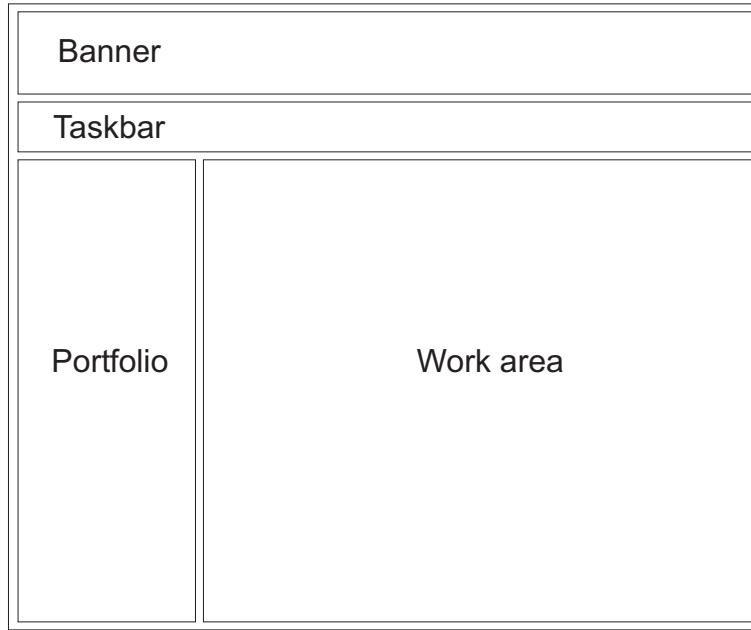
You will use the SAN Volume Controller Console panels to identify SAN Volume Controller clusters in your environment. Once the cluster has been identified, you can use the SAN Volume Controller View Clusters panel to Launch another browser window with specific information for a specific cluster.

---

## SAN Volume Controller Console layout

This topic provides general information about the basic frame layout of the SAN Volume Controller Console.

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



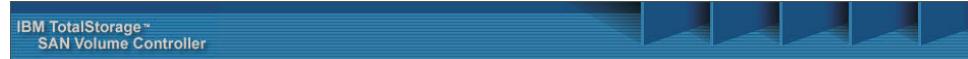
*Figure 17. Basic frame layout*

---

## SAN Volume Controller Console banner area

This topic provides information about the banner area of the SAN Volume Controller Console.

This area is used for product or customer identification.



*Figure 18. Banner area*

---

## SAN Volume Controller task bar

This topic provides information about the task bar of the SAN Volume Controller Console.

The task bar keeps track of all opened primary tasks and allows the user to quickly jump back and forth between them.



Figure 19. Task bar

---

## SAN Volume Controller Console portfolio

This topic provides information about the portfolio area of the SAN Volume Controller Console.

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

---

## SAN Volume Controller Console work area

This topic provides information about the work area of the SAN Volume Controller Console.

The work area of the SAN Volume Controller Console contains product content through panels.

---

## Upgrading the SAN Volume Controller Console software

This topic provides information about upgrading the software for your SAN Volume Controller Console software.

Before upgrading the software for the SAN Volume Controller Console, perform the steps outlined in the procedure called, Changing the master console hostname.

**Related topics:**

- “Changing the master console hostname” on page 67



---

## **Chapter 9. Overview of creating a cluster using the SAN Volume Controller Console**

This topic provides an overview of the panels and information that you will be viewing within the create cluster wizard.

### **Overview:**

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

---

## **Creating a cluster on the SAN Volume Controller Console**

This topic provides information about creating a cluster on the SAN Volume Controller Console.

In order to create a cluster on the SAN Volume Controller Console, you must follow a process of steps.

### **Steps:**

The following list provides the steps that you will need to perform:

1. Create a cluster from the SAN Volume Controller's front panel. A temporary password for the admin user is generated by the node.
2. Access the SAN Volume Controller Console using a web browser.
3. Sign on with the superuser name and password. For first time access, use the superuser name `superuser` and the default password `passw0rd`. You will be required to change this default password. Sign on with the superuser name and password. The Welcome panel is displayed.
4. Add a new cluster to the console.
5. Complete the create cluster wizard:
  - a. Complete the creation of the cluster
  - b. Set up the error logging attributes
  - c. Set up the featurization attributes
  - d. Upload the SSH key

Once these steps are complete and you have exited out of the wizard, you can now use the Web application for the SAN Volume Controller passwords.

### **Related topics:**

- “Accessing the SAN Volume Controller Console” on page 77
- Chapter 9, “Overview of creating a cluster using the SAN Volume Controller Console”

---

## **Prerequisites for creating a cluster using the SAN Volume Controller Console**

This section lists the prerequisites that you must abide by before creating a cluster using the SAN Volume Controller Console.

Ensure that you install the following levels of Web browsers before you connect to the cluster:

- Windows and UNIX operating systems
  - Netscape version 6.2
    - You can get earlier levels from the following Web site:  
<http://wp.netscape.com/downloads/archive.htm>
  - Internet Explorer Version 6+
    - You can get version 6+ from the following Web site:  
<http://www.microsoft.com/windows/ie/downloads/ie6/default.asp>

- AIX operating system
  - You can get AIX Netscape from the following Web site:  
<http://www-1.ibm.com/servers/aix/browsers/>

You must ensure that your proxy setting is disabled. Refer to the appropriate browser and perform the following steps:

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

**Related topics:**

- “Creating a cluster using the SAN Volume Controller Console”

---

## **Creating a cluster using the SAN Volume Controller Console**

This task provides step-by-step instructions about how to configure the cluster.

**Note:** If you are creating a cluster using the SAN Volume Controller Console, then you will need to generate a SSH key pair before performing this task. If you are adding an SSH public key to enable your system to use the command-line interface (CLI), then you will need to have generated an SSH key pair for this system. Once the cluster creation is completed through the SAN Volume Controller Console creation wizard, you can access the same panel directly from the SAN Volume Controller Console by launching the SAN Volume Controller application for the cluster.

**Steps:**

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

1. Start the SAN Volume Controller Console by clicking on the desktop icon or by using your Web browser to go to <http://IPADDRESS:9080/ica>, where *IPADDRESS* is the IP address of the master console. The sign in window displays.
2. The Enter Network Password window is displayed. For the userid, type superuser and password *passw0rd*. The first time you signon as the superuser, you will be required to change the password for the superuser. After you change the password, the welcome window displays.
3. If this is the first time you have accessed the SAN Volume Controller Console proceed as described in 3a otherwise proceed as 3b.
  - a. The welcome screen will be shown as Figure 21 on page 84. Click the **Add SAN Volume Controller Cluster** button.

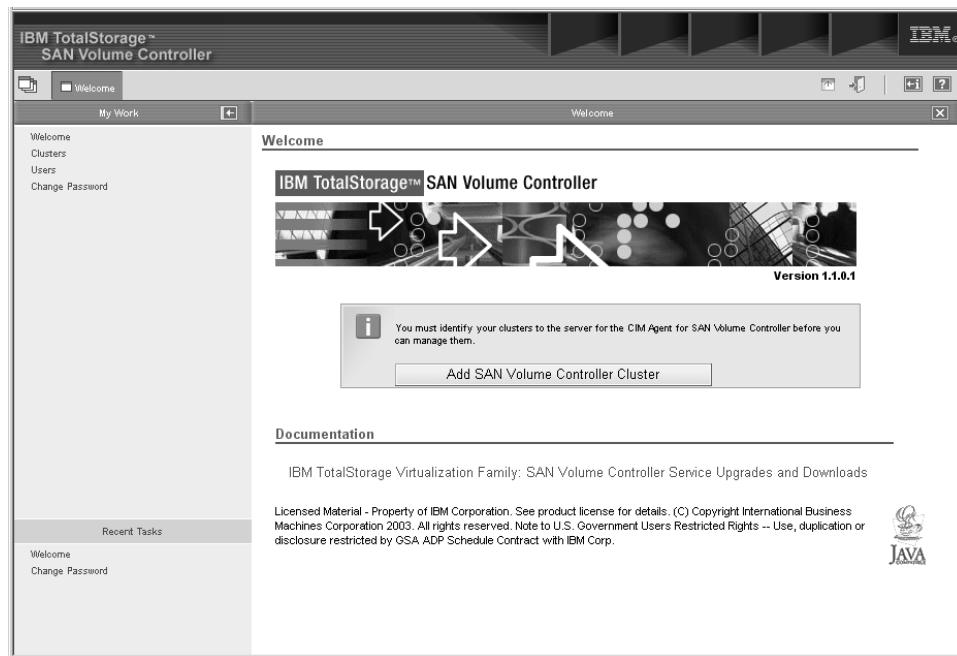


Figure 20. Welcome panel

- b. Select **Clusters** from the portfolio. From the drop down list of tasks select **Add cluster** and click **Go**.

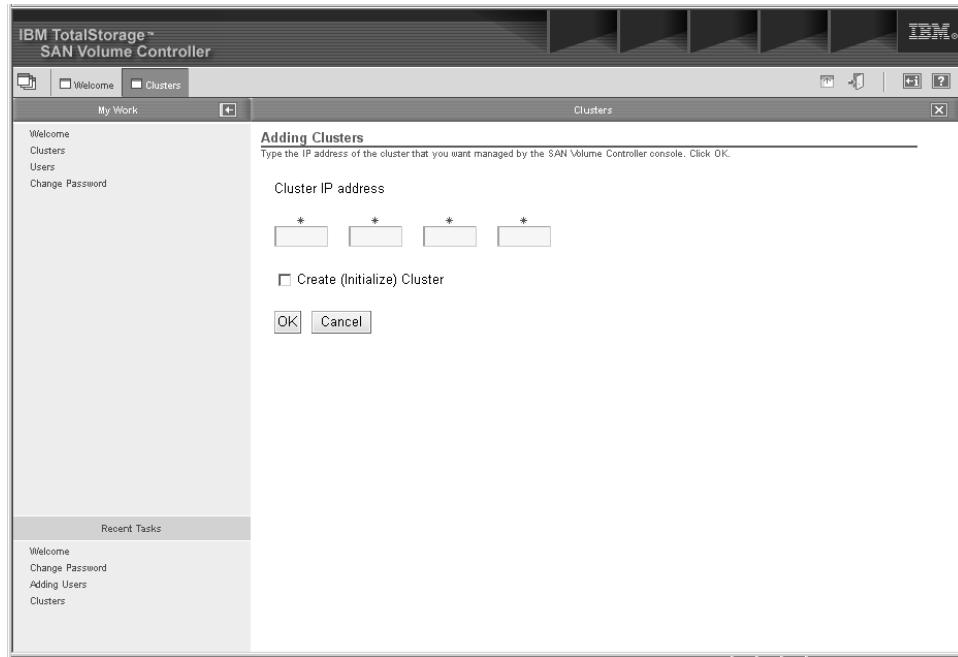


Figure 21. Add Cluster panel

4. Enter the IP address of your cluster.

If the cluster has not been fully created (that is you have just followed the steps in chapter 5 and created the cluster from the front panel) then check the Create (Initialize) Cluster box.

If the cluster is already in use and you are just adding this cluster to the clusters that this installation of the SAN Volume Controller Console is managing, do not check the Create (Initialize) Cluster box. Click **OK**.

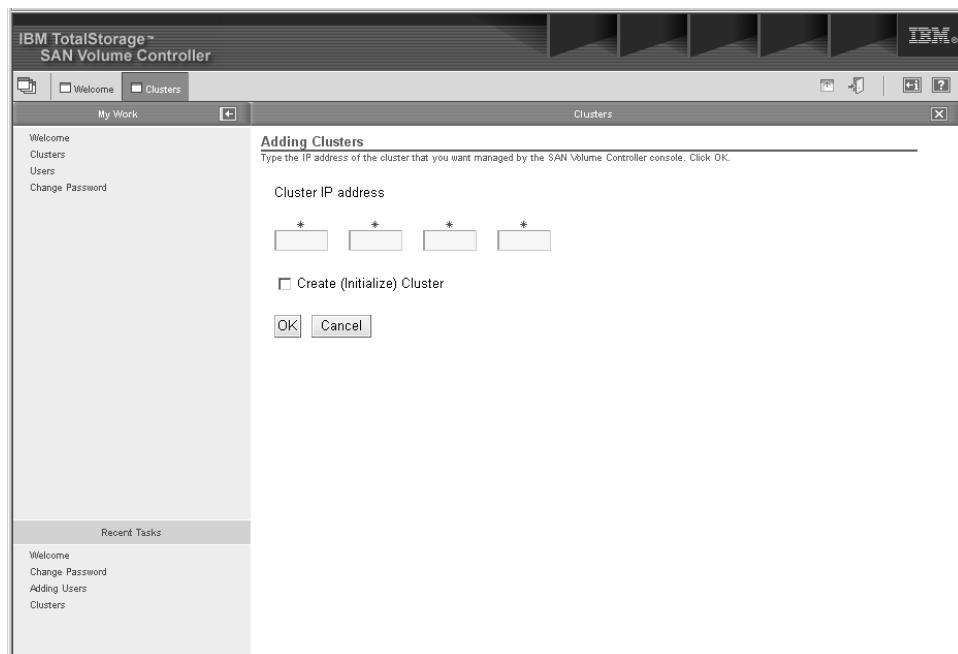


Figure 22. Add Cluster panel

5. You will be prompted to accept the new certificate for the cluster.

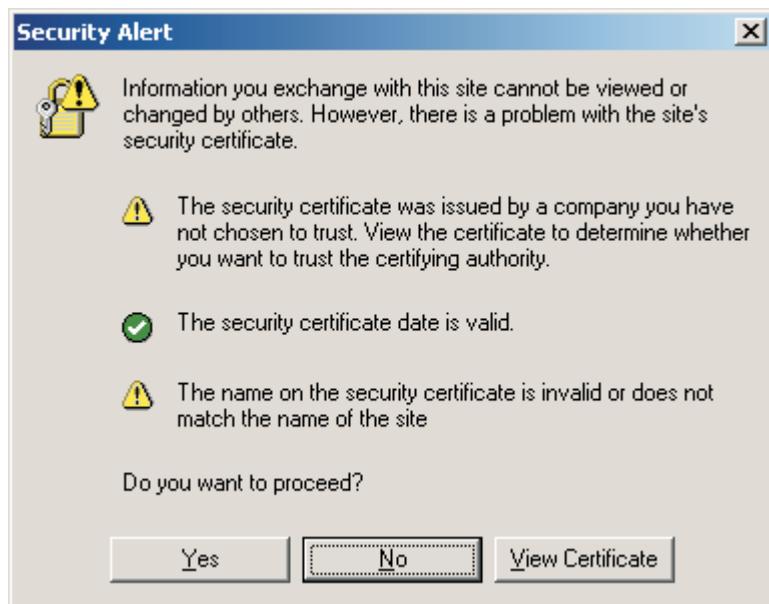


Figure 23. Security alert panel

Click **View Certificate** and on the resulting window click **Install Certificate**.



Figure 24. Certificate Information panel

Click **Next**, **Next**, **Install**, **OK** to complete the install certificate wizard.

Click **OK** to close the Certificate window as shown in Figure 23 on page 85 and click **Yes** to close the Security Alert window as shown in Figure 24.

6. You will be prompted for the cluster user name and password. The username is admin and the password is the one generated by the process in “Overview of passwords” on page 55. Enter the random password that was generated and click **OK**.
7. The create cluster wizard will begin, click **Continue**. If the cluster already existed and you did not check the **Initialize Cluster** checkbox in step 4 on page 84 proceed to step 11 on page 89.
8. Complete the Create New Cluster step by entering a new administrator password and enter a service password. Note these passwords as you will need them in the future to upload new SSH keys via the SAN Volume Controller Console.

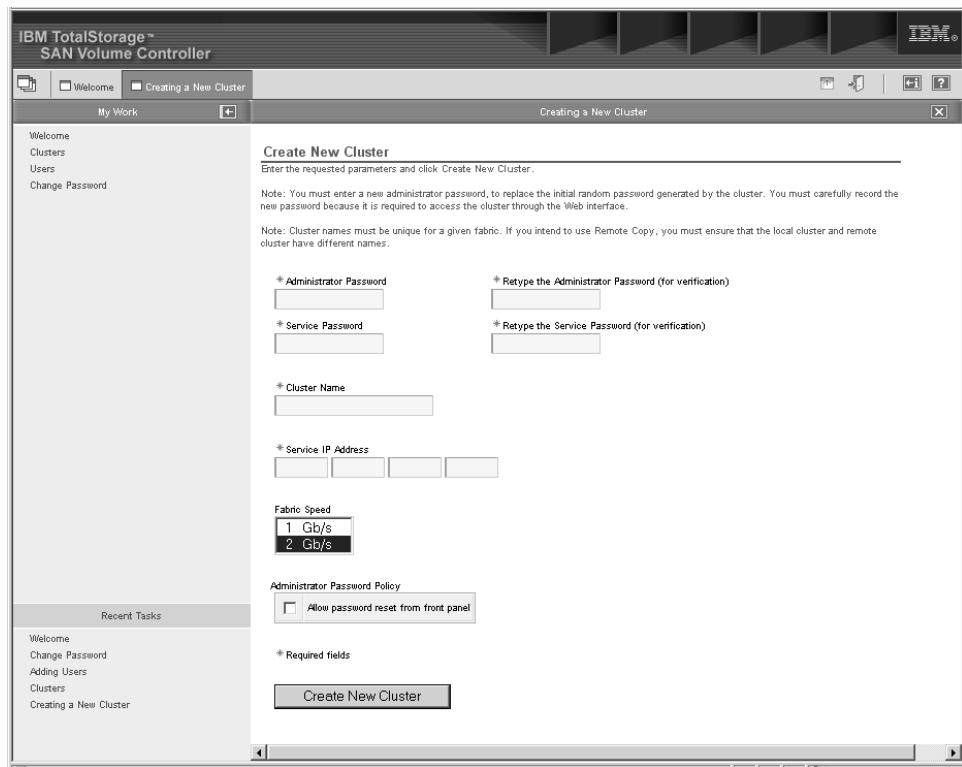


Figure 25. Create New Cluster wizard

- a. Enter a name for your cluster - Note that you cannot subsequently change the name of the cluster so care should be taken when naming it.
- b. Enter the service IP address for the cluster, this is the IP address that will be used if you have to bring a single node up in service mode.
- c. Select the speed of your fabric, either 1 or 2 Gb/s
- d. If you wish to be able to reset the administrator password from the front panel then check the box.
- e. When complete click the Create New Cluster button. The cluster will then be created - this will take a few seconds. When the web page returns, click "Continue"
9. You will then be notified that the password has been changed. Click **Continue** to proceed to the **Error Notification Settings** panel.

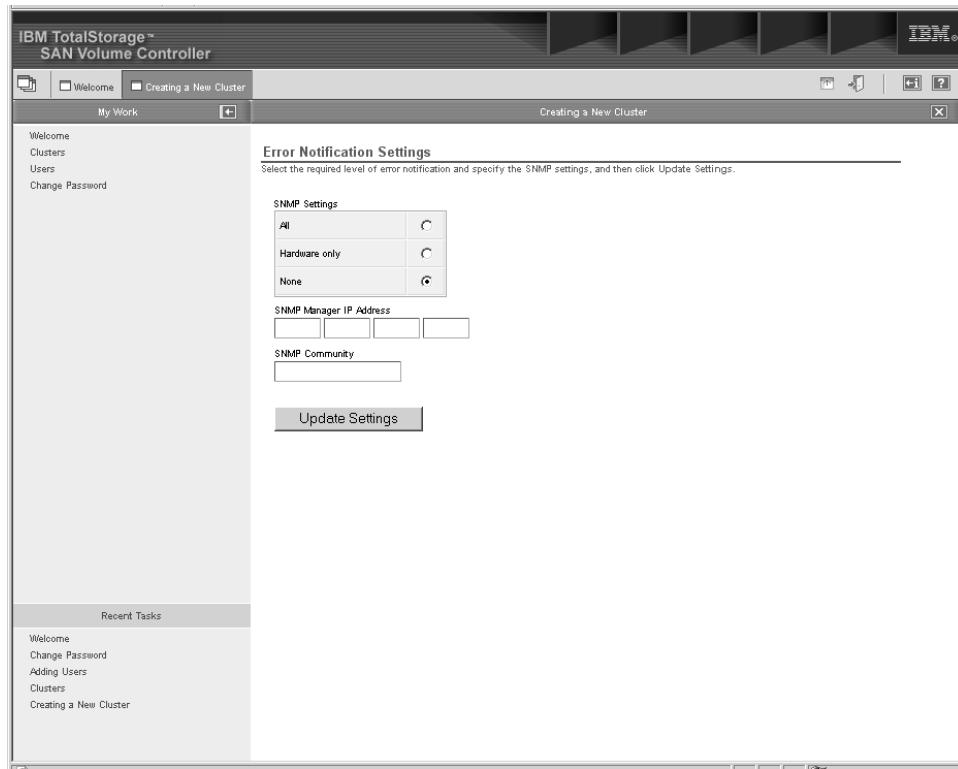
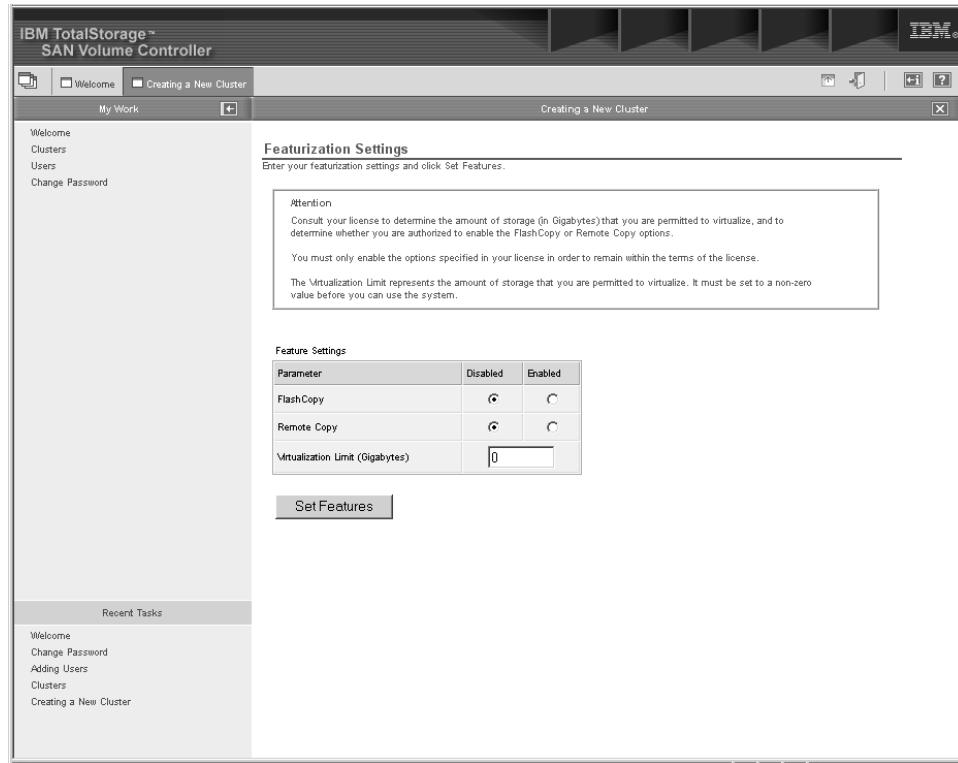


Figure 26. Error Notification Settings panel

- a. If you wish errors to be forwarded as SNMP traps, select either the "All" or "Hardware only" selection boxes. Hardware only will only send SNMP traps for hardware related errors, all will send SNMP traps for all errors, hardware and software.
  - b. Enter the IP address of the machine that is running your SNMP management software (note if you are using IBM Director on the master console to collect SNMP traps, enter the IP address of the master console here)
  - c. Enter the SNMP community name.
  - d. Click Update Settings to continue.
10. Click **Continue**. The Featurization window is displayed.

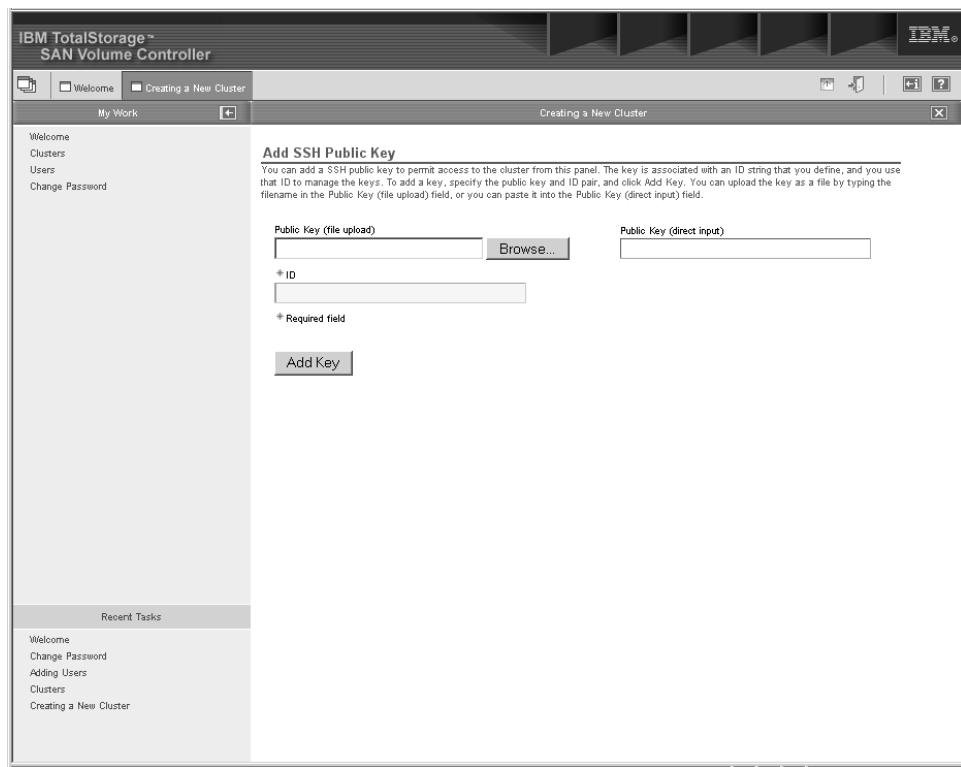


*Figure 27. Featurization Settings panel*

The allowed setting for each of the parameters is specified in your user's license.

- a. Enable the FlashCopy or Remote copy options if they are licensed.
  - b. Enter the virtualization limit as specified in the licence. A zero value is not allowed for this field.
  - c. Click **Set features**. A featurization screen is displayed.
11. Click **Continue** to display the Add SSH Public Key step.

At this point you may be re-prompted for a username and password. Enter **admin** as the user name and enter the new password you gave during 8 on page 86.



*Figure 28. Add SSH public key panel*

Click Browse to locate the public key for the master console.

Enter an ID (label) for this key and click Add Key button

12. Click continue to display the completion page. You should see a picture of a single SAN Volume Controller. Click on the X in the corner of the window to close the wizard.

#### **Result:**

You have now successfully connected and configured the cluster.

#### **Related topics:**

- “Prerequisites for creating a cluster using the SAN Volume Controller Console” on page 81
- Chapter 10, “Scenario: typical usage for the SAN Volume Controller Console”, on page 95

---

## **Launching the SAN Volume Controller Application**

This task provides step-by-step instructions for launching the SAN Volume Controller application.

#### **Steps:**

Perform the following steps to access this panel directly from the SAN Volume Controller Console:

1. Select **Clusters** in the portfolio section of your browser window (the left-hand frame). A new view will be displayed in the work area (main frame).

- Check the small box in the Select column left of the cluster in which you are interested to select that cluster. Select **Launch the SAN Volume Controller Application** in the drop down list box of the work area. Click **Go**. A secondary browser window opens to the SAN Volume Controller (SVC) application.

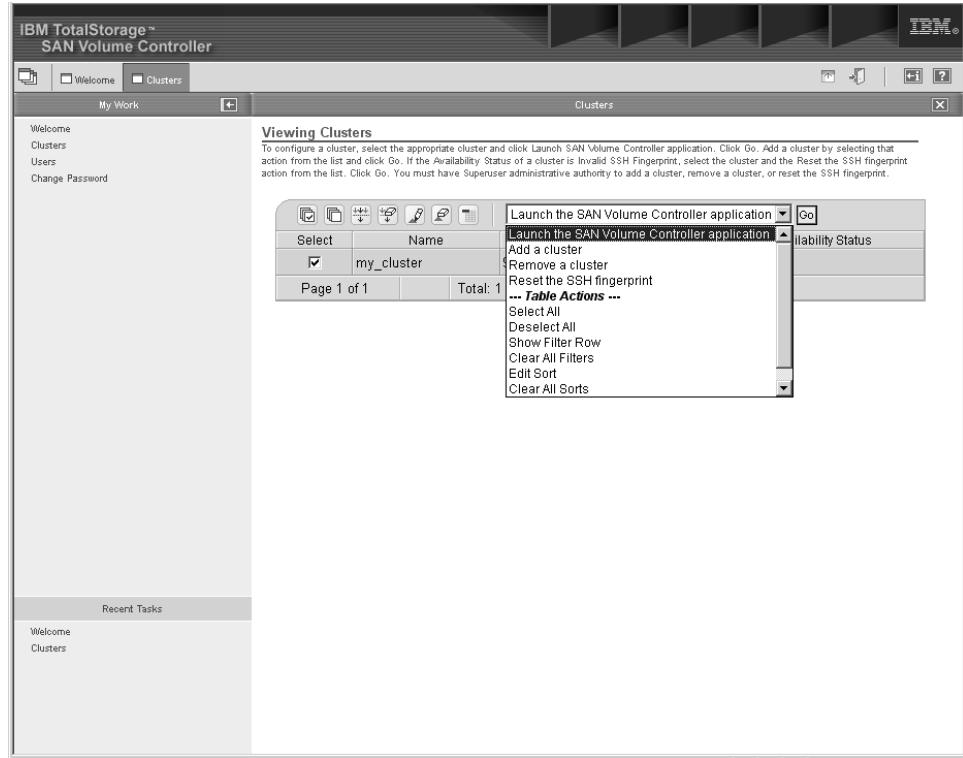


Figure 29. Viewing clusters panels

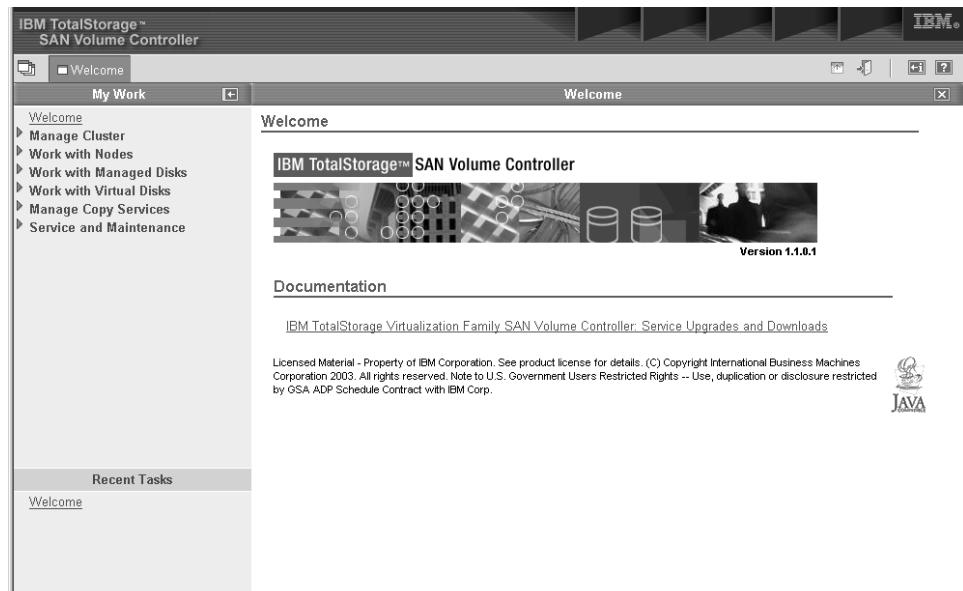


Figure 30. Welcome panel

## Setting the cluster time using the SAN Volume Controller Console

This task provides step-by-step instructions for setting the cluster time using the SAN Volume Controller Console.

### Steps:

Perform the following steps to set the cluster time:

1. Click **Manage Cluster** from the portfolio.
2. Click **Set cluster time** to check the cluster date, time, and time-zone settings.

**Note:** The date and time are used to time-stamp events and errors.

The Cluster date and time settings panel is displayed.

The Set Cluster Time window displays the existing time and time-zone settings

The screenshot shows the 'Cluster date and time settings' panel. At the top, a note says: 'This option displays the existing Cluster Date/Time and Time Zone settings, and allows you to update the values, if required.' Below this is a section titled 'Existing settings' containing a table with three rows: Cluster date (08-Apr-2003), Cluster time (17:29:44), and Cluster time zone (UTC). Below this is a section titled 'New settings' with input fields for Date (01-31), Month (01-12), Year (20xx), Hours (00-23), Minutes (00-59), and Time Zone (UTC). At the bottom are two checkboxes: 'Update cluster time/date' and 'Update cluster time zone', followed by a large 'Update' button.

Figure 31. Cluster date and time settings panel

for the cluster. The time parameters are displayed in a table and are also already entered into several input fields. A list of valid time zones is displayed in a list, and the existing cluster time-zone setting is highlighted in that list.

3. Perform the following steps to change the information in the window:
  - a. Type the changes to any of the input field parameters or select a new time zone from the list.
  - b. When you have made the changes, select the appropriate check-boxes to update the time, the time zone, or both.
  - c. Click **Update** to submit the update request to the node.

## Displaying cluster properties using the SAN Volume Controller Console

This topic provides step-by-step instructions about displaying cluster properties using the SAN Volume Controller Console.

### Steps:

Perform the following steps to display the VPD:

1. Click **Manage Cluster** from the portfolio.
2. Click **View Cluster properties** to view the properties (vital product data (VPD) for the cluster.



Figure 32. View Cluster properties panel

3. Click **IP Addresses** to view the cluster level information such as the IP address, the service IP address, the subnet mask, and the default gateway addresses.
4. Click **Space** to view the space and capacity within the VDisks and MDisk groups.
5. Click **SNMP** to view the SNMP details.
6. Click **Statistics** to view the cluster statistics details.



---

## Chapter 10. Scenario: typical usage for the SAN Volume Controller Console

This topic provides a hypothetical example of configuring your SAN Volume Controller using the SAN Volume Controller Console. The main focus of the following example is to provide storage to your host system. Our hypothetical example is the following:

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

The following steps illustrates how this can be done:

1. Create a cluster.
2. Configure the cluster with an IP address of 9.20.123.456, a fabric speed of 2 Gb/s. Name the cluster examplecluster.
3. Launch the SAN Volume Controller application for the cluster. A secondary browser window opens to the SAN Volume Controller (SVC) Web application. Now you can work with the specific SAN Volume Controller cluster which you selected.
4. Add nodes
  - knode and lnode to the I/O group called io\_group0 in the examplecluster cluster
  - mnode and nnode to the I/O group called io\_group1 in the examplecluster cluster
5. Create the managed disk (MDisk) groups maindiskgroup and bkpdiskgroup
6. Create 4 virtual disks (VDisks)
  - 2 VDisks from maindiskgroup
  - 2 VDisks from bkpdiskgroup
7. Create 2 host objects
  - a host object called demohost1 with HBAs that have WWPNs of 210100e08b251dd4, and 210100e08b251dd5
  - a host object called demohost2 with HBAs that have WWPNs of 210100e08b251dd6, and 210100e08b251dd7
8. Create the VDisk-to-host mappings
  - Create a VDisk-to-host mapping for demohost1
  - Create a VDisk-to-host mapping for demohost2

Once this step is complete, you have then successfully created storage on your host system.

9. Create a FlashCopy consistency group called maintobkpfcopy and add the 2 FlashCopy mappings to it

**Related topics:**

- “Creating a cluster using the SAN Volume Controller Console” on page 82
- “Adding nodes to the cluster”
- “Create managed disk (MDisk) groups” on page 101
- “Create virtual disks (VDisks)” on page 103
- “Creating host objects” on page 104
- “Create VDisk-to-host mappings” on page 105
- “Create a FlashCopy mapping” on page 106
- “Create a FlashCopy consistency group” on page 106

---

## Adding nodes to the cluster

This task provides step-by-step instructions about how to add nodes to the cluster using the SAN Volume Controller Console. You must add at least one other node to the cluster so it has a minimum of two nodes.

**Prerequisites:**

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

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

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

**Special procedures when adding a node to a cluster:**

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

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

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

```
svcinfo lsnode
```

or, if this information is not available, then

- *Before* the node is added back into the cluster all the hosts using the cluster must be shut down.

The node must then be added before the hosts are rebooted or, if the I/O group information is not available and it is inconvenient to shutdown and reboot all of the hosts using the cluster, then

- On all the hosts connected to the cluster, unconfigure the Fibre Channel adapter device driver, the disk device driver, and the SDD device driver, before you add the node to the cluster.

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

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

#### Hypothetical scenarios where the special procedures may apply:

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

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

#### Background:

Applications on host systems direct I/O operations to filesystems or logical volumes which are mapped by the operating system to vpaths which are pseudo disk objects supported by the SDD driver. See the *IBM Subsystem Device Driver (SDD) User's Guide*, SC26-7540.

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

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

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

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

#### Data Corruption Scenario:

Consider a four-node SAN Volume Controller configuration.

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

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

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

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

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

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

#### Context:

Assume that the cluster has been created.

#### Steps:

Perform the following steps to add nodes to the cluster:

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



Figure 33. Nodes panel

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

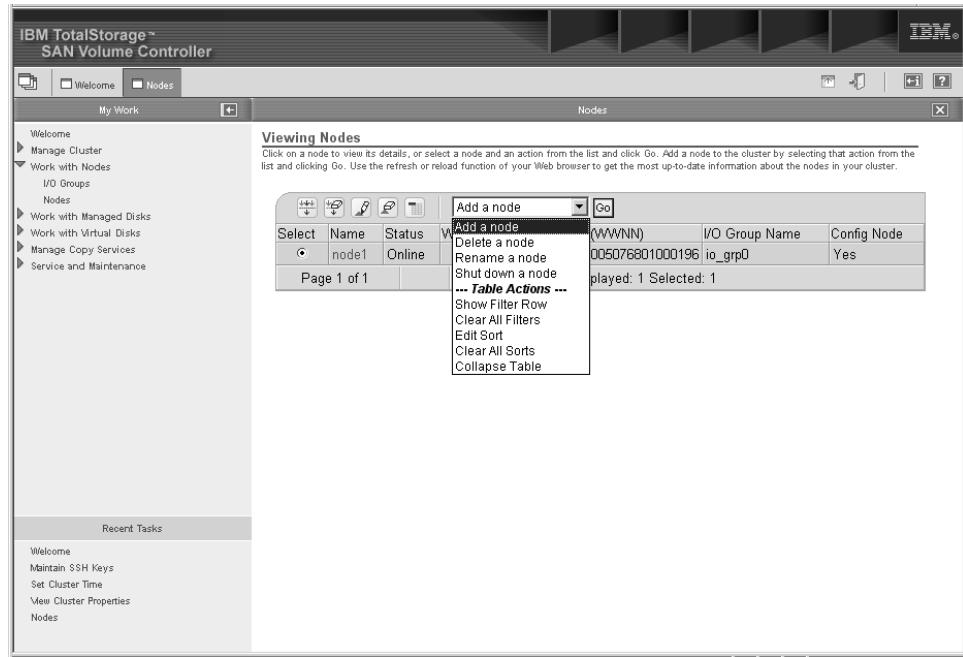


Figure 34. Add Node drop down list

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

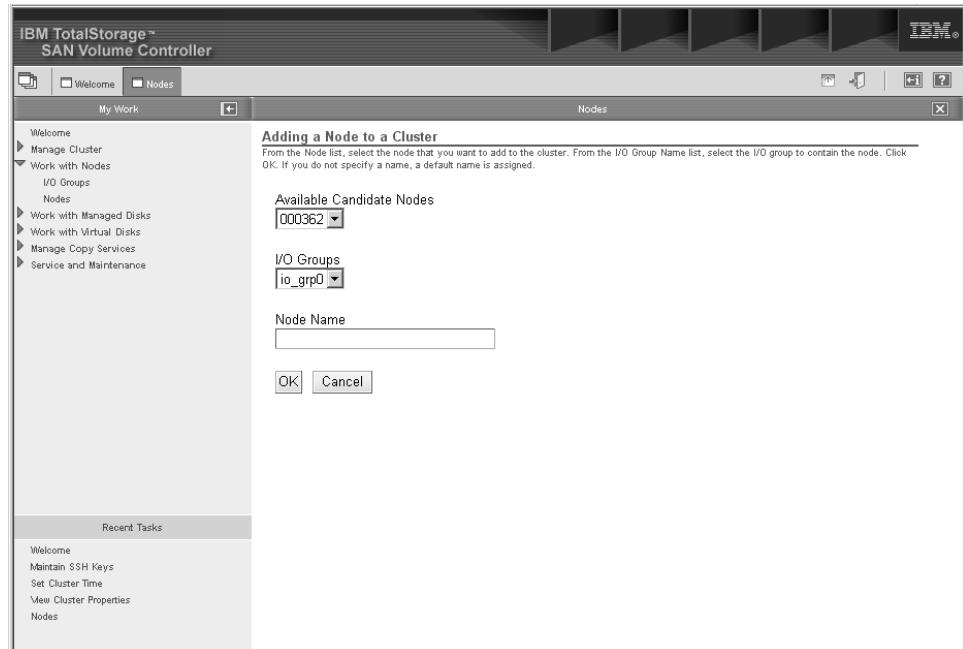


Figure 35. Add Node to Cluster panel

#### Notes:

- a. Each node in an I/O group must be connected to a different uninterruptible power supply.

- 
- b. If you do not supply a name, the cluster will assign a default name to the object. It is recommended that wherever possible you provide a meaningful name for objects to aid object determination in the future.

**Example:**

In our hypothetical scenario, the nodes are called:

knod and lnode

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

io\_group0

In our hypothetical scenario, the nodes are called:

mnode and nnode

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

io\_group1

- 
- 
- 
- 
5. Repeat step 4 on page 99 for each of the nodes that you want to add to the cluster.

**Related topics:**

- Chapter 10, “Scenario: typical usage for the SAN Volume Controller Console”, on page 95

---

## Displaying node properties using the SAN Volume Controller Console

This topic provides step-by-step instructions about displaying the node properties using the SAN Volume Controller Console.

**Steps:**

Perform the following steps to display the node properties:

1. Click **Work with Nodes** from the portfolio.
2. Click **Nodes** from the portfolio. The Nodes panel is displayed. See Figure 33 on page 98 for more information.
3. Select the name of the node that you want to view the details for. The Viewing general details panel is displayed.

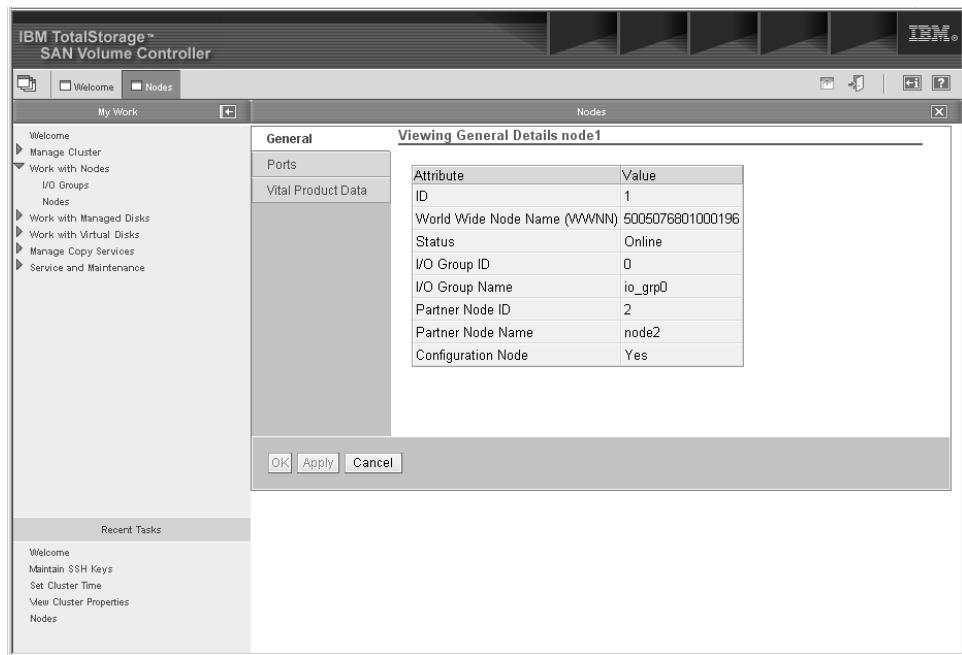


Figure 36. Viewing general details panel

4. Click **Ports** to view the WWPN port details. The Viewing port details node1 panel is displayed.
5. Click **Vital Product Data** to view the node hardware details. The Viewing vital product data panel is displayed.

## Create managed disk (MDisk) groups

This task provides step-by-step instructions about how to create managed disk (MDisk) groups using the SAN Volume Controller Console.

### Steps:

Perform the following steps to create a MDisk group:

1. Click **Work with Managed Disks** in the portfolio.
2. Click **Managed Disk Groups** in the portfolio. The Filtering managed disk (MDisk) groups panel is displayed.

**Note:** The filter panels can be used to pre-filter the list of objects that are displayed. This will reduce the number of objects returned to the SAN Volume Controller Console. This can be useful when you have a large number of objects (for example 4096 Mdisks or 1024 VDisks) and do not wish to display them all. You can bypass the filtering and display all objects by clicking the **Bypass Filter** button.

3. Specify the filter criteria that you want to use and click **OK** or click the **Bypass Filter** button to display all objects of this type. The Managed Disk (MDisk) Groups Filter panel is displayed.

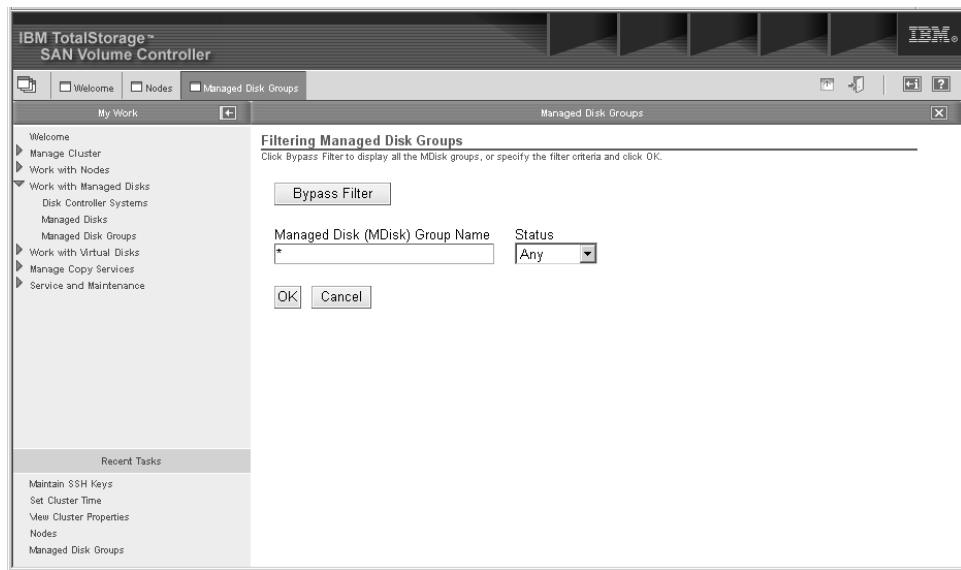


Figure 37. Managed Disk (MDisk) Groups Filter panel

4. Select **Create MDisk Group** from the list. Click **Go**. The Creating Managed Disk (MDisk) Groups panel is displayed.
5. Type the name of the MDisk group, add MDisks from the **MDisk Candidates** list.

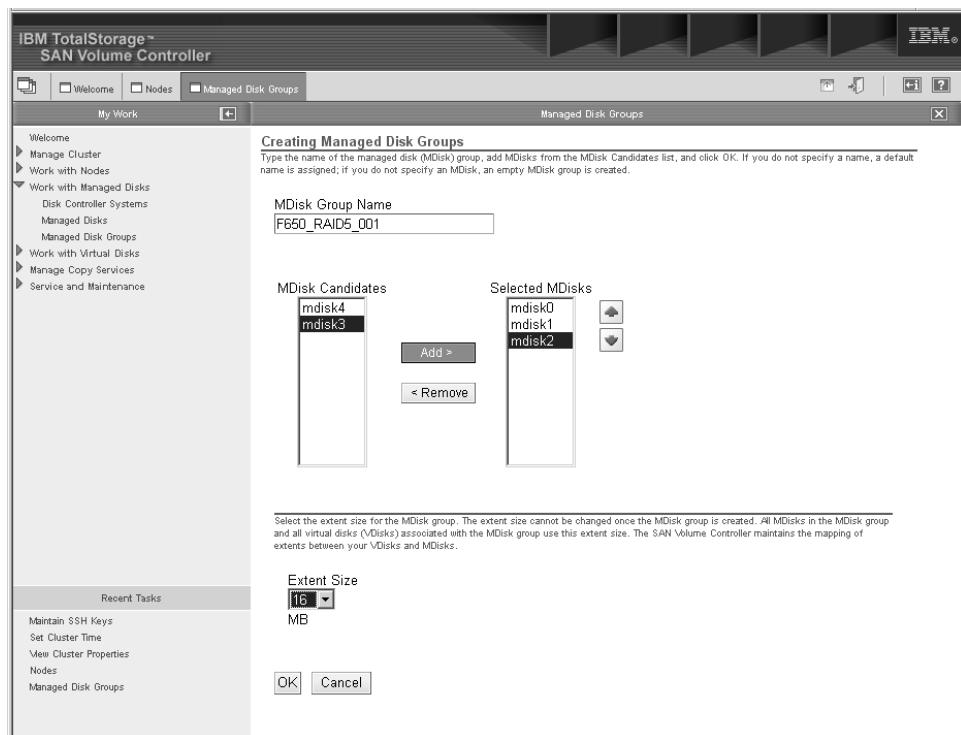


Figure 38. Creating MDisk groups panel

#### Example:

In our hypothetical scenario, type  
maindiskgroup

- add MDisks  
 mdsk0, mdsk1, mdsk2, mdsk3
- from the **MDisk Candidates** list.
6. Select the extent size from the list.
- Example:**
- In our hypothetical scenario, select  
 32
- for the extent size used within this MDisk group and click **OK**.
7. Repeat steps 4 on page 102 through 6 for all of the MDisk groups that you want to create.
- Example:**
- In our hypothetical scenario, repeat steps 4 on page 102 through 6, in which the second MDisk group is named  
 bkpdiskgroup
- with the following MDisks attached,  
 mdsk4, mdsk5, mdsk6, mdsk7
- The extent size will be  
 32
- MB.

**Related topics:**

- Chapter 10, “Scenario: typical usage for the SAN Volume Controller Console”, on page 95

## Create virtual disks (VDisks)

This task provides step-by-step instructions about how to create virtual disks (VDisks) using the SAN Volume Controller Console.

**Steps:**

Perform the following steps to create VDisks:

1. Click **Work with Virtual Disks** from the portfolio.
2. Click **Virtual Disks** from the portfolio. The Filtering virtual disks (VDisks) panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The Viewing VDisks panel is displayed.
4. Select **Create VDisk** from the list. Click **Go**. The Create virtual (VDisks) wizard is displayed.
5. Complete the Create virtual (VDisk) wizard. Perform the following to complete the Create virtual (VDisk) wizard:
  - specify the name of the VDisk

**Example:**

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

mainvdisk1

- select the I/O group and MDisk group that the VDisk will be a part of

**Example:**

In our hypothetical scenario, the I/O group that the VDisk uses is:

`iogroup0`

In our hypothetical scenario, the MDisk group that the VDisk uses is:

`maindiskgroup`

- select the type of VDisk you want to create

**Example:**

In our hypothetical scenario, the VDisk type is:

`striped`

- assign certain MDisks to the VDisk

**Example:**

In our hypothetical scenario, our VDisk type is striped, therefore we need to select both MDisks:

`mdsk0` and `mdsk1`

6. Complete the wizard for every VDisk that you need to create.

**Example:**

In our hypothetical scenario, you will create three more VDisks with the name:  
`mainvdisk2`, `bkpvdisk1`, `bkpvdisk2`

The I/O group that the VDisks use are:

`iogroup0`, `iogroup1`, `iogroup1`

The MDisk group that the VDisks use are:

`maindiskgroup`, `bkpdiskgroup`, `bkpdiskgroup`

The VDisk type for all three VDisks are:

`striped`

The MDisks that uses the VDisks are:

`mdsk2`, `mdsk3...mdsk4`, `mdsk5...mdsk6`, `mdsk7`

**Related topics:**

- Chapter 10, “Scenario: typical usage for the SAN Volume Controller Console”, on page 95

---

## Creating host objects

This task provides step-by-step instructions about how to create host objects using the SAN Volume Controller Console.

**Steps:**

Perform the following steps to create host objects:

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

- Type the name of the logical host object, assign World Wide Port Name (WWPN) to it, and click **OK**. If you don't specify a name, a default name is assigned.

**Example:**

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

host0

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

210100e08b251dd4, 210100e08b251dd5

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

- Repeat steps 4 on page 104 through 5 for each host object that you want to create.

**Example:**

In our hypothetical scenario, repeat steps 4 on page 104 through 5 and name the host:

demohost2

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

210100e08b251dd6, 210100e08b251dd7

**Related topics:**

- Chapter 10, "Scenario: typical usage for the SAN Volume Controller Console", on page 95

---

## Create VDisk-to-host mappings

This task provides step-by-step instructions about how to create virtual disk (VDisks) to host mappings using the SAN Volume Controller Console.

**Steps:**

Perform the following steps to create VDisks to host mappings:

- Click **Work with Virtual Disks** from the portfolio.
- Click **Virtual Disks** from the portfolio. The Filtering Virtual Disks (VDisks) panel is displayed.
- Specify the filter criteria that you want to use. Click **OK**. The Viewing VDisks panel is displayed.
- Select the VDisk that you want to map to your host.
- Select **Map VDisk to Host** from the list. Click **Go**. The Map VDisk to Host panel is displayed.
- Select the host to which you want to map this VDisk, and click **OK**. If you don't specify a SCSI LUN ID, the SAN Volume Controller assigns the next available SCSI LUN ID on the host adapter.
- Repeat steps 3 through 6 for every VDisk that you want to map to your host.

**Related topics:**

- Chapter 10, "Scenario: typical usage for the SAN Volume Controller Console", on page 95

---

## Create a FlashCopy consistency group

This task provides step-by-step instructions about how to create a FlashCopy consistency group using the SAN Volume Controller Console.

### Steps:

Perform the following steps to create a FlashCopy consistency group:

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

### Example:

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

`maintobkpfcopy`

The mappings that should be added are:

`main1copy, main2copy`

**Note:** You could have created the FlashCopy consistency group before you created the mappings and then added the FlashCopy mappings to the consistency group.

### Related topics:

- Chapter 10, “Scenario: typical usage for the SAN Volume Controller Console”, on page 95

---

## Create a FlashCopy mapping

This task provides step-by-step instructions about how to create a FlashCopy mapping using the SAN Volume Controller Console.

### Steps:

Perform the following steps to create a FlashCopy mappings:

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

### Example:

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

`main1copy`

6. Select the source VDisk from the list.

**Example:**

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

7. Select the target VDisk from the list.

**Example:**

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

8. Select the priority for the background copy.

**Example:**

In our hypothetical scenario, the priority for the background copy is:  
`75`

percent.

9. Repeat steps 4 on page 106 through 8 for each FlashCopy that you want to create.

**Example:**

In our hypothetical scenario, the second FlashCopy mapping is named:  
`main2copy`

The name of the source VDisk is:

`maindisk2`

The name of the target VDisk is:

`bkpdisk2`

The priority for the background copy is:

`50`

percent.

**Related topics:**

- Chapter 10, “Scenario: typical usage for the SAN Volume Controller Console”, on page 95



---

# Chapter 11. Advanced function FlashCopy overview

This topic provides an overview about the advanced function FlashCopy overview.

## Overview:

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

## Related topics:

- Appendix B, "Valid combinations of FlashCopy and Remote Copy functions", on page 263

---

## Starting FlashCopy mappings

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

Perform the following steps to start a FlashCopy mapping:

1. Click **Manage Copy Services** in the portfolio.
2. Click **FlashCopy Mappings** in the portfolio. The Filtering FlashCopy mappings panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The FlashCopy mappings panel is displayed.
4. Click **Start Mapping**. The Starting FlashCopy mappings` panel is displayed.

---

## Stopping FlashCopy mappings

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

## Steps:

Perform the following steps to stop a FlashCopy mapping:

1. Click **Manage Copy Services** in the portfolio.
2. Click **FlashCopy Mappings** in the portfolio. The Filtering FlashCopy Mappings panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The FlashCopy Mappings panel is displayed.
4. Click **Stop Mapping**. The Stopping FlashCopy mappings panel is displayed.

---

## Deleting FlashCopy mappings

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

Perform the following steps to delete a FlashCopy mapping:

1. Click **Manage Copy Services** in the portfolio.
2. Click **FlashCopy mappings** in the portfolio. The Filtering FlashCopy mappings panel is displayed.

3. Specify the filter criteria that you want to use. Click **OK**. The FlashCopy mappings panel is displayed.
4. Click **Delete Mapping**. The Deleting FlashCopy mapping panel is displayed.

---

## Starting FlashCopy consistency groups

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

### Steps:

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

1. Click **Manage Copy Services** in the portfolio.
2. Click **FlashCopy Consistency Groups** in the portfolio. The Filtering FlashCopy Consistency Groups panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The FlashCopy Consistency Groups panel is displayed.
4. Click **Start Consistency Group**. The Starting FlashCopy Consistency Groups panel is displayed.

---

## Stopping FlashCopy consistency groups

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

### Steps:

Perform the following steps to stop a FlashCopy consistency group:

1. Click **Manage Copy Services** in the portfolio.
2. Click **FlashCopy Consistency Groups** in the portfolio. The Filtering FlashCopy Consistency Groups panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The FlashCopy Consistency Groups panel is displayed.
4. Click **Stop Consistency Group**. The Stopping Consistency Groups panel is displayed.

---

## Deleting consistency groups

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

### Steps:

Perform the following steps to delete a FlashCopy consistency groups:

1. Click **Manage Copy Services** in the portfolio.
2. Click **FlashCopy consistency groups** in the portfolio. The Filtering FlashCopy consistency groups panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The FlashCopy consistency groups panel is displayed.
4. Click **Delete Consistency Groups**. The Delete Consistency Groups panel is displayed.

---

## Chapter 12. Advanced functions overview for the SAN Volume Controller Console

This topic provides overview information about the advanced functions that you are enable to perform using the SAN Volume Controller Console. More specifically, it provides step-by-step instructions for the following advanced functions:

- “Guidelines for MDisk group creation”
- “Determining a nodes WWPNs using the SAN Volume Controller Console” on page 112
- “Determining a storage controller name from its SAN Volume Controller name” on page 112
- “Determining the relationship between VDisks and MDisks using the SAN Volume Controller Console” on page 112
- “Determining the relationship between MDisks and RAID arrays or LUNs using the SAN Volume Controller Console” on page 113
- “Increasing the size of your cluster using the SAN Volume Controller Console” on page 113
- “Adding a node to increase the size of your cluster” on page 114
- “Migrating a VDisk to a new I/O group” on page 114
- “Replacing a faulty node in the cluster using the SAN Volume Controller” on page 115
- “Recovering from offline VDisks after a node or an I/O group failed” on page 116
- “Replacing an HBA in a host using the SAN Volume Controller Console” on page 117
- “Adding a new storage controller to a running configuration” on page 118
- “Removing a storage controller” on page 119
- “Expanding a VDisk using the SAN Volume Controller Console” on page 120
- “Shrinking a VDisk using the SAN Volume Controller Console” on page 122
- “Migrating VDisks between MDisk groups” on page 123
- “Advanced function Remote Copy overview” on page 125
- “Advanced function cluster overview” on page 126

---

### Guidelines for MDisk group creation

This topic provides information about the guidelines that you will need to follow to create an MDisk group.

When creating managed disk groups and adding managed disks to the groups, you should note the following:

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

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

---

## Determining a nodes WWPNs using the SAN Volume Controller Console

This task provides step-by-step instructions for determining a nodes WWPNs using the SAN Volume Controller.

### Steps:

Perform the following steps to determine a nodes WWPNs:

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

---

## Determining a storage controller name from its SAN Volume Controller name

This task provides step-by-step instructions for determining a storage controller name from its SAN Volume Controller name.

### Steps:

Perform the following steps to determine the storage controller name:

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

---

## Determining the relationship between VDisks and MDisks using the SAN Volume Controller Console

This task provides step-by-step instructions for determining the relationship between VDisks and MDisks using the SAN Volume Controller Console.

### Steps:

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

1. Click **Work with VDisks** from the portfolio.
2. Select the VDisk that you want to view the relationship between this VDisk and its MDisks.
3. Select the **Show MDisks** task. The Work with MDisks panel is displayed. This panel lists the MDisks and make up the selected VDisk.

#### **Steps:**

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

1. Click **Work with MDisks** from the portfolio.
2. Select the VDisk that you want to view the relationship between this VDisk and its MDisks.
3. Select the **Show VDisks** task. The Work with VDisks panel is displayed. This panel lists the VDisks and make up the selected MDisk.

---

## **Determining the relationship between MDisks and RAID arrays or LUNs using the SAN Volume Controller Console**

This task provides step-by-step instructions for determining the relationship between MDisks and RAID arrays or LUNs using the SAN Volume Controller Console.

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

#### **Steps:**

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

1. Click **Work with MDisks** from the portfolio.
2. Select the MDisk to view the details. Write down the controller name and controller LUN number.
3. Click the **Work with Disk Controllers** panel.
4. In the filter screen, enter the controller name in the **Name** field. The panel displayed should show just the one controller.
5. Select the name to show the detailed view of the given controller. Write down the vendor ID and the product ID and WWNN and use these to determine the controller that is presented to the MDisk.
6. From the native user interface for the given controller, list the LUNs it is presenting and match the LUN number with that noted in 2. This will tell you the exact RAID array and partition that corresponds with the MDisk.

---

## **Increasing the size of your cluster using the SAN Volume Controller Console**

This task provides step-by-step instructions for increasing the size of your cluster.

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

#### **Steps:**

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

1. Perform the steps in the “Adding a node to increase the size of your cluster” section and repeat this procedure for the second node.
2. If you wish to balance the load between the existing I/O groups and the new I/O groups, follow the “Migrating a VDisk to a new I/O group” procedure. Repeat this procedure for all VDisks you want to assign to the new I/O group.

## **Adding a node to increase the size of your cluster**

This task provides step-by-step instructions for adding a node to increase the size of your cluster.

### **Steps:**

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

1. Click **Work with I/O groups** to determine which I/O group you wish to add the nodes to.
2. Look for the first I/O group listed that has a node count = 0. Write down the I/O group name or ID. You will need it in the following step.
3. Add the node back into the cluster by selecting **Add Node** task from the **Work with Nodes** panel.
4. Select the node from the list of candidate nodes and select the I/O group from the list.
5. Optionally enter a node name for this node.
6. Verify that the node is online by refreshing the **Work with Nodes** panel. You may need to close the panel and reopen it for the refresh to take effect.
7. You may also need to modify the configuration of your disk controllers. If your controller uses a mapping technique to present its RAID arrays or partitions to the cluster, you will need to modify the port groups that belong to the cluster because the WWNN and WWPN’s of the node have changed. See Chapter 25, “Configuring disk controllers”, on page 223 for more information.

## **Migrating a VDisk to a new I/O group**

This task provides step-by-step instructions for migrating a VDisk to a new I/O group to increase the size of your cluster using the SAN Volume Controller Console.

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

**Attention:** This is a disruptive procedure, access to the VDisk will be lost while you follow this procedure.

### **Steps:**

Perform the following steps to migrate a single VDisk:

1. Quiesce all I/O operations for the VDisk. You may need to determine the hosts that are using this vdisk. See “Determining the relationship between VDisks and MDisks using the SAN Volume Controller Console” on page 112 for more information.
2. Before migrating the VDisk, it is essential that for each vpath presented by the VDisk you intend to move, the SDD configuration is updated to remove the

vpaths in question. Failure to do this may result in data corruption. See *IBM Subsystem Device Driver (SDD) User's Guide* for details about how to dynamically reconfigure SDD for the given host operating system.

3. Any FlashCopy mappings or Remote Copy relationships that use this VDisk should be stopped and deleted. To check if the VDisk is part of a mapping or relationship, perform the following steps:
  - a. Click **Work with VDisks**.
  - b. Click on the VDisk name to view the details.
  - c. Look for the **FlashCopy ID** and **Remote Copy ID** fields. If these are not blank then the VDisk is part of a mapping or relationship.
- See "Advanced function Remote Copy overview" on page 125 for details about how to stop or delete the mapping or relationship.
4. Migrate the VDisk by selecting the VDisk from the **Work with VDisks** panel and selecting the **Modify** task. Change only the I/O group to the new I/O group name.
5. It is now necessary to follow the SDD procedure to discover the new vpaths and to check that each vpath is now present with the correct number of paths. See the *IBM Subsystem Device Driver (SDD) User's Guide* for details on how to dynamically reconfigure SDD for the given host operating system.

---

## Replacing a faulty node in the cluster using the SAN Volume Controller

This task provides step-by-step instructions for replacing a faulty node in the cluster using the SAN Volume Controller.

### Steps:

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

1. Verify the name and ID of the node that you wish to remove. Click **Work with Nodes**.
  - a. If the node is faulty, it will be shown as offline. Ensure the partner node in the I/O group is online.
    - 1) If the other node in the I/O group is offline, start Directed Maintenance Procedures to determine the fault.
    - 2) If you have been directed here by the DMPs, and subsequently the partner node in the I/O group has failed, see "Recovering from offline VDisks after a node or an I/O group failed" on page 116.
  - b. If you are replacing the node for other reasons, determine the node you wish to replace and again ensure the partner node in the I/O group is online.
    - 1) If the partner node is offline, you will lose access to the VDisks that belong to this I/O group if you continue. Start the Directed Maintenance Procedures and fix the other node before proceeding.
2. Having noted the <nodename> in step 1, remove the node from the cluster by selecting the node in the list and click on **Remove Node** task.
3. If the nodes were 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, it will be necessary to follow the SDD procedure to discover the new paths and to check that each vpath is now

- presenting the correct number of paths. See the *IBM Subsystem Device Driver (SDD) User's Guide* for details about adding paths to existing vpaths.
- b. You may also need to modify the configuration of your disk controllers. If your controller uses a mapping technique to present its RAID arrays or partitions to the cluster, you will need to modify the port groups that belong to the cluster because the WWNN or WWPN's of the node have changed. See Chapter 25, "Configuring disk controllers", on page 223 for more information.
  - 4. Add the node back into the cluster. Select the **Add Node** task from the drop down list. Select the candidate node from the list. Select the I/O group you wish to add this node to and optionally provide a name for the node.
  - 5. Verify that the node is online from the **Work with Nodes** panel.

---

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

This task provides step-by-step instructions for recovering from an offline 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, then you need to perform one of the following procedures to recover the VDisks and regain access to them. Depending on the failure type, you may have lost data that was cached for these VDisks, therefore, they have gone offline.

### Steps:

Perform the following steps to recover from an offline VDisk:

### Example:

**Data loss scenario 1** One node in an I/O group failed and failover started on the second node. During this time, the second node in the I/O group fails before the cache has become write through mode. The first node is successfully repaired but its cache data is stale, therefore, it cannot be used. The second node is repaired or replaced and has lost its hardend data, therefore, the node does not think that it is part of the cluster. You need to perform the following:

- 1. Perform the recovering of the node back into the cluster procedure.
- 2. For each VDisk, move the offline VDisks to the recovery I/O group procedure.
- 3. For each VDisk, move the offline VDisks back to their original I/O group procedure.

### Example:

**Data loss scenario 2** Both nodes in the I/O group have failed and have been repaired. The nodes have lost their hardend data, therefore, the nodes do not think that they are part of the cluster. You need to perform the following:

- 1. For each VDisk, move the offline VDisk to the recovery I/O group.
- 2. For both nodes, move the recovered nodes back in to the cluster.
- 3. For each VDisk, move the offline VDisks back to their original I/O group.

**Recover node back into cluster** Perform the following to recover a node back in to the cluster:

- 1. Verify that the node is offline by viewing the **Work with Nodes** panel.

2. Remove the old instance of the offline node from the cluster by selecting the node and selecting the **Delete Node** task.
3. Verify that the node can be seen on the fabric.
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 it will be necessary to follow the SDD procedure to discover the new paths and to check that each vpath is now presenting the correct number of paths. For information about adding paths to existing vpairs. see the *Subsystem Device Driver (SDD) User's Guide*.
  - b. You may also need to modify the configuration of your disk controllers. If your controller uses a mapping technique to present its RAID arrays or partitions to the cluster you will need to modify the port groups that belong to the cluster because the WWNN or WWPN's of the node have changed. See Chapter 25, "Configuring disk controllers", on page 223 for more information.
5. Add the node back into the cluster by selecting the **Add Node** task from the **Work with Nodes** panel. Select the node from the list of candidate nodes and select the I/O group from the list. Optionally enter a node name for this node.
6. Verify that the node is online by refreshing the **Work with Nodes** panel.

**Note:** You may need to close the panel and re-open it for the refresh to take effect.

**Move the offline VDisk to the recovery I/O group** Perform the following steps to move the offline VDisks to the recovery I/O group:

1. List all VDisks that are offline and that belong to the I/O group in question by selecting **Work with VDisks** from the portfolio. In the filter panel, enter the <iogrpname> in the I/O group filter box and select offline as the status.
2. For each VDisk returned, select the VDisk and select the **Modify** task. In the modify panel, only change the I/O group to **Recovery I/O group**. You may be asked to confirm and force the move, select to force the move.

**Move the offline VDisks back to their original I/O group** Perform the following steps to move the offline VDisks back to their original I/O group:

1. For each VDisk, select the VDisk and select the **Modify** task. In the modify panel, only modify the I/O group back to the original <iogrpname>.
2. Verify that the VDisks are now online by closing the Work with VDisks panel and opening it again. This time in the filter panel only enter the <iogrpname> in the I/O group filter box. The VDisks should all be online.

---

## Replacing an HBA in a host using the SAN Volume Controller Console

This task provides step-by-step instructions for replacing an HBA in a host using the SAN Volume Controller Console.

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

### Prerequisites:

Ensure your switch is zoned correctly. See Chapter 25, “Configuring disk controllers”, on page 223 for more information.

#### Steps:

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

1. Locate the host object that corresponds with the host in which you have replaced the HBA. Click **Work with Hosts** from the portfolio. Select the host object and select the **Add Ports** task.
2. Add the new ports to the existing host object. Select the candidate WWPNs from the list and click **Add**. Complete the task by clicking **OK**.
3. Remove the old ports from the host object. Select the host object and select the **Delete Ports** task. Select the WWPNs you wish to remove (the ones that correspond with the old HBA that was replaced). Click **Add** to add them to the list of WWPNs to be deleted. Complete the task by clicking **OK**.
4. Any mappings that exist between the host object and VDisks will automatically be applied to the new WWPNs. Therefore, the host should see the VDisks as the same SCSI LUNs as before. See *Subsystem Device Driver (SDD) User’s Guide* for adding paths to existing vpaths.

---

## Adding a new storage controller to a running configuration

This task provides step-by-step instructions for adding a new storage controller to a running configuration.

#### Prerequisites:

You can add a new storage controller to your SAN at any time. You should follow the switch zoning guidelines provided in Chapter 26, “Overview about zoning a switch”, on page 233 and also ensure the controller is setup correctly for use with the SAN Volume Controller Chapter 25, “Configuring disk controllers”, on page 223.

You should create one or more arrays on the new controller. It is recommend that you use, RAID-5, RAID-1 or RAID-0+1 (sometimes called RAID-10) for maximum redundancy and reliability. Generally 5+P arrays are recommend. If your controller provides array partitioning we recommend that you create a single partition from the entire capacity available in the array, remember the LUN number that you assign to each partition as you will need this later. You should also follow the mapping guidelines (if your storage controller requires LUN mapping) to map the partitions or arrays to the SAN Volume Controller ports. You can determine the SAN Volume Controller ports by following the “Determining a nodes WWPNs using the SAN Volume Controller Console” on page 112.

#### Steps:

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

1. To ensure that the cluster has detected the new storage (MDisks) click **Work with MDisks** and select the **Detect MDisks** task.
2. The controller itself will have automatically been assigned a default name. If you are unsure which controller is presenting the MDisks click **Work with Disk Controllers**. You should see a new controller listed (the one with the highest numbered default name). Remember the controller name and follow the

"Determining a storage controller name from its SAN Volume Controller name" on page 112 procedure to validate that this is the correct controller.

3. Close and re-open the **Work with MDisks** panel using the filter panel to select a mode of **unmanaged** and a controller name that corresponds with the new controller's name. The MDisks shown should correspond with the RAID arrays or partitions you have created. Remember the field controller LUN number, this corresponds with the LUN number you assigned to each of the arrays or partitions.
4. It is recommended that you create a new managed disk group and add only the RAID arrays that belong to the new controller to this MDisk group. You should also avoid mixing RAID types, so for each set of RAID array types (for example, RAID-5, RAID-1) you should create a new MDisk group. Give this MDisk group an appropriate name, so if your controller is called FAST650-fred, and the MDisk group contains RAID-5 arrays, call it something like F600-fred-R5.
5. Click **Work with MDisk group** from the portfolio. Select the **Create MDisk group** task. On the new panel, enter the name you wish to give this group, select the MDisks you wish to add from the list and click **Add**. Select the extent size you wish this group to have and click **OK**.

---

## Removing a storage controller

This task provides step-by-step instructions for removing a storage controller.

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

This function can also be performed by migrating all the VDisks that are using storage in this MDisk group to another MDisk group. This procedure has an advantage if you wish to consolidate the VDisks in a single or new group. However, you can only migrate a single VDisk at a time. The procedure outlined below will migrate all the data through a single command. If you wish to migrate the VDisks however, follow the "Migrating VDisks between MDisk groups" on page 123 procedure for all VDisks that are using this group. You can determine the relationship between VDisks and MDisks by following the "Determining the relationship between VDisks and MDisks using the SAN Volume Controller Console" on page 112 procedures.

This procedure can also be used to remove or replace a single MDisk in a group. If an MDisk has suffered a partial failure, such as a degraded array, and you can still read the data from the disk but cannot write to it, you can follow this procedure to replace just one MDisk. In steps 1 on page 120 and 3 on page 120 only add or remove a single MDisk rather than a list of MDisks.

### Prerequisites:

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

### Steps:

Perform the following steps to remove a storage controller:

1. Perform steps 1 on page 180 through 4 on page 180 from the “Adding a new storage controller to a running configuration using the CLI” on page 180 section.
2. Select the MDisk group that contains the old MDisks you are decommissioning. Select the **Add MDisk** task. On the task dialog, select the new MDisks from the list and click **Add**. Click **OK** to complete the task.
3. You should now have an MDisk group that contains the old MDisks (those to be decommissioned) and the new MDisks (those that are replacing them). Ensure that the capacity of the new MDisks is the same or exceeds that of the old MDisks before proceeding.
4. Force delete the old MDisks from the group. This will migrate all the data from the old MDisks to the new MDisks. Select the **Remove MDisks** task. Select the MDisks you wish to remove and click **Add**. Click **OK** to complete the task. When prompted click **Forced Delete**. Depending upon the number and size of the MDisks, and the number and size of the VDisks that are using these MDisks, this operation will take some time to complete although the task will complete immediately.
5. The only way to check progress is by using the command line interface. Issue the following command:

```
svctools lsmigrate
```

6. When all the migration tasks have completed, for example, the command in step 5 returns no output, you can safely remove the old controller from the SAN.
7. Once you have removed the old controller from the SAN, re-run the detect MDisks task to remove the entries for the old MDisks.

## Expanding a VDisk using the SAN Volume Controller Console

This task provides step-by-step instructions for expanding a VDisk using the SAN Volume Controller Console.

VDisks can be expanded should it be required. However, if the VDisk contains data that is being used, only AIX and Windows 2000 hosts can cope with a VDisk being expanded. A VDisk that is not yet mapped to any hosts and hence does not contain any customer data can be expanded at any time.

This feature can be used in two ways:

1. To increase the capacity available on a particular VDisk that is already mapped to a host. The procedure below outlines using this feature in this way. The following matrix shows the supported platforms and requirements if this feature is to be used:

*Table 11. Supported platforms and requirements*

Platform	Supported	Requirement
AIX	Yes	AIX 5.2 onwards only
HP-UX	No	
Linux	No	
SUN Solaris	No	
Windows NT	No	
Windows 2000	Yes	

2. To increase the size of a VDisk so that it matches the size of the source or master VDisk and can be used in a FlashCopy mapping or Remote Copy relationship. The following procedure does not need to be followed to use this feature in this way. Ensure the target or auxiliary VDisk is not mapped to any host and issue the following command:

```
svctask expandvdisksize
```

**Note:** You can determine the exact size of the source or master VDisk by issuing the following command:

```
svcinfo lsvdisk -bytes <vdiskname>
```

#### Steps:

Perform the following steps to expand a VDisk that is mapped to a AIX host:

1. Determine the VDisk you wish to expand and remember its <vdiskname>.
2. Determine the host that this VDisk is mapped to. If any of the hosts are running HP-UX, Linux, Solaris or Windows NT then you cannot continue.
3. Determine the volume group that the VDisk belongs to (it is assumed you know the VDisk to HDisk relationship as determined in step 1).
4. Quiesce all I/O operations to *all* volumes that belong to the volume group and sync the filesystems mounted on this volume group.
5. Check the current type of the VDisk by viewing the VDisk details in the **Work with VDisks** panel. If the VDisk has a type of *image* it cannot be expanded. If the VDisk has a type of *sequential* it will become striped when you expand it.
6. You must first deactivate the volume group to which this VDisk belongs. Issue the following:

```
<AIX_HOST_PROMPT> varyoffvg <volume_group>
```

7. From the **Work with VDisks** panel, select the VDisk and select the **Expand** task. Enter the capacity by which you wish to extend this VDisk and the select the appropriate units. Select one, more, or all of the MDisks from the list. These will be the MDisks that provide the extra capacity. Optionally, select the format checkbox if you want this extra capacity to be formatted before use.
8. Re-activate the volume group so that the change in size is detected by the HBA device driver. Issue the following:

```
<AIX_HOST_PROMPT> varonvg <volume_group>
```

9. Run the change volume group to notify the LVM that the size has changed. Issue the following:

```
chvg -g <volume_group>
```

10. Expand the filesystem(s) that are mounted on this VDisk (or use the new capacity as required).
11. Restart I/O operations to the volume group.

#### Steps:

Perform the following steps to expand a VDisk that is mapped to a Windows 2000 host:

1. Ensure that you have run Windows Update and have applied all recommended updates to your system prior to attempting to expand a VDisk that is mapped to a Windows 2000 host.

**Note:** VDisks can be expanded under Windows 2000 concurrently with I/O operations.

2. From the **Work with VDisks** panel, select the VDisk and select the **Expand** task. Enter the capacity by which you wish to extend this VDisk and the select the appropriate units. Select one, more, or all of the MDisks from the list. These will be the MDisks that provide the extra capacity. Optionally, select the format checkbox if you want this extra capacity to be formatted before use.

3. On the Windows Host, start the Computer Management application and open the Disk Management window under the Storage branch.

Restart the Computer Management application if it was opened prior to expanding the VDisk and the disk.

4. You will see the disk that you expanded now has some unallocated space at the end of the disk.

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 etc) or add it to an existing volume.

Dynamic disks can be expanded without stopping IOs in most cases. However, in some applications the operating system may report IO errors. When this problem occurs, either of the following entries may be recorded in the System event log:

```
Event Type: Information  
Event Source: dmio  
Event Category: None  
Event ID: 31  
Description:  
dmio: Harddisk0 write error at block ##### due to  
disk removal  
  
Event Type: Information  
Event Source: dmio  
Event Category: None  
Event ID: 34  
Description:  
dmio: Harddisk0 is re-online by PnP
```

**Attention:** This is a known problem with Windows 2000 and is documented at the Microsoft knowledge base as article Q327020. If either of these errors are seen, run Windows Update and apply the recommended fixes to resolve the problem.

---

## Shrinking a VDisk using the SAN Volume Controller Console

This task provides step-by-step instructions for shrinking a VDisk using the SAN Volume Controller Console.

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 will arbitrarily reduce the capacity of the VDisk by removing a partial, one or more extents from those allocated to the VDisk. You cannot control which extents are removed and so you cannot guarantee that it is unused space that is removed.

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

**Steps:**

Perform the following steps to shrink a VDisk:

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

```
svcinfo lsvdisk -bytes <vdiskname>
```

**Note:** It is not possible to determine the exact size using the SAN Volume Controller Console.

3. Shrink the VDisk by the required amount by selecting the VDisk in the **Work with VDisks** panel and selecting the **Shrink** task. Enter the capacity and units by which to shrink the VDisk.

---

## Migrating VDisks between MDisk groups

This task provides step-by-step instructions for migrating VDisks between MDisk groups.

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. This can only be performed using the CLI.
2. Migrating VDisks from one MDisk group to another. This can be used to remove hot MDisk groups, for example, you can reduce the utilization of a group of MDisks.

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

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

1. Add additional MDisks to the target MDisk group. This will provide additional extents in the group and will allow the migrations to be restarted (by marking the error as fixed).

2. Migrate one or more 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).

#### Steps:

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, select **Manage Cluster** from the portfolio, then select the **Start statistics collection** task. Enter 15 minutes for the interval and click **OK**.
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. Select **Service and Maintenance** from the portfolio and then select the **List dumps** task.
3. Click on the **I/O Statistics logs** link in the panel displayed. This will list the I/O statistics files that have been generated. These are prefixed with **m** for MDisk statistics and **v** for VDisk statistics. Click on one of the filenames to view the contents.
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.
5. Stop the statistics collection again by selecting **Manage Cluster** from the portfolio and then select **Stop statistics collection** task.

Once you have analyzed the I/O statistics data, you can determine which 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.

1. After having determined which VDisk you wish to migrate, and the new MDisk group you wish to migrate it to, click **Work with VDisks**. Select the VDisk you wish to migrate and select the **Migrate VDisk** task. In the dialog returned, select the MDisk group you wish to migrate to and select the number of threads to devote to this migrate. The more threads, the quicker the migrate will complete.
2. The only way to check progress is by using the command line interface. Issue the following command:

```
svcinfo lsmigrate
```

---

## Creating image mode virtual disks

This task provides step-by-step instructions for creating image mode virtual disks.

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

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

1. Managed disks that contain existing data cannot be differentiated from managed disks that are blank. Therefore, it is vital that you control the

introduction of these disks to the cluster. It is recommended that you introduce these disks one at a time. For example, map a single LUN from your RAID controller to the cluster and refresh the view of managed disks. The newly detected disk is displayed.

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

Go to the following Web site for more information:

[www.ibm.com/storage/support/2145](http://www.ibm.com/storage/support/2145)

#### Steps:

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

1. Map a single RAID array or LUN from your RAID controller to the cluster. You can do this either through a switch zoning or a RAID controller based on your host mappings.
2. Rescan the list of managed disks from the SAN Volume Controller Console. Click **Work with Managed Disks -> Managed Disks**. You can then filter through the unmanaged mode disks.  
Optionally, if the new managed disk is not listed you may need to run a fabric level discovery. From the SAN Volume Controller Console, select **Work with Managed Disks** and choose **Task discovery**. After a few minutes, refresh the view of managed disks and the new managed disk should be displayed.
3. Convert the managed disk into an image mode virtual disk. In the SAN Volume Controller Console, select the specific managed disk and choose task **Create VDisk in Image Mode**. This will bring up the create image mode virtual disk wizard. You can select the managed disk group to add this managed disk to, and the I/O group that will provide the upstream data path for the virtual disk.
4. Map the new virtual disk to the hosts that were previously using the data that the MDisk contains. In the SAN Volume Controller Console, select **Work with Virtual Disks --> Virtual Disks**. On the Filtering Virtual Disks (VDisks) panel, enter the filter criteria or click **Bypass filter**. On the Viewing Virtual Disks panel, choose **Map a VDisk to a host**, and click **Go**.

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

#### Related topics:

- “Creating image mode virtual disks using the CLI” on page 189

---

## Advanced function Remote Copy overview

This topic provides an overview about the advanced function FlashCopy and Remote Copy overview.

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

[www.ibm.com/redbooks](http://www.ibm.com/redbooks)

**Related topics:**

- Appendix B, "Valid combinations of FlashCopy and Remote Copy functions", on page 263

---

## Advanced function cluster overview

This topic provides an overview about advanced functions for your cluster.

**Overview:**

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

### Deleting nodes from a cluster

You can delete a node from the cluster with the Deleting a Node from Cluster panel.

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

**Prerequisites:**

**Attention:** If you are deleting a single node, and the other node in the I/O group is online, be aware that the cache on the partner node will go into write-through mode and that you are exposed to a single point of failure should the partner node fail. Proceed to 2.

**Attention:** If you are deleting a node, and this is the last node in the I/O group, you will lose access to all VDisks served by this I/O group. Ensure that all VDisks are not being accessed or contain data that you wish to continue to access, or ensure that they have been migrated to a different (online) I/O group.

1. Begin by determining the VDisks that are still assigned to this I/O group:
  - a. Determine the VDisks in question by requesting a filtered view of VDisks where the filter attribute is the I/O group in question.
  - b. Once you have a list of VDisks, determine the hosts that they are mapped to by following the procedure called, Determining the hosts that a VDisk is mapped to.
  - c. Once you have determined the hosts and are sure that you do not wish to maintain access to these VDisks proceed to 2.
  - d. If you determine that some or all of the VDisks assigned to this I/O group do contain data that you wish to continue to access, you should follow the procedure called, Migrating a VDisk to a new I/O group.
2. Before deleting the node, it is essential that for each vpath presented by the VDisks you intend to remove, the SDD configuration is updated to remove the vpaths in question. Failure to do this may result in data corruption. See the *IBM TotalStorage: Subsystem Device Driver User's Guide* for details about how to dynamically reconfigure SDD for the given host operating system.

3. Proceed to 1.

**Steps:**

Perform the following steps to delete a node from the cluster:

1. Click **Work with Nodes** from the portfolio.
2. Click **Nodes** from the portfolio. The Nodes panel is displayed.
3. Select the node you want to delete and Select **Delete a Node** from the list.  
Click **Go**. The Deleting Nodes from Cluster panel is displayed

**Related topics:**

- “Determining the host that a VDisk is mapped to” on page 171
- “Migrating a VDisk to a new I/O group” on page 114

## **Setting up the cluster features using the SAN Volume Controller Console**

This task provides step-by-step instructions about how to set up cluster features using the SAN Volume Controller Console.

**Steps:**

Perform the following steps to set up the cluster feature settings:

1. Click **Service and Maintenance** from the portfolio.

- Click **Feature settings** to check or update the cluster featurization settings. The Featurization settings panel is displayed.
- Because the feature settings are entered when the cluster is first created, you

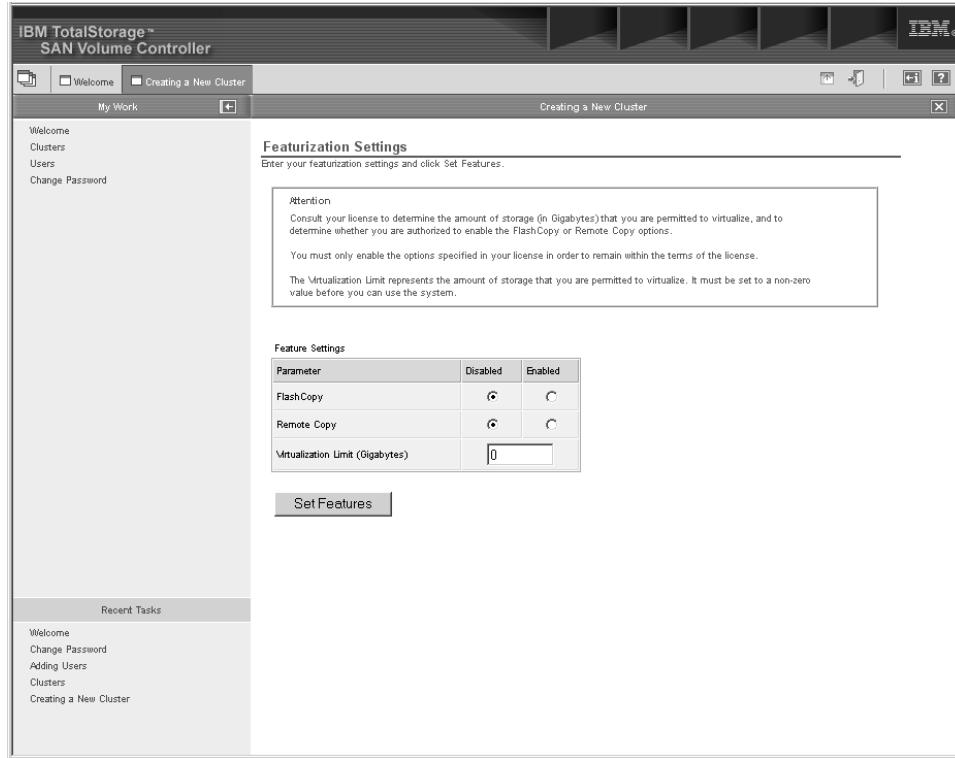


Figure 39. Featurization settings panel

need to update the settings only if you have changed your license. You can change the following values:

- FlashCopy: disabled or enabled
  - Remote copy: disabled or enabled
  - Virtualization limit: number, in gigabytes (1073741824 bytes)
- Click **Update feature settings** once you have selected the new settings. If you enable a function that was previously disabled, or if the virtualization size is increased, a legal notice is displayed. This notice asks you to ensure that you have a valid license for the new settings.

## Enabling the cluster maintenance procedure using the SAN Volume Controller Console

This task provides step-by-step instructions for enabling the cluster maintenance procedure using the SAN Volume Controller Console.

### Steps:

Perform the following steps to enable the maintenance procedure:

- Click **Service and Maintenance** from the portfolio.
- Click **Run Maintenance Procedures** to start the online maintenance procedures. The Maintenance procedures panel is displayed. A pop-up window is also displayed, requesting you to enter the user name and password for the SAN Volume Controller cluster. The Maintenance procedures window enables you to run the maintenance procedure on the cluster.

3. Click **Start Analysis** to analyze the cluster error log. A table of unfixed errors is displayed. The errors are sorted so that the most serious errors (those with the lowest error code) are listed first. The Maintenance panel is displayed. If you click the error code of a particular error log entry, you are guided through a series of actions that help you to estimate the state of the cluster and determine if the error was an isolated event or if a component has failed. 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 **unfixed error** to **fixed error**.

## Maintaining passwords using the SAN Volume Controller Console

This task provides step-by-step instructions about how to maintain passwords using the SAN Volume Controller Console.

### Steps:

Perform the following steps to maintain passwords:

1. Click **Manage Cluster** from the portfolio.
2. Click **Maintain Passwords** to modify the admin or service passwords that control access to the create cluster wizard. The Maintain passwords panel is displayed. The Maintain passwords window enables you to update the passwords that control access to the Web application for admin and service users. Passwords must be typed twice to allow verification. Passwords can consist of A - Z, a - z, 0 - 9, and underscore.
3. Type your admin or service user password and then click **Maintain Passwords** to change the password. If the admin password is changed, a password prompt is displayed and you must re-authenticate the password by entering the new admin password in the password prompt. Make a careful note of the admin password, because without it, you cannot access the cluster through the SAN Volume Controller Console.

## Adding subsequent SSH public keys to the SAN Volume Controller

This task provides step-by-step instructions for adding an SSH public key on to the SAN Volume Controller.

### Steps:

During the cluster creation wizard, you will have added an SSH key to the cluster that allows the master console (where the SAN Volume Controller Console is running) to access the cluster. If you wish to add more SSH keys, that is, grant SSH access to other servers you need to follow the procedure below.

1. After a cluster has been created with the SAN Volume Controller Console creation wizard, you can add further SSH public keys for other systems from which you plan to use the SAN Volume Controller Console using this process.

**Note:** You can also add subsequent SSH keys using the Command Line Interface. See *IBM TotalStorage Virtualization Family SAN Volume Controller Command-Line Interface User's Guide* svctask addsshkey command. Note you can only run this command from a machine that already has SSH access, in this case the master console.

2. Select Service and Maintenance → Maintain SSH Keys in the portfolio frame.

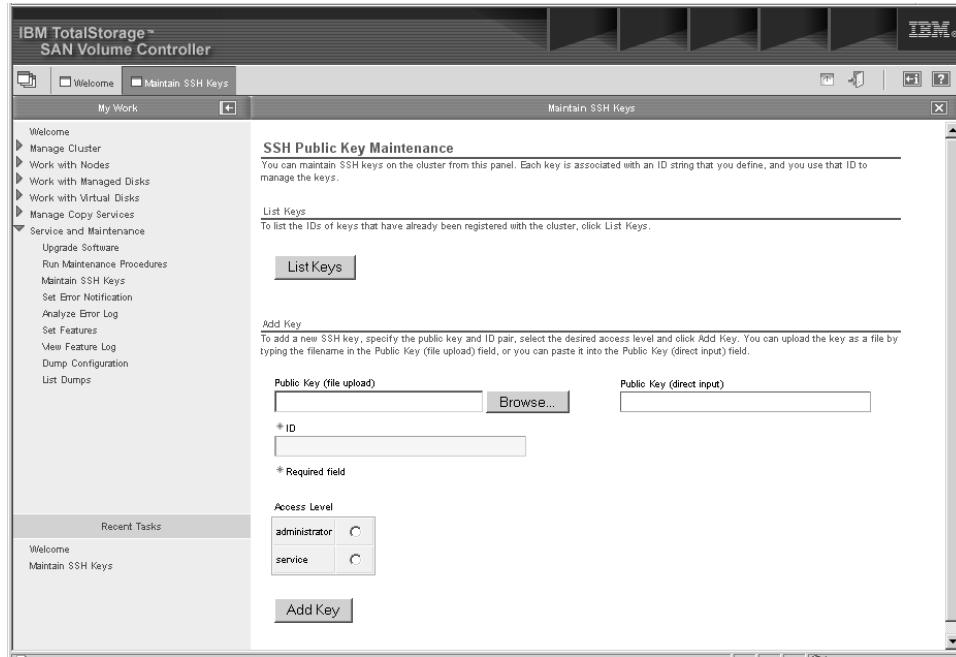


Figure 40. SSH Public Key Maintenance panel

3. Click the **Maintain SSH Keys** option. The window appears to enable you to enter the client SSH public key information to be stored on the cluster. At the SSH key maintenance window, perform the following steps:
  - If you are adding the SSH client key for the master console then click **Browse** and locate the public key you generated earlier. If you are adding a SSH client key for another system then either click **Browse** and locate the public key or cut and paste the public key into the direct input field.
  - Select the **Administrator** radio button.
  - Type a name of your choice in the ID field which uniquely identifies the key to the cluster.
  - Click the **Add Key** button.
  - Click the **Maintain SSH Keys** option.

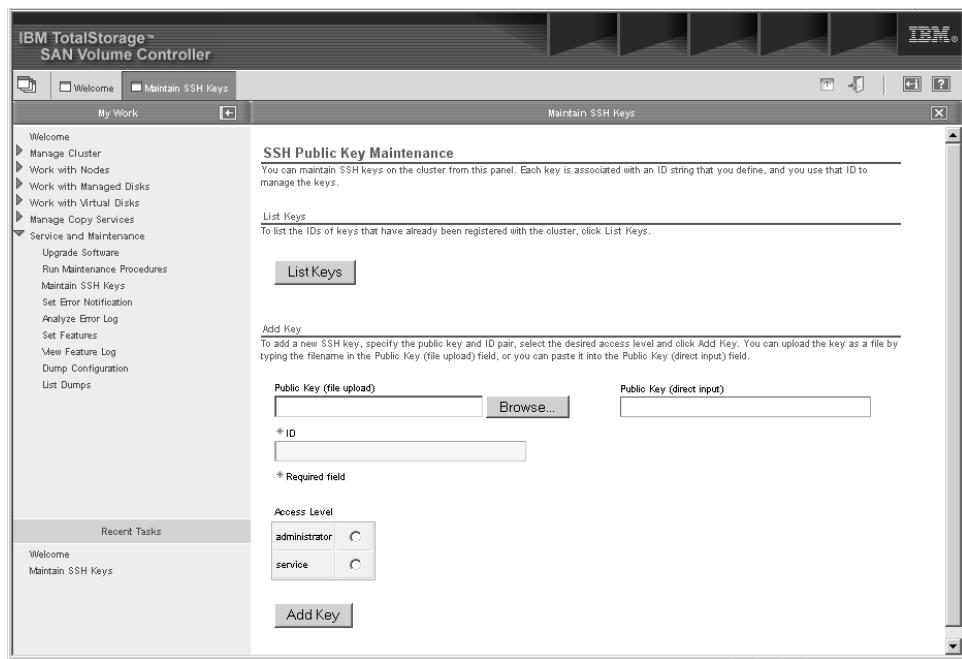


Figure 41. SSH Public Key Maintenance panel

- f. Click the **Show IDs** button to see all key IDs loaded on the SAN Volume Controller.

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 may either be performed using the SAN Volume Controller Console or the Command Line Interface (CLI).

#### Related topics:

- “Secure Shell (SSH) configuration” on page 58

## Setting up error notifications using the SAN Volume Controller Console

This task provides step-by-step instructions about enabling error notifications using the SAN Volume Controller Console.

#### Steps:

Perform the following steps to change the error notification settings:

1. Click **Service and Maintenance** from the portfolio.
2. Click **Error settings** to display the existing error notification settings and to change them. The Modify error notification settings panel is displayed. The Modify Error Notification Settings window enables you to update your error notification settings. You can select whether the cluster raises an SNMP trap, issues an e-mail notification for entries that are added to the cluster error or event log, or both. Three levels of notification are possible:
  - **None** No error or status changes will be sent.
  - **hardware\_only** You will be notified of errors, but you will not be notified of status changes.
  - **All** You will be notified of all errors and status changes.

If you have an SNMP manager installed or if you want to be notified by e-mail of errors or events, you should enable error notification. The notification levels for SNMP and e-mail alerts can be set independently. If you select **All** or **hardware\_only** notification, you must specify a destination for the notification.

3. Click **Modify settings** to update the settings.

## **Resetting a cached SAN Volume Controller cluster SSH host fingerprint on the SAN Volume Controller Console overview**

This topic provides an overview about how to reset a cached SAN Volume Controller cluster SSH host fingerprint on the SAN Volume Controller Console.

### **Overview:**

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. The SSH protocol details that respective credentials are exchanged at the onset of communication between the SSH client and server. The SSH client caches the accepted SSH host server fingerprint and any change to the SSH server fingerprint in future exchanges will result in a challenge to the end user to accept the new fingerprint. When a new code load is performed on the SAN Volume Controller Console, new SSH server keys can be produced which will result in the SSH client flagging the SSH host fingerprint as having changed and therefore being invalid.

The SAN Volume Controller Console externalizes the status of an invalid SAN Volume Controller cluster SSH server key in the Availability Status column of the Cluster dialogue. You can use the actions drop down list box to reset the Invalid SSH Fingerprint status.

### **Related topics:**

- “Resetting a cached SAN Volume Controller cluster SSH host fingerprint on the SAN Volume Controller Console”

## **Resetting a cached SAN Volume Controller cluster SSH host fingerprint on the SAN Volume Controller Console**

This task provides step-by-step instructions to resetting a cached SAN Volume Controller cluster SSH host fingerprint on the SAN Volume Controller Console. A cached SAN Volume Controller cluster SSH fingerprint can become invalid if the cluster microcode is reloaded which results in a new SSH host fingerprint.

### **Steps:**

Perform the following steps to reset an invalid SSH fingerprint:

1. Click **Clusters** from the portfolio.
2. Check the **Select** box for the cluster with a status of **Invalid SSH Fingerprint**.
3. Click **Reset the SSH Fingerprint** in the actions selection list.
4. Click **Go**.
5. Select **OK** when prompted with the message CMMVC3201W.
6. The Viewing Cluster panel will refresh.
7. Check the Availability Status for **OK**.

## **Resetting an refused SSH key relationship between the SAN Volume Controller Console and the SAN Volume Controller cluster overview**

This topic provides an overview about resetting an refused SSH key relationship between the SAN Volume Controller Console and the SAN Volume Controller cluster.

### **Overview:**

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 at key generation time. The SSH client public key is stored on each SAN Volume Controller cluster with which the SAN Volume Controller Console will communicate. 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 will fail.

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 dialogue.

Because the client SSH key pair must be coordinated across two systems, you may have to take one or more actions to reset the pair of keys. Perform one or more of the following steps 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

## **Modifying IP addresses using the SAN Volume Controller Console**

This task provides step-by-step instructions about modifying IP addresses using the SAN Volume Controller Console.

### **Steps:**

Perform the following steps to modify an IP address:

1. Click **Manage Cluster** from the portfolio.
2. Click **IP Addresses** to check or change the IP address settings for the cluster. The Modify IP addresses panel is displayed. The Modify IP Addresses window 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
  - Gateway

Fill in all four fields for any 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 perform the IP address update. When you specify a new cluster IP address, the existing communication with the cluster is broken. You must use the new cluster IP address to reestablish your browser connection. A new SSL certificate is generated by the cluster (to show the new IP address). This new certificate is displayed when the Web browser first connects to the cluster.

## **Listing log or dump files using the SAN Volume Controller Console**

This task provides step-by-step instructions for listing log or dump files using the SAN Volume Controller Console.

### **Steps:**

Perform the following steps to list the log or dump files:

1. Click **Service and Maintenance** from the portfolio.
2. Click **List dumps** to display the log files or dumps that are available on the configuration node of the cluster and on other nodes. 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 (as is usual), the **Check other nodes** button is displayed. If you click this button, the log files and dumps for all nodes that are part of the cluster is displayed. Dumps and logs on all nodes in the cluster can be deleted or copied to the node.

If you click on one of the file types, all the files of that type are listed in a table. Note that for error logs and software dumps, the file names include the node name and time and date as part of the file name.

You can copy the files to a local workstation by right-clicking on the filename and using the **Save target as** (Netscape) or **Save file as** (Internet Explorer) option from the Web browser.

The file types that the **List dumps** option supports are:

- Error logs
- Configuration logs
- I/O statistic logs
- I/O trace logs
- Feature logs
- Software dumps

The software dump files contain dumps of the SAN Volume Controller memory. Your service representative might ask for these dumps to debug problems. The software dumps are large files (approximately 300 MB). Consider copying these files to your host using secure copy methods.

## **Changing the language settings**

This task provides step-by-step instructions about how to change the language settings.

### **Steps:**

Perform the following steps to change the language settings:

1. Click **View Clusters** and select a cluster that you want to change the language setting for.
2. Click **Launch SAN Volume Controller application**.
3. Click **Manage Cluster**.
4. Click **General Properties**. From this panel you can change the locale setting to the appropriate language.

## **Viewing the feature log using the SAN Volume Controller Console**

This task provides step-by-step instructions about viewing the feature log using the SAN Volume Controller Console.

### **Steps:**

Perform the following steps to view the feature log:

1. Click **Service and Maintenance** from the portfolio.
2. Click **Feature Log** to view a log that contains events and errors that are related to the featurization options. The Feature log panel is displayed. The Feature log panel enables you to view the feature log. 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.

## **Analyzing the error log using the SAN Volume Controller Console**

This task provides step-by-step instructions for analyzing the error log using the SAN Volume Controller Console.

### **Steps:**

Perform the following steps to analyze the error log file:

1. Click **Service and Maintenance** from the portfolio.
2. Click **Error Log** to access and display the cluster error log. The Error log analysis panel is displayed. The Error log analysis panel enables you to analyze the cluster error log. You can display the whole log or filter the log so that only errors, events, or unfixed errors are displayed. In addition, you can request the table to be sorted either by error priority or by time. For error priority, the most serious errors are the lowest-numbered errors. They are, therefore, displayed first in the table. For time, either the older or the latest entry can be displayed first in the table. You can also select how many error log entries are to be displayed on each page of the table. The default is set to 10.
3. After selecting the options, click **Process** to display the filtered error log in the table. The Analyze error log continued panel is displayed. Forward and Backward scroll buttons are displayed, depending on the existing page number and the total number of pages that are in the table. If the table contains more than two pages of entries, a **Go to** input area is displayed in the table footer. This input area enables you to skip to a particular page number.  
If you click on the sequence number of a particular table record, more information about that error log entry is displayed. If the record is an error (instead of an event), you can change fixed or unfixed status of the record; that is, you can mark an unfixed error as fixed or a fixed error as unfixed.
4. Click **Clear log** to erase the whole cluster error log.

**Note:** If you click **Clear log**, this will *not* fix the existing errors.

## Shutting down a cluster

This topic provides information about shutting down a cluster.

### Prerequisites:

If all input power to a SAN Volume Controller cluster is to be removed for more than a few minutes, (for example, if the machine room power is to be shutdown for maintenance), it is important that the cluster is shutdown before the power is removed. The reason for this is that if the input power is removed from the uninterruptible power supply units without first shutting down the cluster and the uninterruptible power supplies, the uninterruptible power supply units will remain operational and eventually become drained of power.

When input power is restored to the uninterruptible power supplies they will start to recharge but the SAN Volume Controllers will not permit any I/O activity to be performed to the virtual disks until the uninterruptible power supply 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 uninterruptible power supply units will prevent the battery power being drained and will make it possible for I/O activity to be resumed as soon as input power is restored.

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

**Attention:** If you are shutting down the entire cluster, you will lose access to all VDisks being provided by this cluster.

### Shutting down the cluster:

#### Steps:

Perform the following steps to shut down the cluster:

1. Begin the process of quiescing all I/O to the cluster by stopping the applications on your hosts that are using the VDisks provided by the cluster.
  - a. If you are unsure which hosts are using the VDisks provided by the cluster, follow the procedure called, Determining the hosts that a VDisk is mapped to.
  - b. Repeat the previous step for all VDisks.
2. When all I/O has been stopped, issue the Shut Down cluster task from the Manage Cluster tree in the portfolio.
3. Close the SAN Volume Controller Console window.

### Shutting down a single node:

**Attention:** If you are shutting down a single node, and the other node in the I/O group is online, be aware that the cache on the partner node will go into write-through mode and that you are exposed to a single point of failure should the partner node fail while this node is shut down. Proceed to 2 on page 137.

**Attention:** If you are shutting down a single node, and this is the last node in the I/O group, you will lose access to all VDisks being served by this I/O group.

**Steps:**

Perform the following steps to shut down a single node:

1. Begin the process of quiescing all I/O to the VDisks being served by this nodes I/O group.
  - a. Determine the VDisks in question by requesting a filtered view of VDisks where the filter attribute is the I/O group in question.
  - b. Once you have a list of VDisks, determine the hosts that these are mapped to by following the procedure called, Determining the hosts that a VDisk is mapped to.
2. When all I/O has been stopped issue the Shut down a node task from the Nodes panel by selecting the node you wish to shutdown and selecting the Shut down a node task. Click **Go**.

**Related topics:**

- “Determining the host that a VDisk is mapped to” on page 171
-



---

## **Part 4. Command-Line Interface**

This part provides detailed information about using the command-line interface. More specifically, it provides information about the following:

- Chapter 13, “Getting started with the Command-Line Interface”, on page 141
- Chapter 14, “Scenario: typical usage for the command-line interface”, on page 149
- Chapter 15, “Advanced functions overview for the CLI”, on page 169
- Part 5, “Software upgrade strategy using the CLI and the SAN Volume Controller Console”, on page 199



---

# Chapter 13. Getting started with the Command-Line Interface

This topic provides information about getting started with the Command-Line Interface (CLI).

## Overview:

The SAN Volume Controller cluster Command-Line Interface (CLI) is a collection commands which enable you to manage the SAN Volume Controller. The vehicle for these commands is 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 using the CLI, you must have performed the following initial steps to create and configure a cluster:

- Create a cluster from the front panel
- Complete the creation of a cluster using the SAN Volume Controller Console
- Perform the initial configuration of the cluster using the SAN Volume Controller

In order to use the CLI from a client system you must:

- Install and setup SSH client software on each system which you are going to issue command lines from.
- Generate an SSH key pair on each SSH client.
- Store the SSH public key for each SSH client on to the SAN Volume Controller using the SAN Volume Controller Console.

**Note:** After the first SSH public key has been stored further SSH public keys may be added using either the SAN Volume Controller Console or the CLI.

The functions that can be performed with the Command-Line Interface (CLI) is:

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

## Related topics:

- Chapter 5, “Create cluster from the front panel”, on page 51
- “Creating a cluster on the SAN Volume Controller Console” on page 81
- “Creating a cluster using the SAN Volume Controller Console” on page 82
- “Adding subsequent SSH public keys to the SAN Volume Controller” on page 129

---

## Preparing the SSH client system overview

This topic provides an overview about how to prepare the SSH client system to enable you to issue CLI commands from the host to the cluster.

### **Windows operating systems::**

The master console is a Windows 2000 system which is equipped with the PuTTY Secure Shell (SSH) client software. You can install the PuTTY SSH client software on another Windows host using the PuTTY Installation program putty-0.53b-installer.exe which is in the SSHClient\PuTTY directory of the SAN Volume Controller Console CD-ROM. Or, you can download PuTTY from the following Web site:

<http://www.chiark.greenend.org.uk/~sgtatham/putty/>

The following Web site offers SSH client alternatives for Windows:

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

Cygwin software has an option to install an OpenSSH client. You can download cygwin from the following Web site:

<http://www.cygwin.com/>

### **AIX operating systems::**

For AIX 5L Power 5.1 and 5.2, you can get OpenSSH from the Bonus Packs and you will also need its prerequisite, OpenSSL, from the AIX toolbox for Linux applications for Power Systems. For AIX 4.3.3, you can get the software from the AIX toolbox for Linux applications.

You can also get the AIX installation images from IBM developer Works at the following Web site:

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

### **Linux operating systems::**

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

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

OpenSSH is able to run on a wide variety of additional operating systems. For more information visit the following Web site:

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

---

## Preparing the SSH client system to issue CLI commands

This task provides step-by-step instructions about how to prepare the SSH client system to issue CLI commands.

In order to issue CLI commands to the cluster from a host, you must prepare the Secure Shell (SSH) client on the host so that the host will be accepted by the SSH server on the cluster, and allowed to connect.

The detailed instructions, outlined in “Secure Shell (SSH) configuration” on page 58, cover the PuTTY SSH client, as this is preinstalled on the master console, and is also suitable for other Windows hosts. If you wish to use a host which requires a different type of SSH client, for example OpenSSH, you should follow the instructions for that software.

**Steps:**

Perform the following steps to issue CLI commands from the master console:

For the master console and Windows hosts:

1. Generate a SSH key pair using the PuTTY key generator.
2. Store the SSH clients public key on the cluster (using a browser pointing to the SAN Volume Controller Console).
3. Configure the PuTTY session for the command-line interface

For other types of hosts:

1. Follow the instructions specific to the SSH client to generate an SSH key pair.
2. Store the SSH clients public key on the cluster (using a browser pointing to the SAN Volume Controller Console or the Command Line Interface from an already established host).
3. Follow the instructions specific to the SSH client to establish an SSH connection to the SAN Volume Controller cluster.

**Related topics:**

- “Secure Shell (SSH) configuration” on page 58
- “Generating an SSH key pair using the SSH client called PuTTY” on page 60
- “Configuring the PuTTY session for the command-line interface” on page 64
- “Storing keys in the SAN Volume Controller Console software” on page 61
- “Adding SSH keys for hosts other than the master console” on page 62

---

## Issuing CLI commands from a PuTTY SSH Client system

This task provides step-by-step instructions to issue CLI commands.

**Steps:**

Perform the following steps to issue CLI commands:

1. Open a Command Prompt to open the SSH connection to issue the CLI commands.
2. Make the PuTTY executables available by performing the following:
  - a. Change directory into the PuTTY executables directory. For example, on the master console type the following:

```
C:\Support Utils\putty
```

On another host, which has installed PuTTY in the default location type the following:

C:\Program Files\Putty

- b. Set the path environment variable to include the PuTTY executables directory. For example, type the following:

Set path=c:\Support Utils\putty;%path%

3. Use the PuTTY plink utility to connect to the SSH server on the cluster.

**Related topics:**

- “Running the PuTTY and plink utilities”

---

## Running the PuTTY and plink utilities

This topic provides step-by-step instructions to run the PuTTY plink utility.

All CLI commands are run in an SSH session. You can run the commands in either of the two following modes:

- an interactive prompt mode
- in a single shot mode, which is entered one time to include all parameters.

**Interactive mode:**

For interactive mode, you use the PuTTY executable to open the SSH restricted shell. Type the following:

```
C:\support utils\putty>putty admin@<svcconsoleip>
```

See “Configuring the PuTTY session for the command-line interface” on page 64 for how to setup the PuTTY interactive session.

If you were to issue the **svcinfo lsshkeys** command, which lists the SSH client public keys that are stored on the SAN Volume Controller cluster, the following output is displayed:

```
IBM_2145:admin>svcinfo lsshkeys -user all -delim :  
id:userid:key identifier  
1:admin:smith  
2:admin:jones
```

Type **exit** and press **Enter** to escape the interactive mode command.

The SSH protocol specifies that the first access to a new host server will result in a *challenge* to the SSH user to accept the SSH server public key. Because this is the first time that you connect to an SSH server, the server is not included in the SSH client list of known hosts. Therefore, there is a fingerprint challenge, which asks you do you accept the responsibility of connecting with this host. If you type **y**, the host fingerprint and IP address is saved by the SSH client. For PuTTY, you answer by typing **y** to accept this host fingerprint. This information is stored in the registry for the user name which is logged onto Windows.

The following is an example of the host fingerprint challenge when running in interactive mode:

```
C:\Program Files\IBM\svcconsole\cimom>plink admin@9.43.225.208
```

The server's host key is not cached in the registry. You have no guarantee that the server is the computer you

think it is.  
The server's key fingerprint is:  
ssh-rsa 1024 e4:c9:51:50:61:63:e9:cd:73:2a:60:6b:f0:be:25:bf  
If you trust this host, enter "y" to add the key to  
PuTTY's cache and carry on connecting.  
If you want to carry on connecting just once, without  
adding the key to the cache, enter "n".  
If you do not trust this host, press Return to abandon the  
connection.  
Store key in cache? (y/n) y  
Using username "admin".  
Authenticating with public key "imported.openssh-key"  
IBM\_2145:admin>

### Single line command:

For single line command mode, you can type the following all on one command line:

```
C:\Program Files\IBM\svcconsole\cimom>
plink admin@9.43.225.208 svcinfo lsshkeys
-user all -delim :
Authenticating with public key "imported.openssh-key"
/bin/ls: id:userid:key identifier
1:admin:smith
2:admin:jones
```

```
C:\Program Files\IBM\svcconsole\cimom>
```

The SSH protocol specifies that the first access to a new host server will result in a *challenge* to the SSH user to accept the SSH server public key. Because this is the first time that you connect to an SSH server, the server is not included in the SSH client list of known hosts. Therefore, there is a fingerprint challenge, which asks you do you accept the responsibility of connecting with this host. If you type *y*, the host fingerprint and IP address is saved by the SSH client. For PuTTY, you answer by typing *y* to accept this host fingerprint. This information is stored in the registry for the user name which is logged onto Windows.

The following is an example of the host fingerprint challenge when running in single line command mode:

```
C:\Program Files\IBM\svcconsole\cimom>
plink admin@9.43.225.208 svcinfo lsshkeys
-user all -delim :
The server's host key is not cached in the registry. You
have no guarantee that the server is the computer you
think it is.
The server's key fingerprint is:
ssh-rsa 1024 e4:c9:51:50:61:63:e9:cd:73:2a:60:6b:f0:be:25:bf
If you trust this host, enter "y" to add the key to
PuTTY's cache and carry on connecting.
If you want to carry on connecting just once, without
adding the key to the cache, enter "n".
If you do not trust this host, press Return to abandon the
connection.
Store key in cache? (y/n) y
```

```
Authenticating with public key "imported-openssh-key"
/bin/ls: /proc/20282/exe: Permission denied
dircolors: `/etc/DIR_COLORS': Permission denied
id:userid:key identifier
1:admin:smith
2:admin:jones
```

```
C:\Program Files\IBM\svcconsole\cimom>
```

**Note:** If you are submitting a CLI command with all parameters in single line command mode, you will be challenged upon first appearance of the SSH server host fingerprint. Be careful to ensure that the SSH server host fingerprint is accepted before submitting a batch script file.

The SSH protocol also specifies that once the SSH server public key is accepted, another challenge will be presented if the fingerprint of an SSH server changes from the one previously accepted. In this case, you will need to decide whether to accept this changed host fingerprint. For PuTTY, you answer by typing *y* to accept this host fingerprint. PuTTY stores this information in the registry for the user name which is logged onto Windows.

**Note:** The SSH server keys on the SAN Volume Controller will be regenerated when a microcode load is performed on the cluster. Due to this behavior, you will see a *challenge* presented because the fingerprint of the SSH server has changed.

---

## Configuring the cluster using the CLI

This task provides step-by-step instructions for configuring the cluster using the command-line interface (CLI). The initial steps in creating and configuring a cluster must be performed using the front panel and the SAN Volume Controller Console. Once the cluster has been created and a SSH public key has been added all further tasks can be accomplished using the command-line interface (CLI).

### Steps:

Perform the following steps to configure the cluster:

1. Open a command prompt window.
2. To change your time-zones and set your cluster time, you can issue the **svctask settimezone** and **svctask setcluster** time commands.
3. If you wish to use the Command Line Interface (CLI) from additional systems then use the **svctask addsshkey** to add further SSH public keys.
4. If you choose you can review and modify the initial configuration of the cluster that was performed using the front panel and SAN Volume Controller Console:
  - a. Issue the command **svctask lscluster** to display the cluster properties.
  - b. To modify the passwords, fabric speed or cluster IP address issue the command **svctask chcluster**.
  - c. Issue the command **svctask setpwdreset** to view and change the status of the password reset feature for the front panel.
  - d. To review and modify your featurization settings you can issue the **svctask lslicense** and **svctask chlicense** commands.
  - e. If you wish to modify the set up for error notifications to help manage your errors from the cluster, you can issue the **svctask seevent** command to set up SNMP traps.

#### **Related topics:**

- “Setting the cluster time using the CLI”
- “Maintaining SSH keys using the CLI” on page 192
- “Displaying cluster properties using the CLI” on page 148
- “Modifying passwords using the CLI” on page 148
- “Modifying IP addresses using the CLI” on page 193
- “Maintaining passwords using the CLI” on page 192
- “Reviewing and setting the cluster features using the CLI”
- “Setting up error notifications using the CLI” on page 193

---

## **Setting the cluster time using the CLI**

This task provides step-by-step instructions for setting the cluster time using the command-line interface.

#### **Steps:**

Perform the following steps to set the cluster time:

1. Open a command prompt window.
2. Issue the **svcinfo showtimezone** command to display the current time-zone settings for the cluster. The cluster ID and its associated time-zone are displayed.
3. Issue the **svcinfo lstimezones** command to list the time-zones available on the cluster. A list of valid time-zones settings are displayed in a list. The specific cluster ID and its assigned time-zone are indicated in the list.
4. Issue the **svctask settetimezone** command to set the time zone for the cluster.
5. Issue the **svctask setclustertime** command to set the time for the cluster.

---

## **Reviewing and setting the cluster features using the CLI**

This task provides step-by-step instructions for setting up the cluster features using the command-line interface (CLI).

#### **Steps:**

Perform the following steps to set up the cluster features:

1. Open a command prompt window.
2. Issue the **svctask lslicense** command to return the current license (featurization) settings for the cluster.
3. Issue the **svctask chlicense** command to change the licensed settings of the cluster. Because the feature settings are entered when the cluster is first created, you need to update the settings only if you have changed your license. You can change the following values:
  - FlashCopy: disabled or enabled
  - Remote Copy: disabled or enabled
  - Virtualization limit: number, in gigabytes (1073741824 bytes)

The output displayed lists the feature functions in a list and displays whether they are enabled or disabled.

---

## Displaying cluster properties using the CLI

This task provides step-by-step instructions for displaying cluster properties using the command-line interface (CLI).

### Steps:

Perform the following steps to display cluster properties:

1. Open a command prompt window.
2. Issue the **svcinfo lscluster** command to display a concise view of the cluster.

```
svcinfo lscluster -delim : 10030a007e5
```

where *10030a007e5* is the name of the cluster. The output from this command will display the following: The output from this command will include the following for each cluster on the fabric:

- cluster ID
- cluster name
- cluster IP address
- cluster service mode IP address

---

## Modifying passwords using the CLI

This task provides step-by-step instructions for modifying the admin and service passwords using the command-line interface (CLI). Note that the passwords only affect access to the cluster via the SAN Volume Controller Console. To restrict access to the command line interface (CLI) you must control the list of SSH client keys installed on the cluster.

### Steps:

Perform the following steps to modify the passwords:

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

```
svtask chcluster -admpwd <admin_password>
```

to change the administrator users password.

3. Issue the following command:

```
svtask chcluster -servicepwd <service_password>
```

to change the service users password.

**Note:** If you do not wish the password to be displayed as you enter the command line then you can omit the new password. The command line tool will then prompt you to enter and confirm the password without the password being displayed.

### Related topics:

- “Maintaining passwords using the CLI” on page 192
- “Maintaining SSH keys using the CLI” on page 192

---

## Chapter 14. Scenario: typical usage for the command-line interface

This topic provides a hypothetical example of configuring your SAN Volume Controller using the command-line interface (CLI). The main focus of the following example is to provide storage to your host system.

Our hypothetical example is the following:

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

The following steps illustrates how this can be done:

1. Create a cluster.
2. Configure the cluster with an IP address of 9.20.123.456, a fabric speed of 2 GB. Name the cluster examplecluster.
3. Add nodes
  - knode and lnode to the I/O group called io\_grp0 in the examplecluster cluster
  - mnode and nnode to the I/O group called io\_grp1 in the examplecluster cluster
4. Create the MDisk groups maindiskgroup and bkpdiskgroup
5. Create 4 VDisks
  - 2 VDisks from maindiskgroup
  - 2 VDisks from bkpdiskgroup
6. Create 2 host objects
  - a host object called demohost1 with HBAs that have WWPNs of 10000000C92AD7E5, and 10000000C92F5123
  - a host object called demohost2 with HBAs that have WWPNs of 210000E08B0525D4, and 210100E08B2525D4
7. Create the VDisk-to-host mappings
  - Map the two VDisks from maindiskgroup to demohost1
  - Map the two VDisks from bkpdiskgroup to demohost2
8. Create FlashCopy mappings
  - Create a FlashCopy mapping called main1copy that has a background copy rate of 75
  - Create a FlashCopy mapping called main2copy that has a background copy rate of 50

9. Create a FlashCopy consistency group called maintobkpfcopy and add the 2 FlashCopy mappings to it
10. Prepare and trigger (start) the FlashCopy Consistency Group that contains these mappings

**Note:** Once this step is complete, you have created and allocated storage to your host systems. You have made two VDisks available to demohost1 and then used FlashCopy to make backup copies on two VDisks which are accessible to demohost2.

**Related topics:**

- “Create managed disk (MDisk) groups using the CLI” on page 156
- “Create virtual disks (VDisks)” on page 159
- “Creating host objects using the CLI” on page 162
- “Create VDisk-to-host mappings using the CLI” on page 163
- “Create a FlashCopy consistency group and add mappings using the CLI” on page 165
- “Create FlashCopy mappings using the CLI” on page 164
- “Prepare and trigger a FlashCopy Consistency Group using the CLI” on page 166

---

## Adding nodes to the cluster using the CLI

This task provides step-by step instructions you will need to perform to add nodes to the cluster

**Prerequisites:**

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

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

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

**Special procedures when adding a node to a cluster:**

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

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

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

```
svctools lsnode
```

or, if this information is not available, then

- Before the node is added back into the cluster all the hosts using the cluster must be shut down.

The node must then be added before the hosts are rebooted or, if the I/O group information is not available and it is inconvenient to shutdown and reboot all of the hosts using the cluster, then

- On all the hosts connected to the cluster, unconfigure the Fibre Channel adapter device driver, the disk device driver, and the SDD device driver, before you add the node to the cluster.

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

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

#### Hypothetical scenarios where the special procedures may apply.:

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

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

#### Background:

Applications on host systems direct I/O operations to filesystems or logical volumes which are mapped by the operating system to vpaths which are pseudo disk objects supported by the SDD driver. See the *IBM Subsystem Device Driver (SDD) User's Guide*, SC26-7540.

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

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

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

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

#### Data Corruption Scenario:

Consider a four-node SAN Volume Controller configuration.

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

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

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

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

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

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

#### Context:

Assume that the cluster has been created, that initial configuration has been performed using the SAN Volume Controller Console and that the necessary setup has been performed to use the Command Line Interface (CLI).

The following examples are all based on our hypothetical scenario of setting up a four node cluster. The first node has already been used to create a cluster and therefore there are a further three nodes to add to the cluster.

#### Steps:

Perform the following steps to add nodes to the cluster:

1. Open a command prompt window.
2. Type the **svcinfo lsnode** command to list the nodes that are currently part of the cluster.

#### Example:

```
svcinfo lsnode -delim :  
id:name:UPS_serial_number:WWNN:status:I0_group_id:  
I0_group_name:config_node:UPS_unique_id  
1:node1:10L3ASH:500507680100002C:online:0:io_grp0:yes:202378101C0D18D8
```

The cluster has only just been created so there is only one node in the cluster.

3. Type the **svcinfo lsnodecandidate** command to list nodes that are not assigned to a cluster.

#### Example:

```
svcinfo lsnodecandidate -delim :  
id:panel_name:UPS_serial_number:UPS_unique_id  
5005076801000001:000341:10L3ASH:202378101C0D18D8  
5005076801000009:000237:10L3ANF:202378101C0D1796  
50050768010000F4:001245:10L3ANF:202378101C0D1796
```

There are a total of four nodes, one of which has been used to create a cluster. Therefore there are three candidate nodes which can be added to the cluster.

- Type the **svctask addnode** command to add a node to the cluster. Use the output from the previous commands to choose which I/O group to add the node to and to make sure that when adding a second node to a I/O group that it is attached to a different UPS.

**Notes:**

- a. When adding a node to a cluster you can specify a name for the node. You can also change the name of nodes that are already part of a cluster using the svctask chnode command.
- b. When adding a node to a cluster the node can be identified either using the front panel name which is also printed on a label on the front of the SAN Volume Controller or by using the world wide node name of that node. This example shows both ways.

**Example:**

Add a second node to the first I/O group. Note from the output from step 1 that the node that is already in I/O group 0 is attached to a UPS with the serial number 10L3ASH. Each node in an I/O group must be attached to a different UPS and therefore only the nodes with front panel IDs 000237 and 001245 are suitable candidates.

```
svctask addnode -panelname 000237 -iogrp io_grp0 -name group1node2
```

This command will add the node, identified by the front panel name 000237 to the cluster. The node will be added to I/O group, io\_grp0, and called group1node2.

Next add two nodes to the second I/O group. Check the output from step 3 to make sure that each node is attached to a different UPS.

```
svctask addnode -wwnodename 5005076801000001 -iogrp io_grp1 -name group2node1
svctask addnode -wwnodename 50050768010000F4 -iogrp io_grp1 -name group2node2
```

These commands will add the nodes, identified by the WWNN 5005076801000001 and the WWNN 50050768010000F4 to the cluster. The nodes will be added to I/O group, io\_grp1 and called group2node1 and group2node2.

Finally change the name of the first node from the default name node1 so that it conforms with your naming convention.

```
svctask chnode -name group1node1 node1
```

- Verify the final configuration using the **svcinfo lsnode** command.

**Example:**

In our hypothetical scenario, the command to list the nodes is:

```
svcinfo lsnode -delim :
```

```
id:name:UPS_serial_number:WWNN:status:I0_group_id:
I0_group_name:config_node:UPS_unique_id
1:group1node1:10L3ASH:500507680100002C:online:0:io_grp0:yes:202378101C0D18D8
2:group1node2:10L3ANF:5005076801000009:online:0:io_grp0:no:202378101C0D1796
3:group2node1:10L3ASH:5005076801000001:online:1:io_grp1:no:202378101C0D18D8
4:group2node2:10L3ANF:50050768010000F4:online:1:io_grp1:no:202378101C0D1796
```

**Note:** If this command is issued quickly after adding nodes to the cluster the status of the nodes may be adding rather than online indicating that the process of adding the nodes to the cluster is still in progress. You do not however have to wait for all the nodes to become online before continuing with the configuration process.

### **Result:**

You have now added four nodes to one cluster. The nodes are split into two I/O groups.

### **Related topics:**

- Chapter 14, “Scenario: typical usage for the command-line interface”, on page 149

---

## **Displaying node properties using the CLI**

This task provides step-by-step instructions for displaying node properties using the command-line interface (CLI).

### **Steps:**

Perform the following steps to display the node properties:

1. Open a command prompt window.
2. Issue the **svcinfo lsnode** command to display a concise list of nodes in the cluster.

### **Example:**

Type the following command:

```
svcinfo lsnode -delim :
```

This command displays the following:

```
id:name:UPS_serial_number:WWNN:status:I0_group_id:  
I0_group_name:config_node:UPS_unique_id  
1:group1node1:10L3ASH:500507680100002C:online:0:io_grp0:yes:202378101C0D18D8  
2:group1node2:10L3ANF:5005076801000009:online:0:io_grp0:no:202378101C0D1796  
3:group2node1:10L3ASH:5005076801000001:online:1:io_grp1:no:202378101C0D18D8  
4:group2node2:10L3ANF:50050768010000F4:online:1:io_grp1:no:202378101C0D1796
```

3. Issue the **svcinfo lsnode** command again, however, this time, specify the node ID or name of a node to receive the detailed output.

### **Example:**

For example, to provide a detailed view of the node named group1node1 type the following:

```
svcinfo lsnode -delim : group1node1
```

This command displays the following:

```
id:1  
name:group1node1  
UPS_serial_number:10L3ASH  
WWNN:500507680100002C  
status:online  
I0_group_id:0  
I0_group_name:io_grp0  
partner_node_id:2  
partner_node_name:group1node2  
config_node:yes  
UPS_unique_id:202378101C0D18D8  
port_id:500507680110002C  
port_status:active
```

```
port_id:500507680120002C
port_status:active
port_id:500507680130002C
port_status:active
port_id:500507680140003C
port_status:active
```

The output includes:

- node ID
- node name
- WWNN
- details about the uninterruptible power supply that the node is attached to
- details about the I/O group which the node is a member of
- detailed fibre channel port status information.

---

## Discovering MDisks using the CLI

This task provides step-by-step instructions for discovering MDisks using the command-line interface (CLI).

### Context:

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 will automatically discover the back-end controller and will integrate the controller to determine what storage it is presented to the SAN Volume Controller. The SCSI LUs presented by the back-end controller will be displayed as unmanaged MDisks. If however the configuration of the back-end controller is modified after this has occurred then the SAN Volume Controller may be unaware of these configuration changes. This task allows a user to request the SAN Volume Controller to re-scan the fibre-channel SAN to update the list of unmanaged Mdisks.

**Note:** The automatic discovery performed by SAN Volume Controller does not write anything to a unmanaged MDisk. It is only when a user instructs the SAN Volume Controller to add a MDisk to a managed disk group or use a Mdisk to create an image mode virtual disk that the storage will actually be used.

### Steps:

Perform the following steps to display MDisks:

1. Open a command prompt window.
2. Check to see which MDisks are available by issuing the **svctask detectmdisk** command to manually scan the fibre-channel network for any MDisks.
3. Issue the **svcinfo lsmdiskcandidate** command to show the unmanaged MDisks. These MDisks have not been assigned to an MDisk group. Alternatively, you can issue the **svcinfo lsmdisk** command to view all of the MDisks.

### Example:

In our hypothetical scenario we have a single back-end controller that is presenting eight SCSI LUs to the SAN Volume Controller. Issue the following command:

```
svctask detectmdisk  
svcinfo lsmdiskcandidate
```

This command displays the following:

```
id  
0  
1  
2  
3  
4  
5  
6  
7
```

Issue the following command:

```
svcinfo lsmdisk -delim : -filtervalue mode=unmanaged
```

This command displays the following:

```
id:name:status:mode:mdisk_grp_id:mdisk_grp_name:  
capacity:ctrl_LUN_#:controller_name  
0:mdisk0:online:unmanaged:::273.3GB:0000000000000000:controller0  
1:mdisk1:online:unmanaged:::273.3GB:0000000000000001:controller0  
2:mdisk2:online:unmanaged:::273.3GB:0000000000000002:controller0  
3:mdisk3:online:unmanaged:::273.3GB:0000000000000003:controller0  
4:mdisk4:online:unmanaged:::136.7GB:0000000000000004:controller0  
5:mdisk5:online:unmanaged:::136.7GB:0000000000000005:controller0  
6:mdisk6:online:unmanaged:::136.7GB:0000000000000006:controller0  
7:mdisk7:online:unmanaged:::136.7GB:0000000000000007:controller0
```

#### **Result:**

You have now shown that the back-end controllers and switches have been setup correctly and that the SAN Volume Controller can see the storage being presented by the back-end controller.

#### **Related topics:**

- Chapter 14, “Scenario: typical usage for the command-line interface”, on page 149

---

## **Create managed disk (MDisk) groups using the CLI**

This task provides step-by-step instructions about how to create a MDisk group.

**Attention:** If you add a MDisk to a MDisk group as a managed disk, any data on the MDisk will be lost. If you want to keep the data on a MDisk (for example because you want to import storage that was previously not managed by a SAN Volume Controller) then you should create image mode VDisks instead. See “Creating image mode virtual disks using the CLI” on page 189 for details on this task.

#### Context:

Assume that the cluster has been setup and that a back-end controller has been configured to present some new storage to the SAN Volume Controller.

#### Prerequisites:

Before creating managed disk groups consider how you are going to use your storage. The SAN Volume Controller allows you to create up to 128 managed disks groups and to add up to 128 Mdisks to a Mdisk group. Consider the following factors when deciding how many managed disk groups to create:

- A virtual disk can only be created using the storage from one managed disk group. Therefore if you create small managed disk groups then you may lose the benefits provided by virtualization, namely more efficient management of free space and a more evenly distributed workload to provide better performance.
- If any managed disk in a managed disk group goes offline then all the virtual disks in the managed disk group will go offline. Therefore you might want to consider using different managed disk groups for different back-end controllers or for different applications.
- If you anticipate regularly adding and removing back-end controllers or storage then this task will be made simpler by grouping all the managed disks presented by a back-end controller into one managed disk group.
- All the managed disks in a managed disk group should have similar levels of performance or reliability, or both. If a managed disk group contains managed disks with different levels of performance then the performance of the virtual disks in this group will be limited by the performance of the slowest managed disk. If a managed disk group contains managed disks with different levels of reliability then the reliability of the virtual disks in this group will be that of the least reliably managed disk in the group.

Even with the best planning, circumstances may change and you may wish to reconfigure your managed disk groups after they have been created. The data migration facilities provided by the SAN Volume Controller will allow you to move data without disrupting I/O.

**Choosing a managed disk group extent size:** You must specify the extent size when you create a new MDisk group. You cannot change the extent size later; it must remain constant throughout the lifetime of the MDisk group. MDisk groups may have different extent sizes however this will place restrictions on the use of data migration. The choice of extent size affects the total amount of storage that can be managed by a SAN Volume Controller Cluster. Table 12 on page 158 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 virtual disk that is created, using a larger extent size may increase the amount of wasted storage at the end of each virtual disk. Larger extent sizes also reduces the ability of the SAN Volume Controller to distribute sequential I/O workloads across many managed disks and hence may reduce the performance benefits of virtualization.

Table 12. Extent size

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

**Attention:** You can specify different extent sizes for different managed disk groups, however you will not be able to migrate virtual disks between managed disk groups with different extent sizes. Therefore if possible create all your managed disk groups with the same extent size.

#### Steps:

Perform the following steps to create a MDisk group:

1. Open a command prompt window.
2. Type the **svctask mkmdiskgrp** command to create a MDisk group.

#### Example:

In our hypothetical scenario, the command to create a MDisk group is:

```
svctask mkmdiskgrp -name maindiskgroup -ext 32 -mdisk mdk0:mdk1:mdk2:mdk3
```

This command will create a MDisk group called *maindiskgroup*. The extent size used within this group will be 32 MB, and there are four MDisks *mdk0*, *mdk1*, *mdk2*, *mdk3* added to the group.

#### Example:

In our hypothetical scenario, the command to create a second MDisk group is:

**Note:** In this example we will create a second MDisk group first and add MDisks later.

```
svctask mkmdiskgrp -name bkpmdiskgroup -ext 32
```

This command will create a MDisk group called *bkpmdiskgroup*. The extent size used within this group will be 32 MB.

#### Example:

To add MDisks to the MDisk group, issue the **svctask addmdisk** command. In our hypothetical scenario, the command to add MDisks to the MDisk group is:

```
svctask addmdisk -mdisk mdk4:mdk5:mdk6:mdk7 bkpdiskgroup
```

This command will add four MDisks *mdk4*, *mdk5*, *mdk6*, *mdk7* to the MDisk group called *bkpdiskgroup*.

#### Related topics:

- Chapter 14, “Scenario: typical usage for the command-line interface”, on page 149
- “Adding MDisks to MDisk groups using the CLI” on page 159

---

## Adding MDisks to MDisk groups using the CLI

This task provides step-by-step instructions for adding MDisks to MDisk groups using the command-line interface (CLI).

### Steps:

Perform the following steps to add MDisks to MDisk groups:

1. Open a command prompt window.
2. Type the **svcinfo lsmdiskgrp** command to list the existing Mdisk groups.

#### Example:

In our hypothetical scenario, we have two Mdisk groups, one with four managed disks and one with no managed disks. Type the following command:

```
svcinfo lsmdiskgrp -delim :
```

This command displays the following:

```
id:name:status:mdisk_count:vdisk_count:  
capacity:extent_size:free_capacity  
0:mainmdiskgroup:online:4:0:1093.2GB:32:1093.2GB  
1:bkpmdiskgroup:online:0:0:0:32:0
```

3. To add MDisks to the MDisk group, issue the svctask addmdisk command.

#### Example:

In our hypothetical scenario, the command to add MDisks to the MDisk group is:

```
svctask addmdisk -mdisk mdisk4:mdisk5:mdisk6:mdisk7 bkpmdiskgroup
```

This command will add four MDisks mdisk4, mdisk5, mdisk6 and mdisk7 to the MDisk group called bkpmdiskgroup.

### Related topics:

- Chapter 14, “Scenario: typical usage for the command-line interface”, on page 149

---

## Create virtual disks (VDisks)

This task provides step-by-step instructions about how to create a VDisk.

**Note:** If you want to keep the data on a MDisk (for example because you want to want to import storage that was previously not managed by a SAN Volume Controller) then you should create image mode VDisks instead.

This task only deals with creating VDisks with a striped virtualization policy. For details of other virtualization policies refer to the *IBM TotalStorage Virtualization Family SAN Volume Controller: Command-Line Interface User’s Guide*.

### Context:

Assume that the cluster has been setup and that you have created managed disk groups.

### **Steps:**

Perform the following steps to create VDisks:

1. Open a command prompt window.
2. Decide which managed disk group will provide the storage for the vdisk. Use the **svcinfo lsmdiskgrp** command to list the available managed disk groups and the amount of free storage in each group.

### **Example:**

In our hypothetical scenario, issue the following:

```
svcinfo lsmdiskgrp -delim :
```

This command displays the following:

```
id:name:status:mdisk_count:vdisk_count:  
capacity:extent_size:free_capacity  
0:mainmdiskgroup:online:4:0:1093.2GB:32:1093.2GB  
1:bkpmdiskgroup:online:4:0:546.8GB:32:546.8GB
```

3. Decide which I/O group the VDisk should be assigned to. This determines which SAN Volume Controller nodes in the cluster process the I/O requests from the host systems. If you have more than one I/O group then make sure you distribute the VDisks between the I/O groups so that the I/O workload is shared evenly between all SAN Volume Controller nodes. Use the **svcinfo lsiogrp** command to show the I/O groups and the number of virtual disks assigned to each I/O group.

**Note:** It is normal for clusters with more than one I/O group to have MDisk groups that have VDisks in different I/O groups. FlashCopy can be used to make copies of VDisks regardless of whether the source and destination VDisk are in the same I/O group. If however you plan to use intra-cluster remote copy then make sure that both the master and auxiliary VDisk are in the same I/O group.

### **Example:**

In our hypothetical scenario, there are two I/O groups each with two nodes. Neither I/O group has any virtual disks yet. Issue the following command:

```
svcinfo lsiogrp -delim :
```

This command displays the following:

```
id:name:node_count:vdisk_count  
0:io_grp0:2:0  
1:io_grp1:2:0  
2:io_grp2:0:0  
3:io_grp3:0:0  
4:recovery_io_grp:0:0
```

4. Type the **svctask mkvdisk** command to create a virtual disk (VDisk).

### **Example:**

In our hypothetical scenario, the command to create a VDisk is:

```
svctask mkvdisk -name mainvdisk1 -iogrp 0  
-mdiskgrp 0 -vtype striped -size 256 -unit gb
```

This command will create a VDisk called mainvdisk1, the VDisk will use I/O group 0 and MDisk group 0 (the ID of maindiskgroup as shown in the output from step 2). The VDisk capacity is 256GB and will be made up of extents from the MDisks in the MDisk group.

**Example:**

In our hypothetical scenario, the command to create a second VDisk is:

**Note:** This command is the same as the above example, however, here we are specifying the names of the objects instead of the IDs.

```
svctask mkvdisk -name mainvdisk2 -iogrp io_grp0  
-mdiskgrp maindiskgroup -vtype striped -size 256 -unit gb
```

This command will create a VDisk called mainvdisk2, the Vdisk will use the I/O group named io\_grp0 and the MDisk group named maindiskgroup. The VDisk capacity is 256GB and will be made up of extents from the MDisks in the MDisk group.

**Example:**

In our hypothetical scenario, the commands to create a third VDisk is:

**Note:** This virtual disk is created with an ordered list of MDisks within the MDisk group to allocate extents from.

The following command lists the managed disks in the MDisk group with ID 1 (named bkpmdiskgroup):

```
svcinfo lsmdisk -delim : -filtervalue mdisk_grp_id=1
```

This command displays the following:

```
id:name:status:mode:mdisk_grp_id:  
mdisk_grp_name:capacity:ctrl_LUN_#:  
controller_name  
4:mdisk4:online:managed:1:bkpmdiskgroup:  
136.7GB:0000000000000004:controller0  
5:mdisk5:online:managed:1:bkpmdiskgroup:  
136.7GB:0000000000000005:controller0  
6:mdisk6:online:managed:1:bkpmdiskgroup:  
136.7GB:0000000000000006:controller0  
7:mdisk7:online:managed:1:bkpmdiskgroup:  
136.7GB:0000000000000007:controller0
```

Issue the following command:

```
svctask mkvdisk -name bkpvdisk1 -iogrp io_grp1  
-mdiskgrp bkpmdiskgrp -vtype striped -size 256  
-unit gb -mdisk 4:5
```

This command will create a VDisk called bkpvdisk1, the Vdisk will use the I/O group named io\_grp1 and the MDisk group named bkpmdiskgrp. The VDisk capacity is 256GB and will be made up of extents allocated from the mdisks with IDs 4 and 5.

**Example:**

In our hypothetical scenario, the command to create a fourth VDisk is:

```
svctask mkvdisk -name bkpvdisk2 -iogrp io_grp1  
-mdiskgrp bkpmdiskgrp -vtype striped -size 256 -unit  
gb -mdisk mdisk6:mdisk7
```

This command will create a VDisk called bkpvdisk2, the Vdisk will use the I/O group named io\_grp1 and the MDisk group named bkpmdiskgrp. The VDisk capacity is 256GB and will be made up of extents allocated from the mdisks with names mdisk6 and mdisk7.

5. To list all the virtual disks that have been created use the **svcinfo lsvdisk** command.

**Example:**

In our hypothetical scenario we have created four VDisks. Issue the following command:

```
svcinfo lsvdisk -delim :
```

This command displays the following:

```
id:name:I0_group_id:I0_group_name:status:  
mdisk_grp_id:mdisk_grp_name:capacity:type:FC_id:  
FC_name:RC_id:RC_name  
0:mainvdisk1:0:io_grp0:online:0:mainmdiskgroup:  
512.0GB:striped:::::  
1:mainvdisk2:0:io_grp0:online:0:mainmdiskgroup:  
512.0GB:striped:::::  
2:bkpvdisk1:1:io_grp1:online:1:bkpmdiskgroup:  
512.0GB:striped:::::  
3:bkpvdisk2:1:io_grp1:online:1:bkpmdiskgroup:  
512.0GB:striped:::::
```

**Related topics:**

- Chapter 14, “Scenario: typical usage for the command-line interface”, on page 149
- “Creating image mode virtual disks using the CLI” on page 189

---

## Creating host objects using the CLI

This task provides step-by-step instructions about how to create host objects.

**Steps:**

Perform the following steps to create host objects:

1. Open a command prompt window.

2. Type the **svctask mkhost** command to create a logical host object. Assign your WWPN for the HBAs in the hosts.

**Example:**

In our hypothetical scenario, the command to create a host is:

```
svctask mkhost -name demohost1 -hbawwpn 210100e08b251dd4
```

This command will create a host called *demohost1* with the HBA WWPN of *210100e08b251dd4*.

3. Type the **svctask addhostport** command to add ports to the host.

**Example:**

In our hypothetical scenario, the command to add a port to the host is:

```
svctask mkhost -name demohost2 -hbawwpn 210100e08b251dd5
```

This command will add another HBA WWPN called *210100e08b251dd5* to the host that we created in step 2.

**Example:**

In our hypothetical scenario, the command to create a second host is:

```
svctask mkhost -hbawwpn 210100e08b251dd6:210100e08b251dd7 -name demohost2
```

This command will create a second host called *demohost2* with the HBA WWPN of *210100e08b251dd6*, *210100e08b251dd7*.

**Note:** If you were to add a host with a faulty WWPN, or the WWPN had been assigned to the wrong host, you will need to issue the **svctask addhostport** command to add that same host with the correct WWPN, then issue the **svctask rmhostport** command to delete the host with the wrong or faulty WWPN. For example, if you had a host called *demohost1* and its WWPN stopped working, you would need to issue the following:

```
svctask addhostport -hbawwpn 210100e08b251dd4 demohost1
```

This would add the host called *demohost1* with the WWPN, *210100e08b251dd4*. You would then need to issue the **svctask rmhostport** command to delete the host with the WWPN that had stopped working. For example, you would issue the following:

```
svctask rmhostport -hbawwpn 210100e08b251dd5 demohost1
```

From these two commands, you have deleted the host with the WWPN *210100e08b251dd5*, and have added the same host with the WWPN *210100e08b251dd4*.

**Related topics:**

- Chapter 14, “Scenario: typical usage for the command-line interface”, on page 149

---

## Create VDisk-to-host mappings using the CLI

This task provides step-by-step instructions about how to create VDisk-to-host mappings.

**Prerequisites:**

We are going to map the VDisks named, mainvdisk1 and mainvdisk2, to the host named demohost1. We are also going to map the VDisks named, bkpvdisk1 and bkpvdisk2, to the host named demohost2. The VDisks, mainvdisk1 and mainvdisk2, are contained in the managed disk (MDisk) group, mainmdiskgroup; while the VDisks, bkpvdisk1 and bkpvdisk2, are contained in the MDisk group, bkpmdiskgroup.

#### Steps:

Perform the following steps to create VDisk-to-host mappings:

1. Open a command prompt window.
2. Type the **svctask mkvdiskhostmap** to create VDisk-to-host mappings.

#### Example:

In our hypothetical scenario, the commands to create VDisk-to-host mappings are:

```
svctask mkvdiskhostmap -host demohost1 mainvdisk1
svctask mkvdiskhostmap -host demohost1 mainvdisk2
svctask mkvdiskhostmap -host demohost2 bkpvdisk1
svctask mkvdiskhostmap -host demohost2 bkpvdisk2
```

The above set of commands map each VDisk to a host.

#### Related topics:

- Chapter 14, “Scenario: typical usage for the command-line interface”, on page 149

---

## Create FlashCopy mappings using the CLI

This task provides step-by-step instructions on how to create FlashCopy mappings.

#### Prerequisites:

We are going to create mappings that enable us to copy the VDisk, mainvdisk1, to bkpvdisk1 and the VDisk, mainvdisk2 to bkpvdisk2.

#### Steps:

Perform the following steps to create FlashCopy mappings:

1. Open a command prompt window.
2. Type the **svctask mkfcmap** command to create a FlashCopy mapping.

#### Example:

In our hypothetical scenario, the commands to create FlashCopy mappings are:

```
svctask mkfcmap -source mainvdisk1 -target bkpvdisk1
  -name main1copy -copyrate 75
svctask mkfcmap -source mainvdisk2 -target bkpvdisk2
  -name main2copy
```

The above commands create two FlashCopy mappings. For main1copy the background copy rate is 75; for main2copy, because the rate is not specified in the **mkfcmap** command, the priority is the default, 50.

3. To check the attributes of the mappings that have been created, issue the following **svcinfo lsfcmap** command:

```
svcinfo lsfcmap -delim :
```

This command displays the following:

```
id:name:source vdisk id:source vdisk name:target  
    vdisk id:target vdisk name:group id:group  
    name:status:progress:copy rate  
0:main1copy:0:mainvdisk1:1:bkpvdisk1:::idle_copied:::75  
1:main2copy:2:mainvdisk2:3:bkpvdisk2:::idle_copied:::50
```

**Related topics:**

- Chapter 14, “Scenario: typical usage for the command-line interface”, on page 149

---

## Create a FlashCopy consistency group and add mappings using the CLI

This task provides step-by-step instructions on how to create a FlashCopy Consistency Group and add mappings to it.

If you have created several FlashCopy mappings for a group of VDisks that contain elements of data for the same application, you may find it convenient to assign these mappings to a single FlashCopy Consistency Group. Then you can issue a single prepare or trigger command for the whole group, so that, for example, all the files for a particular database are copied at the same time.

**Steps:**

Perform the following steps to create a FlashCopy mappings:

1. Open a command prompt window.
2. Issue the **svctask mkfcconsistgrp** command to create a FlashCopy Consistency Group.

**Example:**

In our hypothetical scenario, the command to create a FlashCopy Consistency group called *maintobkpfcopy* is:

```
svctask mkfcconsistgrp -name maintobkpfcopy
```

Use the **svcinfo lsfcconsistgrp** command to display the attributes of the group you have created.

```
svcinfo lsfcconsistgrp -delim :
```

This command displays the following:

```
id:name:status  
1:maintobkpfcopy:idle_copied
```

3. Use the **svctask chfcmap** command to add the two FlashCopy mappings created in the previous section to the new consistency group.

**Example:**

In our hypothetical scenario, the commands to add the mappings called *main1copy* and *main2copy* to the consistency group called *maintobkpfcopy* are:

```
svctask chfcmap -consistgrp maintobkpfcopy main1copy  
svctask chfcmap -consistgrp maintobkpfcopy main2copy
```

Use the **svcinfo lsfcmap** command to display the new attributes of the mappings.

```
svcinfo lsfcmap -delim :  
id:name:source_vdisk_id:source_vdisk_name:target_vdisk_id:  
target_vdisk_name:group_id:group_name:state:progress:copy_rate  
0:main1copy:28:maindisk1:29:bkpdisk1:1:maintobkpfcopy:idle_copied::75  
1:main2copy:30:maindisk2:31:bkpdisk2:1:maintobkpfcopy:idle_copied::50
```

Notice that the group\_name field displays maintobkpfcopy for both mappings.

Use the **svcinfo lsfcconsistgrp** command with the name of the consistency group to display the detailed attributes of the group. This now includes a list of the IDs and names of the mappings that are in the group.

```
svcinfo lsfcconsistgrp -delim : maintobkpfcopy  
id:1  
name:maintobkpfcopy  
status:idle_copied  
FC_mapping_id:0  
FC_mapping_name:main1copy  
FC_mapping_id:1  
FC_mapping_name:main2copy
```

**Related topics:**

- Chapter 14, “Scenario: typical usage for the command-line interface”, on page 149

---

## Prepare and trigger a FlashCopy Consistency Group using the CLI

This task provides step-by-step instructions on how to prepare and trigger a FlashCopy Consistency Group to start the flash copy process. This will create a point-in-time copy of the data on the source VDisk and write it to the target VDisk for each mapping in the group.

**Steps:**

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

1. Open a command prompt window.
2. Issue the **svctask prestartfcconsistgrp** command to prepare the FlashCopy Consistency Group before the copy process can be started (triggered). When you have assigned several mappings to a FlashCopy Consistency Group, you only have to issue a single prepare command for the whole group, to prepare all the mappings at once.

**Example:**

In our hypothetical scenario, the command to prepare a FlashCopy Consistency group called *maintobkpfcopy* is:

```
svctask prestartfcconsistgrp maintobkpfcopy
```

The group will enter the preparing state, and then move to the prepared state when it is ready. Issue the **svcinfo lsfcconsistgrp** command to check:

```
svcinfo lsfcconsistgrp -delim :  
id:name:status  
1:maintobkpfcopy:prepared
```

3. Issue the **svctask startfcconsistgrp** command to start (trigger) the FlashCopy Consistency Group to make the copy. You only have to issue a single start command for the whole group, to trigger all the mappings at once.

**Example:**

In our hypothetical scenario, the command to trigger a FlashCopy Consistency group called *maintobkpfcopy* is:

```
svctask startfcconsistgrp maintobkpfcopy
```

The group will enter the copying state, and then return to the idle\_copied state when complete. You can issue the **svcinfo lsfcconsistgrp** command to check the state of the group:

```
svcinfo lsfcconsistgrp -delim :  
id:name:state  
1:maintobkpfcopy:copying
```

Use the **svcinfo lsfcmapprogress** command to check the progress of each mapping, main1copy and main2copy:

```
svcinfo lsfcmapprogress -delim : main1copy id:progress 0:100  
svcinfo lsfcmapprogress -delim : main2copy  
id:progress  
1:23
```

Finally issue the **svcinfo lsfcconsistgrp** command to display the detailed view of the group *maintobkpfcopy*, which returns to idle\_copied state when both mappings have reached 100% progress:

```
svcinfo lsfcconsistgrp -delim : maintobkpfcopy  
id:1  
name:maintobkpfcopy  
state:idle_copied  
FC_mapping_id:0  
FC_mapping_name:main1copy  
FC_mapping_id:1  
FC_mapping_name:main2copy
```

You have now made a point-in-time copy of the data on *mainvdisk1* which has been written to *bkpvdisk1*, and a copy of the data on *mainvdisk2* which has been written to *bkpvdisk2*. The data on *bkpvdisk1* and *bkpvdisk2* will be visible to *demohost2* because these VDisks are only mapped to *demohost2*.

**Related topics:**

- Chapter 14, “Scenario: typical usage for the command-line interface”, on page 149



---

## Chapter 15. Advanced functions overview for the CLI

This topic provides overview information about the advanced functions that you are able to perform using the CLI. More specifically, it provides step-by-step instructions for the following advanced functions:

- “Determining a nodes WWPNs using the CLI”
- “Determining a storage controller name from its SAN Volume Controller name” on page 170
- “Determining the VDisk name from the vpath number on the host” on page 170
- “Determining the host that a VDisk is mapped to” on page 171
- “Determining the relationship between VDisks and MDisks using the CLI” on page 171
- “Determining the relationship between MDisks and RAID arrays or LUNs using the CLI” on page 173
- “Increasing the size of your cluster using the CLI” on page 173
- “Adding a node to increase the size of your cluster using the CLI” on page 174
- “Migrating a VDisk to a new I/O group” on page 174
- “Replacing a faulty node in the cluster using the CLI” on page 175
- “Recovering from offline VDisks after a node or an I/O group failed using the CLI” on page 176
- “Replacing an HBA in a host using the CLI” on page 179
- “Adding a new storage controller to a running configuration using the CLI” on page 180
- “Removing a storage controller using the CLI” on page 181
- “Expanding a VDisk using the CLI” on page 182
- “Shrinking a VDisk using the CLI” on page 185
- “Migrating extents using the CLI” on page 185
- “Migrating VDisks between MDisk groups using the CLI” on page 187
- “Migrating a VDisk between I/O groups using the CLI” on page 188
- “Advanced function FlashCopy and Remote Copy overview for CLI” on page 190
- “Advanced function cluster overview using the CLI” on page 190

For guidelines about creating MDisk groups, see “Guidelines for MDisk group creation” on page 111.

---

### Determining a nodes WWPNs using the CLI

This task provides step-by-step instructions for determining a nodes WWPNs using the CLI.

#### Steps:

Perform the following steps to determine a nodes WWPNs:

1. List the nodes in the cluster by issuing the following command:

```
svctools info lsnode
```

**Note:** Remember the node name or ID as you will need it in the next step.

2. For the node or nodes in question, issue the following command:

```
svcinfo lsnode <nodename/id>
```

where *<nodename/id>* is the node name or ID.

**Note:** Remember the four port ID's (WWPNs).

---

## Determining a storage controller name from its SAN Volume Controller name

This task provides step-by-step instructions for determining a storage controller name from its SAN Volume Controller name.

### Steps:

Perform the following steps to determine a storage controller name:

1. List the storage controllers by issuing the following command:

```
svcinfo lscontroller
```

Remember the controller name or ID for the controller you want to determine.

2. For the controller in question, issue the following command:

```
svcinfo lscontroller <controllername/id>
```

where *<controllername/id>* is the controller name or ID. Remember the WWNN for the controller. Make a written record of it. The WWNN can be used to determine the actual storage controller by launching the native controller user interface or using the command line tools it provides to verify the actual controller that has this WWNN.

---

## Determining the VDisk name from the vpath number on the host

This task provides step-by-step instructions about how to determine the VDisk name from the vpath number on the host.

Each VDisk exported by the SAN Volume Controller is assigned a unique vpath number. This number uniquely identifies the VDisk and can be used to determine which VDisk corresponds to the volume that the hosts sees. This procedure can only be performed using the command line interface.

### Steps:

Perform the following steps to determine the VDisk name from the vpath number:

1. For the volume in question, find the vpath serial number by issuing the following command:

```
datapath query device
```

2. Find the host object defined to the SAN Volume Controller that corresponds with the host you are working with.

- a. The WWPNs are an attribute of the HBA. You can find these by looking at the device definitions stored by your operating system. For example, on AIX they will be in the ODM, in Windows they will be in the Device Manager details for the given HBA.
- b. Verify which host object defined to the SAN Volume Controller that these ports belong to. The ports are stored as part of the detailed view, so you will need to list each host in turn by issuing the following:

```
svcinfo lshost <name/id>
```

where *<name/id>* is the name or ID of the host. Check for matching WWPNs.

**Note:** You should name your hosts accordingly, for example, if the actual host is called *orange* you should also name the host object defined to the SAN Volume Controller as *orange*.

3. Now that you have the *<host name>* as defined to the SAN Volume Controller and the *<vpath serial number>*, issue the following command:

```
svcinfo lshostvdiskmap <hostname>
```

where *<hostname>* is the name of the host. A list is displayed.

4. Look for the VDisk UID that matches the *<vpath serial* and remember the VDisk name or ID.

---

## Determining the host that a VDisk is mapped to

This task provides step-by-step instructions for determining the host that a VDisk is mapped to.

### Steps:

Perform the following steps to determine the host that the VDisk is mapped to:

1. Find the VDisk name or ID that you wish to check.
2. List the hosts that this VDisk is mapped, by issuing the following command:

```
svcinfo lsvdiskhostmap <vdiskname/id>
```

where *<vdiskname/id>* is the name or ID of the VDisk. A list is displayed.

3. Look for the host name or ID to determine which host this VDisk is mapped to. If no data is returned, the VDisk is not mapped to any hosts.

---

## Determining the relationship between VDisks and MDisks using the CLI

This task provides step-by-step instructions for determining the relationship between VDisks and MDisks.

Every vdisk is constructed from one or more mdisks. At times you may need to determine the relationship between the two objects. The following procedure allows you to determine the relationships.

### Steps:

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

1. For a given VDisk <vdiskname/id>, issue the following command:

```
svcinfo lsvdiskmember <vdiskname/id>
```

where <vdiskname/id> is the name or ID of the VDisk. This will return a list of IDs that correspond to the MDisks that make up the VDisk.

**Steps:**

Perform the following steps to determine the relationship between VDisks and MDisks and the number of extents provided by each MDisk:

If you wish more details, you can also determine the number of extents that make are being provided by each MDisk. This procedure can only be performed using the command line interface.

1. For a given VDisk <vdiskname/id>, issue the following command:

```
svcinfo lsvdiskextent <vdiskname/id>
```

where <vdiskname/id> is the name or ID of the VDisk. This will return a table of MDisk IDs and the corresponding number of extents each MDisk is providing as storage for the given VDisk.

**Steps:**

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

1. For a given MDisk <mdiskname/id>, issue the following command:

```
svcinfo lsmdiskmember <mdiskname/id>
```

where <mdiskname/id> is the name or ID of the MDisk. This will return a list of IDs that correspond to the VDisks that are using this MDisk.

**Steps:**

Perform the following steps to determine the relationship between MDisks and VDisks and the number of extents used by each VDisk:

If you wish more details, you can also determine the number of extents that this MDisk is providing for each VDisk. This procedure can only be performed using the command line interface.

1. For a given MDisk <mdiskname/id>, issue the following command:

```
svcinfo lsmdiskextent <mdiskname/id>
```

where <mdiskname/id> is the name or ID of the MDisk. This will return a table of VDisk IDs and the corresponding number of extents being used by each VDisk.

---

## Determining the relationship between MDisks and RAID arrays or LUNs using the CLI

This task provides step-by-step instructions for determining the relationship between MDisks and RAID arrays or LUNs using the CLI.

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

### Steps:

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

1. Show the detailed view of the given MDisk <mdiskname>, by issuing the following command:

```
svcinfo lsmdisk <mdiskname>
```

where <mdiskname> is the name of the MDisk.

**Note:** Remember the controller name or controller ID and controller LUN number.

2. Show the detailed view of the controller determined in 1 by issuing the following command:

```
svcinfo lscontroller <controllername>
```

where <controllername> is the name of the controller.

**Note:** Remember the vendor ID, product ID, and WWNN. Use these to determine what is being presented to the MDisk.

3. From the native user interface for the given controller, list the LUNs it is presenting and match the LUN number with that noted in 1. This will tell you the exact RAID array or partition that corresponds with the MDisk.

---

## Increasing the size of your cluster using the CLI

This task provides step-by-step instructions for increasing the size of your cluster.

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

### Steps:

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

1. Perform the steps in the “Adding a node to increase the size of your cluster using the CLI” on page 174 section and repeat this procedure for the second node.

2. If you wish to balance the load between the existing I/O groups and the new I/O groups, follow the “Increasing the size of your cluster using the CLI” on page 173 procedure. Repeat this procedure for all VDisks you want to assign to the new I/O group.

## Adding a node to increase the size of your cluster using the CLI

This task provides step-by-step instructions for adding a node to increase the size of your cluster using the CLI.

### Steps:

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

1. Issue the following command to verify that the node can be seen on the fabric:

```
svcinfo lsnodecandidate
```

You should see the node listed as a candidate.

**Note:** Remember the WWNN’s. You will need it in the following step.

2. Issue the following command to determine the I/O group you wish to add the nodes to:

```
svcinfo lsiogrp
```

3. Select the first I/O group listed that has a node count = 0.

**Note:** Remember the I/O group name or ID. You will need it in the following step.

4. Issue the following command to add the node into the cluster. The <newnodename> is the name you wish to assign to this node.

```
svctask addnode -wwnodename <WWNN> -iogrp <newiogrpname/id>
]
```

5. Issue the following command to verify that the node is online:

```
svcinfo lsnode
```

You may also need to modify the configuration of your Disk Controllers. If your controller uses a mapping technique to present its RAID arrays or partitions to the cluster you will need to modify the port groups that belong to the cluster because the WWNN and WWPN’s of the node have changed. See Chapter 25, “Configuring disk controllers”, on page 223 for more information.

## Migrating a VDisk to a new I/O group

This task provides step-by-step instructions for migrating a VDisk to a new I/O group to increase the size of your cluster using the CLI.

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

**Attention:** This is a disruptive procedure, access to the VDisk will be lost while you follow this procedure.

**Steps:**

Perform the following steps to migrate a single VDisk:

1. Quiesce all I/O operations for the VDisk. You may need to determine the hosts that are using this vdisk. See “Determining the host that a VDisk is mapped to” on page 171 and “Determining the relationship between VDisks and MDisks using the CLI” on page 171 for more information.
2. Before migrating the VDisk, it is essential that for each vpath presented by the VDisk you intend to move, the SDD configuration is updated to remove the vpaths in question. Failure to do this may result in data corruption. See *IBM Subsystem Device Driver (SDD) User’s Guide* for details about how to dynamically reconfigure SDD for the given host operating system.
3. Any FlashCopy mappings or Remote Copy relationships that use this VDisk should be stopped or deleted. Issue the following command, to check if the VDisk is part of a relationship or mapping:

```
svcinfo lsvdisk <vdiskname/id>
```

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

4. Look for the **FC\_id** and **RC\_id** fields. If these are not blank then the VDisk is part of a mapping or relationship. See “Advanced function FlashCopy and Remote Copy overview for CLI” on page 190 and *IBM TotalStorage Virtualization Family SAN Volume Controller Command-Line Interface User’s Guide* for details on how to stop or delete the mapping or relationship.
5. Issue the following command to migrate the VDisk:

```
svctask chvdisk -iogr <newiogrname/id> <vdiskname/id>
```

6. Follow the *IBM Subsystem Device Driver (SDD) User’s Guide* procedure to discover the new vpaths and to check that each vpath is now presenting the correct number of paths. See the *IBM Subsystem Device Driver (SDD) User’s Guide* for details on how to dynamically reconfigure SDD for the given host operating system.

---

## Replacing a faulty node in the cluster using the CLI

This task provides step-by-step instructions for replacing a faulty node in the cluster using the CLI.

**Steps:**

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

1. Issue the following command to verify the name or ID of the node you wish to remove.

```
svcinfo lsnode
```

- a. If the node was faulty it will be shown as offline. Ensure the partner node in the I/O group is online.
  - 1) If the other node in the I/O group is offline, start the Directed Maintenance Procedures to determine the fault.

- 2) If you have been directed here by the DMPs, and subsequently the partner node in the I/O group has failed, see “Recovering from offline VDisks after a node or an I/O group failed using the CLI” for more information.
- b. If you are replacing the node for other reasons, determine the node you wish to replace and again ensure the partner node in the I/O group is online.
  - 1) If the partner node is offline, you will lose access to the VDisks that belong to this I/O group if you continue. Start the Directed Maintenance Procedures and fix the other node before proceeding.
2. Having noted the <nodename> in step 1 on page 175, remove the node from the cluster by issuing the following command:

```
svctask rmnode <nodename/id>
```

3. Issue the following command to verify that the node can be seen on the fabric:

```
svcinfo lsnodecandidate
```

You should see the node listed as a candidate.

**Note:** Remember the WWNN's for each node, you will need it in the following step.

4. If the nodes were 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, it will be necessary to follow the SDD procedure to discover the new paths and to check that each vpath is now presenting the correct number of paths. See the *IBM Subsystem Device Driver (SDD) User's Guide* sections on dynamic reconfiguration, specifically adding paths to existing vpaths.
  - b. You may also need to modify the configuration of your disk controllers. If your controller uses a mapping technique to present its RAID arrays or partitions to the cluster, you will need to modify the port groups that belong to the cluster because the WWNN or WWPN's of the node have changed. See Chapter 25, “Configuring disk controllers”, on page 223 for more information.
5. Issue the following command to add the node back into the cluster:

```
svctask addnode -wwnnode <WWNN> -iogrp  
<IOPNAME/ID> -name <NODENAME>
```

6. Issue the following command to verify that the node is online:

```
svcinfo lsnode
```

## Recovering from offline VDisks after a node or an I/O group failed using the CLI

This task provides step-by-step instructions for recovering from an offline 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, then you need to perform one

of the following procedures to recover the VDisks and regain access to them. Depending on the failure type, you may have lost data that was cached for these VDisks, therefore, they have gone offline.

**Steps:**

Perform the following steps to recover from an offline VDisk:

**Example:**

**Data loss scenario 1** One node in an I/O group failed and failover started on the second node. During this time, the second node in the I/O group fails before the cache has become write through mode. The first node is successfully repaired but its cache data is stale, therefore, it cannot be used. The second node is repaired or replaced and has lost its hardend data, therefore, the node does not think that it is part of the cluster. You need to perform the following:

1. Perform the recovering of the node back into the cluster procedure.
2. For each VDisk, move the offline VDisks to the recovery I/O group procedure.
3. For each VDisk, move the offline VDisks back to their original I/O group procedure.

**Example:**

**Data loss scenario 2** Both nodes in the I/O group have failed and have been repaired. The nodes have lost their hardend data, therefore, the nodes do not think that they are part of the cluster. You need to perform the following:

1. For each VDisk, move the offline VDisk to the recovery I/O group.
2. For both nodes, move the recovered nodes back in to the cluster.
3. For each VDisk, move the offline VDisks back to their original I/O group.

**Recover node back into cluster** Perform the following to recover a node back in to the cluster:

1. Verify that the node is offline. Issue the following command:

```
svcinfo lsnode
```

2. Remove the old instance of the offline node from the cluster. Issue the following command:

```
svctask rmnode <nodename/id>
```

where *<NODENAME>* is the name of the node.

3. Verify that the node can be seen on the fabric. Issue the following command:

```
svcinfo lsnodecandidate
```

You should see the nodes listed as a candidate.

**Note:** Remember the WWNNs for each node, you will need it in the following step.

4. If the 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 it will be necessary to follow the SDD procedure to discover the new paths and to check that each vpath is now presenting the correct number of paths. See the *Subsystem Device Driver (SDD) User's Guide* sections on dynamic reconfiguration, specifically adding paths to existing vpaths.
  - b. You may also need to modify the configuration of your disk controllers. If your controller uses a mapping technique to present its RAID arrays or partitions to the cluster you will need to modify the port groups that belong to the cluster because the WWNN or WWPN's of the node have changed. See Chapter 25, "Configuring disk controllers", on page 223 for more information.
5. Add the node back into the cluster. Issue the following command:

```
svctask addnode -wwnnode <WWNN> -iogrp
<IOGRPNAME/ID> [-name <NODENAME>]
```

where <WWNN> is the worldwide node name <IOGRPNAME/ID> is the I/O group name or ID.

6. Verify that the node is online. Issue the following command:

```
svcinfo lsnode
```

**Move the offline VDisks to the recovery I/O group** Perform the following steps to move the offline VDisks to the recovery I/O group:

- <IOGRPNAME> = the name of the I/O group that failed.
  - <vdiskname/ID> = the name of one of the VDisks that are offline.
1. List all VDisks that are offline and belong to the I/O group in question. Issue the following command:

```
svcinfo lsvdisk -filtervalue iogrpname=
<IOGRPNAME/ID>:status=offline
```

2. For each VDisk returned, move the VDisk to the recovery I/O group. Issue the following command:

```
svctask chvdisk iogrp recovery_io_grp -force
<vdiskname/ID>
```

**Move the offline VDisks back to their original I/O group** Perform the following steps to move the offline VDisks back to their original I/O group:

- <IOGRPNAME> = the name of the I/O group that failed.
  - <vdiskname/ID> = the name of one of the VDisks that are offline.
1. For each VDisk, move the VDisk back into the original I/O group. Issue the following command:

```
svctask chvdisk iogrp <IOGRPNAME/ID> -force
<vdiskname/ID>
```

2. Verify that the VDisks are now online. Issue the following command:

```
svcinfo lsvdisk -filtervalue iogrpname=
<IOGRPNAME/ID>
```

---

## Replacing an HBA in a host using the CLI

This task provides step-by-step instructions for replacing an HBA in a host using the CLI.

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

### Prerequisites:

Ensure your switch is zoned correctly. See Chapter 25, "Configuring disk controllers", on page 223 for more information.

### Steps:

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

- Issue the following command to list the candidate HBA ports:

```
svcinfo lshbaportcandidate
```

You should see a list of the HBA ports that are available to be added to host objects. One or more of these should correspond with the one or more WWPNs that belong to the new HBA.

- Locate the host object that corresponds with the host in which you have replaced the HBA. The following command lists all the defined host objects:

```
svcinfo lshost
```

To list the WWPNs currently assigned to the host, issue the following:

```
svcinfo lshost <hostobjectname>
```

where *<hostobjectname>* is the name of the host object.

- Add the new ports to the existing host object by issuing the following command:

```
svctask addhostport -hbawwpn <one or more existing WWPNs separated by :> <hostobjectname/ID>
```

where *<one or more existing WWPNs separated by :>* correspond with those listed in step 1 and *<hostobjectname/id>* corresponds with the host object located in step 2.

- Remove the old ports from the host object by issuing the following command:

```
svctask rmhostport -hbawwpn <one or more existing WWPNs separated by :> <hostobjectname/ID>
```

where *<one or more existing WWPNs separated by :>* correspond with those listed in step 2 that belong to the old HBA that has been replaced.

- Any mappings that exist between the host object and VDisks will automatically be applied to the new WWPNs. Therefore, the host should see the VDisks as the same SCSI LUNs as before.

6. See the *Subsystem Device Driver (SDD) User's Guide* for additional information about dynamic reconfiguration

---

## Adding a new storage controller to a running configuration using the CLI

This task provides step-by-step instructions for adding a new storage controller to a running configuration.

### Prerequisites:

You can add a new storage controller to your SAN at any time. You should follow the switch zoning guidelines provided in Chapter 26, “Overview about zoning a switch”, on page 233 and also ensure the controller is setup correctly for use with the SAN Volume Controller Chapter 25, “Configuring disk controllers”, on page 223.

You should create one or more arrays on the new controller. It is recommend that you use, RAID-5, RAID-1 or RAID-0+1 (sometimes called RAID-10) for maximum redundancy and reliability. Generally 5+P arrays are recommend. If your controller provides array partitioning we recommend that you create a single partition from the entire capacity available in the array, remember the LUN number that you assign to each partition as you will need this later. You should also follow the mapping guidelines (if your storage controller requires LUN mapping) to map the partitions or arrays to the SAN Volume Controller ports. You can determine the SAN Volume Controller ports by following the “Determining a nodes WWPNs using the CLI” on page 169.

### Steps:

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

1. To ensure that the cluster has detected the new storage (MDisks) issue the following command:

```
svctask detectmdisk
```

2. The controller itself will have automatically been assigned a default name. If you are unsure which controller is presenting the MDisks, list the controllers by issuing the following command:

```
svcinfo lscontroller
```

You should see a new controller listed (the one with the highest numbered default name). Remember the controller name and follow the instructions in the “Determining a storage controller name from its SAN Volume Controller name” on page 170 section.

3. You should give this controller a name that you can easily use to identify it. Issue the following command:

```
svctask chcontroller -name <newname> <oldname>
```

4. List the unmanaged MDisks by issuing the following command:

```
svcinfo lsmdisk -filtervalue mode=unmanaged:controller_name=<new_name>
```

These MDisks should correspond with the RAID arrays or partitions you have created. Remember the field controller LUN number. This corresponds with the LUN number you assigned to each of the arrays or partitions.

5. It is recommended that you create a new managed disk group and add only the RAID arrays that belong to the new controller to this MDisk group. You should also avoid mixing RAID types, so for each set of RAID array types (for example, RAID-5, RAID-1) you should create a new MDisk group. Give this MDisk group an appropriate name, so if your controller is called FAST650-fred, and the MDisk group contains RAID-5 arrays, call it something like F600-fred-R5). Issue the following command:

```
svctask mkmdiskgrp -ext 16 -name <mdisk_grp_name>
-mdisk <colon separated list of RAID-x mdisks returned
in step 4.
```

**Note:** This will create a new MDisk group with an extent size of 16MB.

---

## Removing a storage controller using the CLI

This task provides step-by-step instructions for removing a storage controller.

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

This function can also be performed by migrating all the VDisks that are using storage in this MDisk group to another MDisk group. This procedure has an advantage if you wish to consolidate the VDisks in a single or new group.

However, you can only migrate a single VDisk at a time. The procedure outlined below will migrate all the data through a single command. If you wish to migrate the VDisks however, follow the “Migrating VDisks between MDisk groups using the CLI” on page 187 procedure for all VDisks that are using this group. You can determine the relationship between VDisks and MDisks by following the “Determining the relationship between VDisks and MDisks using the CLI” on page 171 procedures.

This procedure can also be used to remove or replace a single MDisk in a group. If an MDisk has suffered a partial failure, such as a degraded array, and you can still read the data from the disk but cannot write to it, you can follow this procedure to replace just one MDisk. In steps 1 and 3 on page 182 only add or remove a single MDisk rather than a list of MDisks.

### Prerequisites:

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

### Steps:

Perform the following steps to remove a storage controller:

1. Perform steps 1 on page 180 through 4 on page 180 from the “Adding a new storage controller to a running configuration using the CLI” on page 180 section.
2. Issue the following command:

```
svctask addmdisk -mdisk <colon separated mdisk  
list as determined in step 4> <mdisk_grp_name>
```

Where *<mdisk\_grp\_name>* is the name of the MDisk group that contains the MDisks that are being decommissioned.

3. You should now have an MDisk group that contains the old MDisks (those to be decommissioned) and the new MDisks (those that are replacing them). Ensure that the capacity of the new MDisks is the same or exceeds that of the old MDisks before proceeding.
4. Force delete the old MDisks from the group. This will migrate all the data from the old MDisks to the new MDisks. Issue the following command:

```
svctask rmmdisk -force -mdisk <colon separated  
mdisk list of all the old mdisks> <mdisk_grp_name>
```

Depending upon the number and size of the MDisks, and the number and size of the VDisks that are using these MDisks, this operation will take some time to complete although the command will return immediately.

5. Check progress by issuing the following command:

```
svcinfo lsmigrate
```

6. When all the migration tasks have completed, for example, the command in step 4 returns no output, you can safely remove the old controller from the SAN.
7. Once you have removed the old controller from the SAN, re-run the **svctask detectmdisk** command to remove the entries for the old MDisks.

---

## Expanding a VDisk using the CLI

This task provides step-by-step instructions for expanding a VDisk using the CLI.

VDisks can be expanded should it be required. However, if the VDisk contains data that is being used, only AIX and Windows 2000 hosts can cope with a VDisk being expanded. A VDisk that is not yet mapped to any hosts and hence does not contain any customer data can be expanded at any time.

This feature can be used in two ways:

1. To increase the capacity available on a particular VDisk that is already mapped to a host. The procedure below outlines using this feature in this way. The following matrix shows the supported platforms and requirements if this feature is to be used:

*Table 13. Supported platforms and requirements*

Platform	Supported	Requirement
AIX	Yes	AIX 5.2 onwards only
HP-UX	No	
Linux	No	
SUN Solaris	No	
Windows NT	No	
Windows 2000	Yes	

2. To increase the size of a VDisk so that it matches the size of the source or master VDisk and can be used in a FlashCopy mapping or Remote Copy relationship. The following procedure does not need to be followed to use this feature in this way. Ensure the target or auxiliary VDisk is not mapped to any host and issue the following command:

```
svctask expandvdisksize
```

**Note:** You can determine the exact size of the source or master VDisk by issuing the following command:

```
svcinfo lsvdisk -bytes <vdiskname>
```

**Attention:** This procedure allows you to expand a VDisk that is attached to an AIX 5.2 host only. This procedure will deactivate the volume group to which the VDisk belongs to for a short period of time. That is, I/O operations will not be possible to any filesystem created in the volume group to which the VDisk that is being expanded belongs.

#### Steps:

Perform the following steps to expand a VDisk that is mapped to a AIX host:

1. Determine the VDisk you wish to expand and remember its <vdiskname>. You may need to follow the “Determining the VDisk name from the vpath number on the host” on page 170 procedure.
2. Determine the host that this VDisk is mapped to. You may need to follow the “Determining the host that a VDisk is mapped to” on page 171 procedure. If any of the hosts are running HP-UX, Linux, Solaris or Windows NT, then you cannot continue.
3. Determine the volume group that the VDisk belongs to (it is assumed you know the VDisk to HDisk relationship as determined in step 1).
4. Quiesce all I/O operations to *all* volumes that belong to the volume group and sync the filesystems mounted on this volume group.
5. Issue the following command to find out the current type of the VDisk:

```
svcinfo lsvdisk <vdiskname>
```

If the VDisk has a type of *image* it cannot be expanded. If the VDisk has a type of *sequential* it will become striped when you expand it.

6. You must first deactivate the volume group to which this VDisk belongs to. Issue the following command:

```
<AIIX_HOST_PROMPT> varyoffvg <volume_group>
```

7. Expand the VDisk by the required amount by issuing the following command:

```
svctask expandvdisksize -size <capacitytoexpandby>
-unit <unitsforexpansion> [ -mdisk <mdisklist> -fmtdisk ]
<vdiskname/ID>
```

For more information on this command see *IBM TotalStorage Virtualization Family SAN Volume Controller Command-Line Interface User's Guide*.

8. Re-activate the volume group so that the change in size is detected by the HBA device driver. Issue the following command:

- <AIX\_HOST\_PROMPT> varonvg <volume\_group>
9. Run the change volume group to notify the LVM that the size has changed.  
Issue the following command:

```
chvg -g <volume_group>
```

10. Expand the filesystem(s) that are mounted on this VDisk (or use the new capacity as required).
11. Restart I/O operations to the VDisk.

#### **Steps:**

Perform the following steps to expand a VDisk that is mapped to a Windows 2000 host:

1. Ensure that you have run Windows Update and have applied all recommended updates to your system prior to attempting to expand a VDisk that is mapped to a Windows 2000 host.

**Note:** VDisks can be expanded under Windows 2000 concurrently with I/O operations.

2. Expand the VDisk by the required amount by issuing the following command:  
`svctask expandvdisksize -size <capacitytoexpandby> - unit <unitsforexpansion>`
3. On the Windows Host, start the Computer Management application and open the Disk Management window under the Storage branch.  
Restart the Computer Management application if it was opened prior to expanding the VDisk and the disk.
4. You will see the disk that you expanded now has some unallocated space at the end of the disk.

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 etc) or add it to an existing volume.

Dynamic disks can be expanded without stopping IOs in most cases. However, in some applications the operating system may report IO errors. When this problem occurs, either of the following entries may be recorded in the System event log:

```
Event Type: Information
Event Source: dmio
Event Category: None
Event ID: 31
Description:
dmio: Harddisk0 write error at block ##### due to
disk removal

Event Type: Information
Event Source: dmio
Event Category: None
Event ID: 34
Description:
dmio: Harddisk0 is re-online by PnP
```

**Attention:** This is a known problem with Windows 2000 and is documented at the Microsoft knowledge base as article Q327020. If either of these errors are seen, run Windows Update and apply the recommended fixes to resolve the problem.

---

## Shrinking a VDisk using the CLI

This task provides step-by-step instructions for shrinking a VDisk using the CLI.

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

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

### Steps:

Perform the following steps to shrink a VDisk:

1. Validate that the vdisk is not mapped to any host objects. You may need to follow the “Determining the host that a VDisk is mapped to” on page 171 procedure. 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>
```

For more information on this command see *IBM TotalStorage Virtualization Family SAN Volume Controller Command-Line Interface User’s Guide*.

---

## Migrating extents using the CLI

This task provides step-by-step instructions about how to migrate extents to improve performance.

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 VDisks from one MDisk group to another. This can be used to remove hot MDisk groups, for example, reduce the utilization of a group of MDisks.

You can determine the usage of particular MDisks by gathering I/O statistics about MDisks and VDisks. 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:

```
svcinfo lsiostatsdumps
```

This will list the I/O statistics files that have been generated. These are prefixed with **m** 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>scp <clusterip>:/dumps/iostats/m_*
```

This will copy all the MDisk statistics files to the AIX host in the current directory.

4. Analyze the dumps to determine which MDisks are hot. It may be helpful to also determine which VDisks are being heavily utilized as you can spread the data they contain more evenly across all the MDisks in the group using the procedure below.
5. Stop the statistics collection again by issuing the following command:

```
svctask stopstats
```

#### Steps:

Once you have determined which MDisks are hot, you can migrate some of the data onto some less hot MDisks within the same MDisk group.

1. Determine the number of extents that are in use by each VDisk for the given MDisk. Issue the following command:

```
svcinfo lsmdiskextent <mdiskname>
```

This will return the number of extents that each VDisk is using on the given MDisk. You should pick some of these to migrate elsewhere in the group.

2. Determine the other MDisks that reside in the same MDisk group.

- a. To determine the MDisk group that the MDisk belongs to, issue the following command:

```
svcinfo lsmdisk <mdiskname>/ID>
```

Look for the **mdisk\_grp\_name** attribute.

- b. List the MDisks in the group by issuing the following command:

```
svcinfo lsmdisk -filtervalue mdisk_grp_name=<mdiskgrpname>
```

3. Select one of these MDisks as the target MDisk for the extents. You can determine how many free extents exist on an mdisk by issuing the following command:

```
svcinfo lsfreeextents <mdiskname>
```

You can issue the **svcinfo lsmdiskextent <newmdiskname>** command for each of the target MDisks to ensure that you are not just moving the over-utilization to another MDisk. Check that the VDisk that owns the set of extents to be moved, (see step1 on page 186), 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 186. *<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 186 to 4 for all the sets of extents you wish to move.
6. You can check the progress of the migration(s) by issuing the following command:

```
svcinfo lsmigrate
```

---

## Migrating VDisks between MDisk groups using the CLI

This task provides step-by-step instructions for migrating VDisks between MDisk groups.

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

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

1. Add additional MDisks to the target MDisk group. This will provide additional extents in the group and will allow the migrations to be restarted (by marking the error as fixed).
2. Migrate one or more 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).

### Steps:

Perform the following steps to migrate VDisks between MDisk groups:

1. Isolate any VDisks that are overutilized. You can determine this by requesting an I/O statistics dump and analyzing the output. To start I/O statistics gathering, issue the following command:

```
svctask startstats - interval 15
```

2. This will generate a new I/O statistics dump file approximately every 15 minutes. Wait for at least 15 minutes after issuing the **svctask startstats** command and then issue the following command:

```
svcinfo lsiostatsdumps
```

This will list the I/O statistics files that have been generated. These are prefixed with **m** 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>scp <clusterip>:/dumps/iostats/v_*
```

This will copy all the VDisk statistics files to the AIX host in the current directory.

4. Analyze the dumps to determine which 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 using the “Migrating extents using the CLI” on page 185.
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:

```
svcinfo lsmigrate
```

---

## Migrating a VDisk between I/O groups using the CLI

This task provides step-by-step instructions for migrating a VDisk between I/O groups.

**Attention:** These migration tasks are disruptive, in that the cached data held within the cluster must first be written to disk, then the allocation of the VDisk can be changed.

Modifying the I/O group that services the virtual disk cannot be done concurrently with I/O operations. It also requires a rescan at the host level to

ensure that SDD gets notified that the allocation of the preferred node has changed and the ports by which the virtual disk is accessed has changed. This should only be done in the situation where one pair of nodes has become over utilized.

**Steps:**

Perform the following steps to migrate a VDisk between I/O groups:

1. Sync all filesystems that are mounted on the given virtual disk.
2. Stop all I/O operations to the virtual disk.
3. Type the following:

```
svctask chvdisk -iogrp <new_io_grp_name_or_id>  
<vdisk>
```

4. Issue the SDD command to resync the VDisk to host mapping. See the *IBM Subsystem Device Driver (SDD) User's Guide* for more information.
5. Restart the I/O operations to the virtual disk.

---

## Creating image mode virtual disks using the CLI

This task provides step-by-step instructions to convert your virtual disks from image mode to managed mode using the CLI.

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

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

1. Managed disks that contain existing data cannot be differentiated from managed disks that are blank. Therefore, it is vital that you control the introduction of these disks to the cluster. It is recommended that you introduce these disks one at a time. For example, map a single LUN from your RAID controller to the cluster and refresh the view of managed disks. The newly detected disk is displayed.
2. *Do not* add a managed disk that contains existing data to a managed disk group manually. If you do, the data will be lost. When you create an image mode virtual disk from this managed disk, it will be automatically added to the managed disk group. However, it will be added in such a way that the cluster can control how it is added to ensure the data is not lost.

Go to the following Web site for more information:

[www.ibm.com/storage/support/2145](http://www.ibm.com/storage/support/2145)

**Steps:**

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

1. Map a single RAID array or LUN from your RAID controller to the cluster. You can do this either through a switch zoning or a RAID controller based on your host mappings.
2. Rescan the list of managed disks from the SAN Volume Controller Console. Issue the `svcinfo lsmdisk` command to list the available managed disks.

Optionally, if the new managed disk is not listed you may need to run a fabric level discovery. Issue the **svctask detectmdisk** command to manually rescan the fibre-channel network for any new managed disks that might have been added.

3. Convert the managed disk into an image mode virtual disk. Issue the **svctask mkvdisk** command to create an image mode virtual disk object. Once mapped to a host object, these virtual disks are seen as disk drives with which the host can perform I/O operations.
4. Map the new virtual disk to the hosts that were previously using the data that the MDisk contains. Issue the **svctask mkvdiskhostmap** command to create a new mapping between a virtual disk and a host. That is, the virtual disk is made accessible for I/O operations to the specified host.

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

---

## Advanced function FlashCopy and Remote Copy overview for CLI

This topic provides an overview about the advanced function FlashCopy and Remote Copy overview.

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

[www.ibm.com/redbooks](http://www.ibm.com/redbooks)

---

## Advanced function cluster overview using the CLI

This topic provides an overview about advanced functions for your cluster.

### Overview:

The following sections details the advanced cluster functions that you can perform using the CLI.

## Deleting a node from a cluster using the CLI

This task provides step-by-step instructions about how to delete a node from a cluster using the CLI.

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

### Prerequisites:

**Attention:** If you are deleting a single node, and the other node in the I/O group is online, be aware that the cache on the partner node will go into write-through mode and that you are exposed to a single point of failure should the partner node fail.

**Attention:** If you are deleting a node, and this is the last node in the I/O group, you will lose access to all VDisks served by this I/O group. Ensure that all VDisks are not being accessed or contain data that you wish to continue to access, or ensure that they have been migrated to a different (online) I/O group.

1. Begin by determining the VDisks that are still assigned to this I/O group:
  - a. Determine the VDisks in question by requesting a filtered view of VDisks where the filter attribute is the I/O group in question. This can be done using the following command:

```
svcinfo lsvdisk -filtervalue IO_group_name=<name>
```

where <name> is the name of the I/O group in question.
  - b. Once you have a list of VDisks, determine the hosts that they are mapped to by following the procedure called, Determining the hosts that a VDisk is mapped to.
  - c. Once you have determined the hosts and are sure that you do not wish to maintain access to these VDisks proceed to 2.
  - d. If you determine that some or all of the VDisks assigned to this I/O group do contain data that you wish to continue to access, you should follow the procedure called, Migrating a VDisk to a new I/O group.
2. Before deleting the node, it is essential that for each vpath presented by the VDisks you intend to remove, the SDD configuration is updated to remove the vpaths in question. Failure to do this may result in data corruption. See the *IBM TotalStorage: Subsystem Device Driver User's Guide* for details about how to dynamically reconfigure SDD for the given host operating system.
3. Proceed to 1.

#### Steps:

Perform the following steps to delete a node:

1. Open a command prompt window.

#### Notes:

- a. Before removing a node, be sure this is what you want to do. Any VDisks that are assigned to the I/O group that this node belongs to, will be assigned to the other node in the I/O group, that is, the preferred node will be changed. You cannot change this setting back once this has been done. Also, all VDisks will go into write-through cache mode as there is not a redundant node available to duplicate the cached information.
  - b. If this is the last node in the I/O group or the last node in the cluster, you will be asked to force the delete.
  - c. If this is the last node in the cluster or is currently assigned as the configuration node, all connections to the cluster will be lost. The user interface and any open CLI sessions will halt. This may result in a timeout from the command as the command cannot be completed before the node is deleted.
2. Issue the **svctask rmnode** command to delete a node from the cluster. You can enter this command any time after a cluster has been created.

## Enabling the cluster maintenance procedure using the CLI

This task provides step-by-step instructions for enabling the cluster maintenance procedure using the command-line interface (CLI).

#### Steps:

Perform the following steps to enable the maintenance procedure:

1. Open a command prompt window.
2. Issue the **svctask finderr** command to analyze the error log for the highest severity of unfixed errors. This command scans the error log for any unfixed errors. Given a priority ordering defined within the code, the highest priority of unfixed errors is returned.
3. Issue the **svctask dumperlog** command to dump the contents of the error log to a text file. You can also use this command to delete unwanted error log dumps from the cluster.
4. Issue the **svctask clearerrlog** command to clear all entries from the error log including status events and any unfixed errors.  
**Attention:** You should only use this command when you have either rebuilt the cluster, or have fixed a major problem that has caused many entries in the error log that you do not want to manually fix.
5. Issue the **svctask cherrstate** command to mark the state of the error. The state can either be fixed or unfixed.

## Maintaining passwords using the CLI

This task provides step-by-step instructions for maintaining passwords using the command-line interface (CLI).

### Steps:

Perform the following steps to maintain passwords:

1. Open a command prompt window.
2. Issue the **svctask setpwdreset** command to view and change the status of the password-reset feature for the display panel. Passwords can consist of A - Z, a - z, 0 - 9, and underscore. Make a careful note of the admin password, because without it, you cannot access the cluster.

## Maintaining SSH keys using the CLI

This task provides step-by-step instructions for maintaining SSH keys using the command-line interface (CLI).

**Attention:** After you add a cluster, close the Maintaining SSH Keys panel.

### Steps:

Perform the following steps to maintain SSH keys:

1. Open a command prompt window.
2. Issue the **svcinfo lsshkeys** command to list the SSH keys that are available on the cluster.
3. Issue the **svctask addsshkey** 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, type 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.

4. You can issue the **svctask rmsshkey** command to remove an SSH key from the cluster.
5. You can issue the **svctask rmallsshkeys** command to remove all of the SSH keys from the cluster.

## Setting up error notifications using the CLI

This task provides step-by-step instructions for setting up error notifications using the command-line interface.

### Steps:

Perform the following steps to set up error notifications:

1. Open a command prompt window.
2. Issue the **svctask setevent** command to specify what you like to happen when an error or event is logged to the error log. You can select whether the cluster raises an SNMP trap, issues an e-mail notification for entries that are added to the cluster error or event log, or both. Three levels of notification are possible:
  - **None** No error or status changes will be sent.
  - **hardware\_only** You will be notified of errors, but you will not be notified of status changes.
  - **All** You will be notified of all errors and status changes.

If you have an SNMP manager installed or if you want to be notified by e-mail of errors or events, you should enable error notification. The notification levels for SNMP and e-mail alerts can be set independently. If you choose All or **hardware\_only** notification, you must select a destination for the notification.

## Modifying IP addresses using the CLI

This task provides step-by-step instructions for modifying IP addresses using the command-line interface (CLI).

### Steps:

Perform the following steps to modifying IP addresses:

1. Open a command prompt window.
2. Issue the **svcinfo lscluster** command to list the IP address of the cluster.
3. Issue the **svctask chcluster** command to modify the IP address. This command enables you to change the settings of the following:
  - Cluster IP address
  - Service IP address (used when the node is not part of the cluster)
  - Subnet mask
  - Gateway

If you specify a new cluster IP address, the existing communication with the cluster is broken.

## Listing log or dump files using the CLI

This task provides step-by-step instructions for listing log or dump files using the command-line interface (CLI).

### Steps:

Perform the following steps to list log or dump files:

1. Open a command prompt window.
2. You can issue any of the following commands to list error log files:
  - `svcinfo lserrlogbymdisk`
  - `svcinfo lserrlogbymdiskgroup`
  - `svcinfo lserrlogbyvdisk`
  - `svcinfo lserrlogbyhost`
  - `svcinfo lserrlogbynoded`
  - `svcinfo lserrlogbyiogrp`
  - `svcinfo lserrlogbyfcconsistgrp`
  - `svcinfo lserrlogbyfcmap`
  - `svcinfo lserrlogbyrcconsistgrp`
  - `svcinfo lserrlogbyrcrelationship`

These commands will list the error log by type. These commands will return a list of dumps in the appropriate directory. For example, issue the `svcinfo lserrlogbymdisk` command, displays the error log by MDisks.

You can display the whole log or filter the log so that only errors, events, or unfixed errors are displayed. In addition, you can request the output to be sorted either by error priority or by time. For error priority, the most serious errors are the lowest-numbered errors. They are, therefore, displayed first in the table. For time, either the older or the latest entry can be displayed first in the output.

3. You can issue any of the following command to list dump files.
  - `svcinfo lsconfigdumps`
  - `svcinfo lserrlogdumps`
  - `svcinfo lsfeatureddumps`
  - `svcinfo lsiostatsdumps`
  - `svcinfo lsiotracedumps`
  - `svcinfo lssoftwaredumps`
  - `svcinfo ls2145dumps`

These commands will list the dump file by type. These commands will return a list of dumps in the appropriate directory. For example, issue the `svcinfo lsconfigdumps` command, a list of dumps for configurations will be stored in the /dumps/configs destination directory.

The software dump files contain dumps of the SAN Volume Controller memory. Your service representative might ask for these dumps to debug problems. The software dumps are large files (approximately 300 MB). Consider copying these files to your host using secure copy (scp) methods.

## Changing the language setting using the CLI

This task provides step-by-step instructions for changing the language settings.

### Steps:

Perform the following steps to change the language settings:

1. Open a command prompt window.

- Issue the **svcservicetask setlocale** command to change the locale setting for the cluster. It changes all interfaces output to the chosen language. For example, if you wanted to change the English default language to Japanese, type the following:

```
svcservicetask setlocale -locale 3
```

where 3 is the argument that stands for Japanese. The arguments are:

- 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 panel settings.

## Viewing the feature log using the CLI

This task provides step-by-step instructions for viewing the feature log using the command-line interface (CLI).

### Steps:

Perform the following steps to view the feature log:

- Open a command prompt window.
- Issue the **svcinfo lsfeaturedumps** command to return a list of dumps in the /dumps/feature destination directory. The feature log is maintained by the cluster. The feature log records events that are generated when license parameters are entered or when the current license settings have been breached.
- Issue the **svcservicemodeinfo lsfeaturedumps** command to return a list of the files that exist of the type specified on the given node.

## Analyze the error log using the CLI

This task provides step-by-step instructions for analyzing the error log using the command-line interface (CLI).

### Steps:

Perform the following steps to analyze the error log:

- Open a command prompt window.
- You can issue any of the following commands to list error log files:
  - svcinfo lserrlogbymdisk**
  - svcinfo lserrlogbymdiskgroup**
  - svcinfo lserrlogbyvdisk**
  - svcinfo lserrlogbyhost**

- `svcinfo lserrlogbynode`
- `svcinfo lserrlogbyiogrp`
- `svcinfo lserrlogbyfcconsistgrp`
- `svcinfo lserrlogbyfcmap`
- `svcinfo lserrlogbyrcconsistgrp`
- `svcinfo lserrlogbyrcrelationship`

These commands will list the error log by type. These commands will return a list of dumps in the appropriate directory. For example, issue the `svcinfo lserrlogbymdisk` command, displays the error log by MDisks.

You can display the whole log or filter the log so that only errors, events, or unfixed errors are displayed. In addition, you can request the output to be sorted either by error priority or by time. For error priority, the most serious errors are the lowest-numbered errors. They are, therefore, displayed first in the table. For time, either the older or the latest entry can be displayed first in the output.

## Shutting down a cluster using the CLI

This task provides step-by-step instructions for shutting down a cluster using the command-line interface (CLI).

### Prerequisites:

If all input power to a SAN Volume Controller cluster is to be removed for more than a few minutes, (for example, if the machine room power is to be shutdown for maintenance), it is important that the cluster is shutdown before the power is removed. The reason for this is that if the input power is removed from the uninterruptible power supply units without first shutting down the cluster and the uninterruptible power supplies, the uninterruptible power supply units will remain operational and eventually become drained of power.

When input power is restored to the uninterruptible power supplies they will start to recharge but the SAN Volume Controllers will not permit any I/O activity to be performed to the virtual disks until the uninterruptible power supply 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 uninterruptible power supply units will prevent the battery power being drained and will make it possible for I/O activity to be resumed as soon as input power is restored.

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

**Attention:** If you are shutting down the entire cluster, you will lose access to all VDisks being provided by this cluster.

### Shutting down the cluster:

#### Steps:

Perform the following steps to shut down a cluster:

1. Begin the process of quiescing all I/O to the cluster by stopping the applications on your hosts that are using the VDisks provided by the cluster.
  - a. If you are unsure which hosts are using the VDisks provided by the cluster, follow the procedure called, Determining the hosts that a VDisk is mapped to.
  - b. Repeat the previous step for all VDisks.
2. Open a command prompt window.
3. When all I/O has been stopped, issue the **svctask stopcluster** to shut down a single node or the entire cluster in a controller manner. If you specify the node ID or node name, you can shut down a single node.  
 When you enter this command either a node ID or node name argument, the node in question is shut down. After the command completes, the other node in the I/O group destages the contents of its cache and goes into write-through mode until the power to the node is returned and the node rejoins the cluster.  
**Attention:** If this is the last node in an I/O group, you will lose all access to the virtual disks in the I/O group. Before you enter this command, ensure that this is what you want to do. You must specify the force flag.  
 If a shutdown command has been sent to the cluster and both cluster and uninterruptible power supply units have powered off, when input power is restored it will be necessary to restart the uninterruptible power supply units by pressing the power button on the uninterruptible power supply front panel.
4. Close the ssh session if you are using ssh in interactive mode.

#### **Shutting down a single node:**

**Attention:** If you are shutting down a single node, and the other node in the I/O group is online, be aware that the cache on the partner node will go into write-through mode and that you are exposed to a single point of failure should the partner node fail while this node is shut down. Proceed to 2.

**Attention:** If you are shutting down a single node, and this is the last node in the I/O group, you will lose access to all VDisks being served by this I/O group.

#### **Steps:**

Perform the following steps to shut down a single node:

1. Begin the process of quiescing all I/O to the VDisks being served by this nodes I/O group.
  - a. Determine the VDisks in question by requesting a filtered view of VDisks where the filter attribute is the I/O group in question. This can be done using the following command:  
`svcinfo lsvdisk -filtervalue IO_group_name=<name>`

where <name> is the name of the I/O group in question.
  - b. Once you have a list of VDisks, determine the hosts that these are mapped to by following the procedure called, Determining the hosts that a VDisk is mapped to.
2. When all I/O has been stopped issue the following command to shut down the node:  
`svctask stopcluster <nodename/ID>`

where <nodename/ID> is the name or ID of the node that you want to shut down.

**Note:** If this is the last node in the I/O group you also need to specify the -force parameter. For example to force the shutdown of node1:

```
svctask stopcluster -force node1
```

**Related topics:**

- “Determining the host that a VDisk is mapped to” on page 171

---

## Part 5. Software upgrade strategy using the CLI and the SAN Volume Controller Console

This chapter provides information about the software upgrade strategy.

You can upgrade your software while your day-to-day operations are running. You must, however, expect performance to be degraded while the software is being installed.

**Note:** Applying a software update takes approximately one hour. This is in part due to the 30 minute delay which is inserted to allow the 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 is increased each time a new release is made, although only some of the components might have changed. Included in the package are Linux, Apache, and the SAN Volume Controller software.

If you are upgrading through more than one level; for example, from level 1 to level 3, under some circumstances, you might need to install an intermediate level. For example, if you are upgrading from level 1 to level 3, you might need to install level 2 before you install level 3. Details of any prerequisite levels are provided with the source files.

**Attention:** Applying a software upgrade while the node is in service mode results in deleting the node from the cluster. Status information stored within the node will be deleted, and this will cause data loss if the cluster is dependent solely on this node.

**Attention:** Concurrent software upgrade should not be performed while any Remote Copy, FlashCopy or Data Migration operations are active. Quiesce all such operations before starting the software upgrade procedure.

**Attention:** Ensure that you have no unfixed errors in the log. Start the Directed Maintenance Procedures and ensure that you fix any outstanding errors before attempting to concurrently upgrade the software.



---

## Chapter 16. Disruptive software upgrade

This task provides step-by-step instructions about how to perform a disruptive software upgrade using the CLI.

The IBM Total Storage SAN Volume Controller only supports concurrent code upgrades. To ensure that a code upgrade is coordinated across all nodes in the cluster, it is necessary for the nodes to be able to communicate with each other across the fibre-channel SAN. However, some users may prefer to perform a disruptive code upgrade. The following procedure documents how to quiesce I/O to the SAN before performing a concurrent code upgrade to ensure that there is no I/O in progress during the upgrade.

### Steps:

Perform the following steps to complete the disruptive software upgrade process:

1. Stop any host applications and unmount the filesystems that are using storage that is being managed by the SAN Volume Controller. If your hosts are being shutdown then this will occur as the host is shutdown, otherwise it will be necessary to do this manually on each host. This step will ensure that hosts will stop issuing I/O operations and that any data in the filesystem caches is flushed.
2. Shutdown the cluster by issuing the **svctask stopcluster** command. This command will stop the SAN Volume Controllers from issuing I/O to back-end controllers and will flush data from the SAN Volume Controller cache.
3. Re-zone 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 back-end controller (keep the old switch configuration so it can be restored at 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 IBM Total Storage SAN Volume Controller has been isolated from the back-end storage you will get some error logs indicating that this has occurred.

5. Perform the software upgrade in the same manner as for a concurrent code upgrade.
6. Restore the original switch configuration.
7. Clear any error logs produced at step 4 indicating that back-end storage is unavailable. Check that all back-end storage is now online and accessible to the SAN Volume Controllers.
8. Remount filesystems and start host applications.

### Related topics:

- “Shutting down a cluster using the CLI” on page 196
- Chapter 22, “Installing the upgrade using the CLI”, on page 215



---

## Chapter 17. Upgrading software using the SAN Volume Controller Console

This task provides step-by-step instructions about upgrading software using the SAN Volume Controller Console.

Software upgrade files can be quite large, if you experience problems when uploading upgrade files to the cluster reliably, you should disable proxies on the Web browser from where you will upload the file. This should also shorten the file upload time.

**Note:** Note: If you disable proxies, you may not be able to connect to external Web sites. It is therefore advised that prior to disabling proxies, you make a record of your existing settings in case you need to restore access to other Web sites.

### Prerequisites:

If you are using Internet Explorer, perform the following:

1. Click on **Tools** in the menu.
2. Select **Internet Options ->Connections** tab.
3. Click on **LAN Settings...** and ensure that the box marked **Use a proxy server** is unchecked. Click **OK** twice to accept the settings.

If you are using Netscape, perform the following:

1. Click on **Edit** in the menu.
2. Click on **Preferences....** Expand the Advanced section and select **Proxies**.
3. Select the radio button marked **Direct connection to the Internet**. Click **OK** to accept the settings.

### Steps:

Perform the following steps to upgrade the software:

1. Click **Service and Maintenance** from the portfolio.
2. 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.
3. Click **Upload** to copy a new software level from your host to the cluster. (This action uses the upload feature of the Web browser.) The Software upgrade - file upload panel is displayed. You can get new software levels from the IBM Product Support Web site, or from an installation CD.

After a successful copy of the file it is applied as described in Chapter 21, “Secure copy (scp) overview”, on page 213, the install process will fail if all the nodes configured into the cluster are not present. This behavior cannot be overridden using the force flag. If any node configured to be a member of the cluster is not present then in order to upgrade the software the node must either be deleted from the cluster or must be brought online. Furthermore, if a node has been deleted from the cluster such that any IO group has only one member then the software upgrade will also fail. This is because the upgrade

process will result in 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.

Before you begin the software upgrade, make sure that you are aware of the following:

- The code 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 begin executing the new firmware, concurrently with normal cluster activity.
- The procedure to update a single node takes approximately 5 minutes.
- During the update of a node it does not participate in I/O activity in the I/O group. Thus all I/O activity for the virtual disks in the I/O group is directed to the other node in the I/O group by the host multipathing software. During the update of a node the other node in the I/O group will notice that its partner is not participating in the cluster and will as a result attempt to flush the write-back 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 does create a single point of data loss. Should the remaining node in an I/O group experience a failure during a software update of its partner then the only valid copy of dirty data in the write-back cache could be lost.
- All of the nodes connected to one uninterruptible power supply are updated first before any of the nodes connected to the other uninterruptible power supply.
- A 30 minute delay is inserted into the procedure between updating the nodes connected to one uninterruptible power supply and starting to update the nodes on the other uninterruptible power supply. This allows time for the host multipathing software to rediscover paths to the nodes on the first uninterruptible power supply so that when nodes on the second uninterruptible power supply are updated loss of access does not result.
- The update is not committed until all nodes in the cluster have been successfully updated to the new code level. If all nodes successfully re-start with the new code the new version is committed. When this happens, the cluster VPD is updated to reflect the new level of code. After this point downgrade to a package with a lower major number is no longer possible.
- New behaviors or functions in the installed firmware will only be available to be invoked when all member nodes are upgraded and the update is committed.
- Since the software upgrade process takes some time the install command completes as soon as the software package is verified by the cluster. To determine when the upgrade has completed you must either display the software version in the cluster VPD or look for the Software upgrade complete event in the error/event log. If any node fails to re-start with the new code level or fails at any other time during the process the code is backed-off. See Chapter 19, “Automatic upgrade”, on page 209 for more information.
- During a software upgrade the version number of each node is updated when the software has been installed and that node has been restarted. The cluster software version number is updated when the new version of software is committed.
- When code upgrade starts an entry is made in the error or event log and another entry is made when the upgrade completes or fails.

4. Click **Apply upgrade** to display the Applying software upgrade panel. This page enables you to select the upgrade and to apply it to the cluster. This page displays a list of the software levels that you can apply to the cluster. When a new code level is applied, it is automatically installed on all the nodes that are in the cluster.



---

## Chapter 18. Error counts

This topic provides useful information that you will need to know when installing the upgrade.

During the SAN Volume Controller software upgrade, you can expect to see either I/O error counts displayed by the datapath query adapter, or an increase in the number of datapath query device commands if active I/O operations exist between hosts and the SANs. See the *IBM TotalStorage Subsystem Device Driver (SDD) User's Guide* for more information about datapath query commands.

During the software upgrade, each SAN Volume Controller node of a working pair is upgraded sequentially. The SAN Volume Controller node that is being upgraded is temporarily unavailable, and all I/O operations to that SAN Volume Controller fail. As a result, I/O error counts increase. However, failed I/O operations are directed to the other SAN Volume Controller node of the working pair, and applications should not see any I/O failures.



---

## Chapter 19. Automatic upgrade

This topic provides information about upgrading automatically.

New nodes introduced to the cluster normally have software packages downloaded to them from the cluster without any manual intervention. A new node requiring a code version higher than that currently available on the cluster or a node that already contains a code version higher than that on the cluster will not be configured into the cluster. If a node is added to the network that has no code installed, for example because the disk drive has been replaced, or it has such an old code version installed that it cannot advertise itself to the clusters, a re-install of the software is forced by using the Node Rescue procedure.

When new nodes are added to the cluster, the upgrade packages are usually automatically downloaded to them from the SAN Volume Controller cluster. No manual intervention is needed.

If you add a new SAN Volume Controller node that has:

- a code version that is higher than the one that is available on the cluster, or
- that has node that already contains a code version that is higher than the one that is on the cluster

that node is *not* configured into the cluster. It will join the cluster, however the node will be downgraded to the cluster level.

If you add a new SAN Volume Controller node to a network:

- that has no SAN Volume Controller code installed, or
- has such as old version of the code installed that it cannot advertise itself to the clusters

press and hold the left and right navigation buttons on the front panel to force a reinstallation of the software.



---

## **Chapter 20. Automatic recovery from upgrade problems:**

This topic provides information about upgrading automatically.

The cluster will automatically terminate the upgrade process if any of the nodes fail to upgrade to the new software level. In this case, any nodes that have already upgraded to the new software level will downgrade back to the original code level. You should check the error log to determine the reason for the failure before attempting to upgrade the cluster again.



---

## Chapter 21. Secure copy (scp) overview

This topic provides information about using secure copy (scp).

### Overview:

Secure copy (scp) provides a file transfer mechanism for secure shell (SSH) to copy files either between two directories on the SAN Volume Controller configuration node, or between the configuration node and another host. You must have appropriate permissions on the source and destination directories on your respective hosts to be able to use scp. Secure copy is available to you when you install an SSH client on your host system.

The scp interface deliberately limits the permissions to the file systems inside the SAN Volume Controller. If you log on as admin, the only writable file system is the following:

/home/admin/upgrade

and similarly, /home/service/upgrade for the service log in. When copying files to the cluster using scp, /home/admin/upgrade is the default directory destination.

If the cluster is inoperative, the configuration interface is not available.

### Example:

Assume you want to copy a file called svcinfo.trc from the /dumps directory. You want to copy this file from the machine called teststand to your local directory, where you will name the file test.txt.

```
scp admin@teststand:/dumps/svcinfo.trc test.txt
```

Output similar to the following is displayed:

```
svcinfo.trc 100%|*****| 12909 00:00
```

### Example:

Assume you want to copy a file called software\_upgrade.pkg from your local directory to the upgrade directory on the machine called teststand. Issue the following command:

```
scp software_upgrade.pkg admin@teststand:/home/admin/upgrade
```

Output similar to the following is displayed:

```
software_upgrade.pkg 100%|*****| 12909 00:00
```



---

## Chapter 22. Installing the upgrade using the CLI

This topic provides useful information that you will need to know when installing the upgrade.

You can use either secure copy (scp) or the SAN Volume Controller Console to copy the upgrade package to each SAN Volume Controller cluster or issue CLI commands.

If you want to use secure copy, perform the following:

1. Once you have downloaded the software upgrade package, copy the package onto the node where the CLI is running. Issue the following to copy the package:

```
scp filename admin@cluster_address:/home/admin/upgrade
```

where *cluster\_address* is your cluster IP address. You are notified of copy failures by error messages from the CLI and the SAN Volume Controller Console. If there is insufficient space on the cluster to store the software upgrade package then the copy operation will fail. If this occurs, issue the **svctask cleardumps** command to make space for the upgrade package, then repeat the copy operation.

2. After a successful copy of the file, issue the **svcservicetask applysoftware -file filename** command, where *filename* is the name of the file that you copied the software upgrade package too. This command starts the installation of the code. The installation process will fail if a node is not present and if the node is not paired with another node in an I/O group. You can, however, use the -force option to override this restriction if you are prepared to lose access to data during the upgrade.

**Note:** The installation process will *only* fail when some paths between the host systems and the cluster are not available. Data access can be lost temporarily during the upgrading process. You can prevent this if, before you start the installation, you issue a datapath query device on each host system to ensure that all paths are available. See the *IBM TotalStorage Subsystem Device Driver (SDD) User's Guide* for more information about datapath query commands.

**Attention:** The order in which the nodes are upgraded depends on the following:

- The position of the nodes. The code will be transferred to all the nodes in an I/O group.
  - The I/O group ID. The code will be transferred from the lowest I/O group ID that includes nodes on it.
3. To verify that the upgrade was successful, you can perform any one of the following steps:
    - The code level is distributed to all the nodes that are in the cluster. The nodes, in turn, are then restarted. If all the nodes successfully restart with the new code level, the new version is committed and the cluster vital product data (VPD) is updated to new level of code.

- The software upgrade is complete when the cluster verifies the upgrade package. To determine whether the upgrade has completed, you must either display the software version in the cluster VPD, or look for the Software upgrade complete event in the SAN Volume Controller error or event log. If the node does not restart automatically during the upgrade, you should repair or manually delete that node from the cluster to complete the backout process.
- Alternatively, you can also either perform the following steps:
  - a. Issue the **svctask dumperlog** command to dump the contents of the error log to a text file. You can also use this command to delete unwanted error log dumps from the cluster.
  - b. Once you have the contents of the error log dumped into a text file, verify that there were no errors in the text file. If there are no errors, you have successfully upgraded the software and output similar to the following is displayed in the log file:

Upgrade completed successfully
  - c. Issue the **svcinfo lsnodevpd** command for each node. You should see that the software version field has been updated.

**Related topics:**

- Chapter 17, “Upgrading software using the SAN Volume Controller Console”, on page 203

---

## Chapter 23. Accessible CLI commands during the upgrade process

This chapter provides information about the software upgrade.

You can run the upgrade operation concurrently with your day-to-day I/O operations. However, the following SAN Volume Controller commands are *only* the commands that are available while running the software upgrade. The commands listed in this topic are a subset of the commands you can issue during the software upgrade process. All other commands are disabled. Commands that are disabled will fail with a message that indicates that a software upgrade is in progress.

**Note:** These commands are only the prefix to the commands that you will be able to issue.

- `svcinfo lsxxxx`
- `svcinfo lsxxxxcandidate`
- `svcinfo lsxxxxprogress`
- `svcinfo lsxxxxmember`
- `svcinfo lsxxxxextent`
- `svcinfo lsxxxxdumps`
- `svcinfo lserrlogxxxx`
- `svcinfo caterrlog`
- `svcinfo caterrlogbyseqnum`
- `svctask rmnode`

where `xxxx` is the object type that you want to list.

You can perform a software upgrade by issuing CLI commands. For regular software upgrades you can issue the **svcservicetask applysoftware** command. For a service software upgrade, you can issue the service mode software upgrade **svcservicemode task applysoftware** command.



---

## Chapter 24. Manual recovery from software upgrade problems

This task provides step-by-step instructions about how to recover from software upgrade problems.

**Attention:** This procedure causes a loss of *all* data currently configured in the cluster. This is a last resort and should only be done if you have recently backed-up your data.

When a revised version of software is committed, you might not be able to return to a previous software version because some data structures might have been changed such that they cannot be used with the previous software version. Therefore, if you have any problems, you must go forward to a later version of the code. In extreme conditions where you cannot wait for a software update and you need to return to the previous software version, you can use the following procedure.

**Attention:** This procedure, however, causes the total loss of the SAN Volume Controller cluster. This should only be done as a last resort.

### Steps:

Perform the following steps to reset from software upgrade problems:

1. Power-off all but one of the nodes that are in the cluster.
2. Set the powered-on node to the service access mode.
3. Use the service access functions to force the download of the older software package.
4. Repeat the action for each of the failed nodes.
5. From a node that has the new code, create a new cluster.

### Related topics:

- “Resetting a cached SAN Volume Controller cluster SSH host fingerprint on the SAN Volume Controller Console overview” on page 132
- “Resetting an refused SSH key relationship between the SAN Volume Controller Console and the SAN Volume Controller cluster overview” on page 133



---

## **Part 6. Other configurations**

This part provides information about other configurations that you can perform.  
This part provides information about the following configurations:

- Chapter 25, “Configuring disk controllers”, on page 223
- Chapter 26, “Overview about zoning a switch”, on page 233



---

# Chapter 25. Configuring disk controllers

This topic provides overview information about configuring disk controllers.

## Overview:

When configuring back-end storage controllers, it is important to ensure that the storage is configured to provide some kind of redundancy against hard disk failures because when using virtualization a failure of back-end storage can affect a larger amount of storage being presented to the hosts. To provide redundancy, back-end storage should be configured as RAID arrays which 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. The larger the number of disks, the fewer disks are required to provide availability for the same total capacity (1 per array). However, more disks means a longer time is taken to rebuild a replacement disk after a disk failure, and during this period a second disk failure will cause a loss of all array data. More data is affected by a disk failure for a larger number of member disks resulting in reduced performance while rebuilding onto a hot spare and more data being exposed if a second disk fails before the rebuild has completed. The smaller the number of disks, the more likely it is that write operations span an entire stripe (strip size x number of members minus one). In this case, write performance is improved. The number of disk drives required to provide availability may be unacceptable if arrays are too small.

## Notes:

1. If in doubt, arrays with between 6 and 8 member disks is recommended.
2. When creating RAID arrays with mirroring, the number of component disks in each array does not affect redundancy or performance.

The attachment of a given backend controller to a SAN Volume Controller requires that some specific settings be applied to the backend storage, some limitations are also listed for each storage type. There are 2 major steps in this process:

1. Setting the characteristics of the SAN Volume Controller to storage connection(s)
2. Mapping logical unit(s) to these connections such that the SAN Volume Controller can access them.

The following table displays the supported RAID disk controllers:

*Table 14. Supported RAID disk controllers*

Controller	Model
IBM (LSI)	FastT200 FastT500 FastT700 FAStT600 FAStT900
IBM (ESS)	2105-F20 2105-800

#### **Related topics:**

- “Configuring a balanced storage subsystem”
- “Configuring FASST disk controllers for the storage server” on page 228
- “Configuring FASST disk controllers for the storage manager” on page 230
- “Configuring the Enterprise Storage Server (ESS)” on page 230

---

## **Configuring a balanced storage subsystem**

This task provides step-by-step instructions for configuring a balanced storage subsystem.

The virtualization features of the IBM Total Storage SAN Volume Controller enable you to choose how your storage is divided up and presented to hosts. While virtualization provides you with a great deal of flexibility, it also offers the potential for setting up a storage subsystem which can be overloaded. A storage subsystem is overloaded if the quantity of I/O transactions that are being issued by the host systems exceeds the capability of the storage to process those transactions. If a storage subsystem is overloaded, then, at best, it causes delays in the host systems and, at worst, it causes I/O transactions to be timed out in the host which leads to errors being logged by the hosts and I/Os being failed back to applications.

As an extreme example of an overloaded storage subsystem it would be possible to use an IBM Total Storage Volume Controller to virtualize a single RAID array and to divide this storage among sixty-four host systems. Clearly, if all host systems attempt to access this storage at the same time the single RAID array will be overloaded. The following guidelines are provided to help you configure balanced storage subsystems.

#### **Steps:**

Perform the following steps to configure balanced storage subsystems:

1. Calculate the I/O rate for an array. For each RAID array in the storage subsystem use the following table to calculate the approximate number of I/O operations per second that can be processed by the RAID array. Note that the actual number of I/O operations per second that can be processed will vary depending on the location and length of each I/O, whether the I/O is a read or a write operation and on the specifications of the component disks of the RAID array.

*Table 15. Calculate the I/O rate*

Type of RAID Array	Number of component disks in the RAID Array	Approximate I/O rate
RAID-1 (mirrored) arrays	2	150
RAID-3, RAID-4, RAID-5 (striped + parity) arrays	N + 1 parity	150 * N
RAID-10, RAID 0+1, RAID 1+0 (striped + mirrored) arrays	N	150 * N

For example, a RAID-5 array with eight component disks has an approximate I/O rate of  $150 * 7 = 1050$ .

2. Calculate the I/O rate for a managed disk. If there is a one-to-one relationship between back-end arrays and managed disks (this is the recommended

configuration) then the I/O rate for a managed disk is the same as the I/O rate of the corresponding array. If an array is divided up into multiple managed disks then the I/O rate per managed disk is the I/O rate of the array divided by the number of managed disks that are using that array.

3. Calculate the I/O rate for a managed disk group. The I/O rate for a managed disk group is simply the sum of the I/O rates of the managed disk within that group.

For example, a managed disk group contains eight managed disks each of which corresponds to a RAID-1 array. From the table above the I/O rate for each managed disks can be calculated as 300. The I/O rate for the managed disk group is therefore  $300 * 8 = 2400$ .

4. Calculate the impact of FlashCopy relationships. If you are using the FlashCopy feature provided by the IBM Total Storage SAN Volume Controller then you need to consider how much additional I/O will be generated by using this feature as this will reduce the rate at which I/O from host systems that can be processed. When a FlashCopy relationship is copying data any write I/Os from host systems to areas of the source or target virtual disk that have not yet been copied will cause extra I/Os to be generated by the IBM Total Storage SAN Volume Controller to copy the data before the write I/O is actually performed. The effect of using FlashCopy depends on the type of I/O workload being generated by an application:

*Table 16. Calculate the impact of FlashCopy relationships*

Type of application	Impact to I/O rate	Additional Weighting for FlashCopy
Application is doing no I/O	Insignificant impact	0
Application is only reading data	Insignificant impact	0
Application is only issuing random writes	Up to 50 times as much I/O	49
Application is issuing random reads and writes	Up to 15 times as much I/O	14
Application is issuing sequential reads or writes	Up to 2 times as much I/O	1

For each virtual disk that is either the source or target of an active FlashCopy relationship consider the type of application that will be using that virtual disk and record the additional weighting for that virtual disk.

For example, a FlashCopy relationship is being used to provide point in time backups. During the FlashCopy process, a host application generates an I/O workload of random reads and writes to the source virtual disk. A second host application reads the target virtual disk and writes the data to tape to create a backup. The additional weighting for the source virtual disk is 14. The additional weighting for the destination virtual disk is 0.

5. Calculate the I/O rate for virtual disks in a managed disk group. Calculate the number of virtual disks in the managed disk group. Add the additional weighting for each virtual disk that is either the source or a target of an active FlashCopy relationship. Divide the I/O rate of the managed disk group by this number to give an I/O rate per VDisk.

Example 1: A managed disk group has an I/O rate of 2400 and contains 20 virtual disks. There are no FlashCopy relationships. The I/O rate per virtual disk is  $2400 / 20 = 120$ .

Example 2: A managed disk group has an I/O rate of 5000 and contains 20 virtual disks. There are two active FlashCopy relationships which have source virtual disks in this managed disk group. Both source virtual disks are being accessed by applications issuing random reads and write and hence the additional weighting for each of these virtual disks is 14. The I/O rate per virtual disk is  $5000 / (20 + 14 + 14) = 104$ .

6. Determine whether the storage subsystem is overloaded. The figure determined at step 4 on page 225 provides some indication of how many I/O operations per second can be processed by each virtual disk in the managed disk group. If you know how many I/O operations per second your host applications generate you can compare these figures to determine if the system is overloaded. If you do not know how many I/O operations per second your host applications generate then you can either measure this (for example by using the I/O statistics facilities provided by the IBM Total Storage SAN Volume Controller to measure the I/O rate of your virtual disks) or use the following table as a guideline:

*Table 17. Determine if the storage subsystem is overloaded*

Type of Application	I/O rate per virtual disk
Applications that generate a high I/O workload	200
Applications that generate a medium I/O workload	80
Applications that generate a low I/O workload	10

7. Interpret the result. If the I/O rate generated by the application exceeds the I/O rate you calculated per virtual disk this indicates that you could be overloading your storage subsystem and you should monitor the system carefully to see if the back-end storage is actually limiting the overall performance of your system. It is also possible that the calculation above is too simplistic to model your use of storage, for example the calculation assumes that your applications generate the same I/O workload to all virtual disks.

One method you can use to monitor the performance of your storage subsystem is to use the I/O statistics facilities provided by the IBM Total Storage SAN Volume Controller to measure the I/O rate of your managed disks. Alternatively you could use the performance and I/O statistics facilities provided by your back-end controllers.

If you find your storage subsystem is over loaded there are several actions that can be take to resolve the problem:

- a. Adding more back-end storage to the system will allow you to increase the quantity of I/O that can be processed by your storage subsystem. The virtualization and data migration facilities provided by the IBM Total Storage SAN Volume Controller can be used to redistribute the I/O work load of virtual disks across a greater number of managed disks without having to take the storage offline.
- b. Stop any unessential FlashCopy relationships as this will reduce the amount of I/O operations submitted to the back-end storage. If you are making many FlashCopy's in parallel then consider starting less FlashCopy relationships in parallel.
- c. The I/O workload generated by a host can often be limited by adjusting the queue depth (for example, the maximum number of I/O operations that are submitted in parallel). Depending on the type of host and type of host bus adapters it may be possible to limit the queue depth per virtual disk and/or

limit the queue depth per host bus adapter. An alternative method of limiting the I/O workload generated by a host would be to use the I/O governing features provided by the IBM Total Storage SAN Volume Controller. These techniques may be particularly applicable if using a mixture of different host systems to prevent one host system from saturating an I/O subsystem to the detriment of the other host systems. Note that although these techniques may be used to avoid I/O time-outs it still means the performance of your system is being limited by the amount of storage.

---

## Data migration on an existing FAStT installation which contains partitions

This topic provides information about data migration on an existing FAStT installation which contains partitions.

You can enable the SAN Volume Controller to be introduced to an existing SAN environment, so that you have the option of utilizing image mode LUNs to import the existing data into the virtualization environment without requiring a backup and restore cycle. For example, each FAStT partition may contain up to 32 LUNs. Each partition can only access a unique set of HBA ports (as defined by WWPNs). That is, for a single host to access multiple partitions, unique host fibre ports (WWPNs) need to be assigned to each partition. All LUNs within a partition are surfaced to assigned host fibre ports (no sub-partition LUN mapping).

Host A has visibility to LUN 0, 1, 2 in Partition 0

Host B has visibility to LUN 0, 1, 2, 3, 4, 5 in  
Partition 1

Host C has visibility to LUN 0, 1, 2 in Partition 2

To allow Host A to access the LUNs in partition B, it is necessary to 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 save and restore cycles would require that a set of unique SAN Volume Controller HBA port WWPNs for each partition. This would allow the FAStT to surface the LUNs (with your data) to the SAN Volume Controller, which would then configure these LUNs as image-mode LUNs and surface them to the required hosts. Unfortunately, this violates a requirement that all SAN Volume Controller nodes be able to see all backend storage. To work around this problem, change the FAStT to allow more than 32 LUNs in 1 storage partition, so you can move all the LUNs from all the other partitions into 1 partition and map to the SAN Volume Controller cluster.

For example, lets say the FAStT has 8 partitions with 30 LUNs in each, and all need to be migrated to a 4-node SAN Volume Controller cluster with 4 ports on each SAN Volume Controller. Perform the following:

1. Change the mappings for the first 4 partitions on the FAStT such that each partition is mapped to 1 port on each node, this maintains redundancy across the cluster.
2. Create a new partition on the FAStT that is mapped to all 4 ports on all the SAN Volume Controllers (actually not a partition at all)
3. Gradually migrate the data into the MDisks in the target partition, as storage is freed from the source partitions this can be reused as new storage in the target partition. As partitions are deleted new partitions that need to be migrated can

be mapped and migrated in the same way. The host side data access and integrity would be maintained throughout this process.

## Configuring FAStT disk controllers for the storage server

This task provides a list of the support actions when configuring the FAStT disk controllers.

**Attention:** The SAN Volume Controller does not concurrently support I/O operations with the download of ESM (Environmental Services Monitor) firmware. You must quiesce all I/O operations from the hosts that are using storage provided by the FAStT controllers you wish to update before installing new ESM firmware. The FAStT storage server has many options and actions. The following list the supported actions and its impact on the SAN Volume Controller and its configuration.

1. host type:

- a. You must set either the default host type of your FAStT or the host type of the chosen partition to:

IBM TS SAN VCE

You can set the host type in 2 ways:

- 1) Click **Storage Subsystem -> Change -> Default Host Type**, or
- 2) For each host port you can specify the host type of that port or modify existing ports.

- b. The IBM TS SAN VCE setting above was added to the FAStT NVRAM image for the SAN Volume Controller connections. For older versions of Microcode or for users who do not have the latest NVRAM image, the Windows Non-Clustered (SP5 or higher) settings will work.

2. WWNN:

- a. Set the subsystem so that both controllers have the same WWNN. Scripts are available from the FAStT support Web site to change the set up of the FAStT if required.

[www.storage.ibm.com](http://www.storage.ibm.com)

3. auto volume transfer (AVT):

- a. Make sure the auto volume transfer is enabled. The host type selection should have enabled this function already.
- b. View the storage subsystem profile data to confirm that you have the AVT function enabled. This storage profile is presented as a text view in a separate window.
- c. Scripts are available from the FAStT Web site to enable AVT if required.

[www.storage.ibm.com](http://www.storage.ibm.com)

4. limitations:

- a. Only one FAStT storage partition can be created that contains any of the ports of any of the nodes in a single SAN Volume Controller cluster.
- b. You must not map more than one partition to any of the ports on any of the nodes in a SAN Volume Controller cluster. Otherwise, unexpected behavior might result. For example, there will not be any warning messages, however, there will be errors logged in the SAN Volume Controller error log and access to storage may be lost.

5. access LUN:
  - a. The access LUN function, also known as the Universal Transport Mechanism (UTM) LUN, might not be in a partition that contains the SAN Volume Controller ports. It is not required by the SAN Volume Controller. The UTM LUN is a special LUN that allows the SAN Volume Controller to be configured through suitable software over the Fibre channel connection. However, the SAN Volume Controller does not require the UTM LUN, therefore does not generate errors either way.
  - b. The FAStT *must not* have the Access (UTM) LUN presented as Logical Unit Number 0 (zero).
6. logical unit:
  - a. The SAN Volume Controller attempts to follow the FAStT specified preferred ownership. You can specify which controller (A or B) is used to do I/O operations to a given Logical Unit. If the SAN Volume Controller can see the ports of the preferred controller and no error conditions exist, then it will access that Logical Unit through one of the ports on that controller.
  - b. Under error conditions, the ownership is ignored. Meaning, the SAN Volume Controller has found a given path through the fabric to be errant, or there is no connection to a given port.
7. limited support:
  - a. Make sure you are aware of the controller mode.
  - b. A subsystem in Redundant Disk Array Controller (RDAC) mode or controllers can be set in Active/Passive or Active/Offline mode. If the subsystem is set in one of these modes, the SAN Volume Controller cannot do failover operations. Redundancy and availability might be reduced under these conditions.
8. copy services (FlashCopy and Remote copy):
  - a. FAStT copy services *must not* be used when the SAN Volume Controller is attached to the FAStT. Partitioning might allow copy services to be used on other host platforms.
9. subsystem identification:
  - a. The serial number presented by the command-line and Web application on the SAN Volume Controller is the serial number of the controllers. These serial numbers can be viewed on the FAStT Web application by selecting a controller and clicking **Properties**.  
If the serial numbers of the controller are not displayed, the WWNN or WWPN will be displayed. The WWNN or WWPN can be used to identify the different controllers.
10. cache:
  - a. Ensure that you have the following enabled on any logical units mapped to the SAN Volume Controller:
    - read caching
    - write caching
    - write cache mirroring
 Caching without batteries must not be enabled.

**Related topics:**

- Chapter 25, “Configuring disk controllers”, on page 223
- “Configuring a balanced storage subsystem” on page 224
- “Configuring FAStT disk controllers for the storage manager” on page 230

- “Configuring the Enterprise Storage Server (ESS)”

---

## Configuring FAStT disk controllers for the storage manager

This task provides a list of the support actions when configuring the FAStT disk controllers.

The FAStT storage manager has many options and actions. The following shows the supported actions and their impact on the SAN Volume Controller and its configuration.

1. controller run diagnostics:
  - a. The diagnostics should be automatically recovered by the SAN Volume Controller software. The diagnostics main function is to maintain redundancy and availability.
  - b. Check your MDisks to make sure that they have not been set to degraded mode after this action.
2. controller disable data transfer:
  - a. This option is not supported when a SAN Volume Controller is attached to the FAStT. Loss of availability and redundancy may occur if data transfer is disabled.
3. setting an array Offline:
  - a. Do not set an array Offline. If you use this setting, you might lose access to the MDisk group.
4. array increase capacity:
  - a. Increasing capacity is supported but the new capacity is not usable until the MDisk is removed from an MDisk group and then added again. You might have to migrate data to increase the capacity.
5. redistribute logical drives or change ownership of the preferred path:
  - a. These actions are supported but might not take effect until a cluster rediscovery is initiated on the SAN Volume Controller cluster. This can be achieved using the **svctask detectmdisk** command.
6. controller reset
  - a. Controller reset should only be performed if directed to do so by service personnel. The alternate controller is functional and available to the SAN Volume Controller reset should be automatically recovered by the SAN Volume Controller software.
  - b. Check your MDisks to make sure that they have not been set to degraded state during this operation. You can issue the **svctask includemdsk** to repair degraded MDisks.

### Related topics:

- Chapter 25, “Configuring disk controllers”, on page 223
- “Configuring a balanced storage subsystem” on page 224
- “Configuring FAStT disk controllers for the storage server” on page 228
- “Configuring the Enterprise Storage Server (ESS)”

---

## Configuring the Enterprise Storage Server (ESS)

This task provides step-by-step instructions for configuring the ESS.

**Attention:** If you have set your Host Type on your ESS (for the defined SVC Cluster host ports) to any host type other than:

- RS/6000 or
- SAN Volume Controller

You *must not* subsequently change them. If you do you may permanently lose access to your data.

**Steps:**

Perform the following steps to configure the ESS:

1. Click **Storage Allocation**.
2. Click **Open System Storage**.
3. Click **Modify Host Systems**.
4. Create a host entry for every port on every SAN Volume Controller node in your cluster. Complete the following fields:

**Nickname**

Type a unique name for each port (for example, knode or lnode).

**Host Type**

Select **IBM SAN Volume Controller** or **RS/6000** if that is not available.

**Host Attachment**

Select **Fibre Channel attached**.

**Hostname/IP address**

Leave this field blank.

**WWPN**

Either select the WWPN from the list, or type it manually. A configuration command will fail if you use WWPN 0 in the command string.

5. After you are finished adding all of the ports, click **Perform Configuration Update**.
6. Click **Add Volumes** to add the volumes on which you want the SAN Volume Controller to run.
7. From the Add Volumes window, perform the following actions:
  - 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. After you are done creating all the volumes, click **Perform Configuration Update**.
8. Map the volumes to all of your SAN Volume Controller ports by performing the following steps:
  - a. Click **Modify Volume Assignments**.
  - b. Select all of the volumes that you created earlier.
  - c. Click **Assigning selected volumes to target hosts**.
  - d. Select all of the remaining SAN Volume Controller host ports that you created earlier.
  - e. Select the **Use same ID/LUN in source and target** check box.
  - f. Click **Perform Configuration Update**.

**Related topics:**

- Chapter 25, “Configuring disk controllers”, on page 223
- “Configuring a balanced storage subsystem” on page 224
- “Configuring FAStT disk controllers for the storage server” on page 228
- “Configuring FAStT disk controllers for the storage manager” on page 230

---

## Chapter 26. Overview about zoning a switch

This topic provides information about zoning a switch.

### Overview:

The number of virtual paths of each virtual disk (VDisk) by zoning a switch is limited. Implementation of the following rules will help you achieve the correct number of virtual paths.

- Each host (or partition of a host) may have between 1 and 4 fibre channel ports.
- Switch zoning should be used to ensure that each host fibre channel port is zoned to exactly 1 fibre channel port for each SAN Volume Controller node in a cluster.
- To obtain the best performance from a host with multiple fibre channel ports, the zoning should ensure that each fibre channel port of a host is zoned with a different group of SAN Volume Controller ports.
- To obtain the best overall performance of the subsystem, the workload for each SAN Volume Controller port should be equal. This will typically involve zoning roughly the same number of host fibre channel ports to each SAN Volume Controller fibre channel port.

We recommend manually setting the domain IDs prior to building the multiswitch fabric and prior to zoning. One reason is, that when two switches are joined while active, they will determine if the Domain ID is already in use as before, but if there is a conflict it cannot be changed in an active switch. This conflict will cause the fabric merging process to fail. The second reason is that the domain ID is used to identify switch ports when zoning is implemented using the domain and switch port number. If domain ID's are negotiated at every fabric start up, there is no guarantee that the same switch will have the same ID next time, therefore any zoning definitions may become invalid. Additionally, if the domain ID is changed after a SAN is set up, some host systems may have difficulty logging back in with the switch and a host reconfiguration may be required to detect devices on the switch again.

### Related topics:

- “Fibre channel switches” on page 44
- “Switch operations over long distances” on page 238

---

## Zoning a switch

This topic provides two examples of valid configurations for zoning switches.

### Example:

The number of paths from the SAN Volume Controller nodes to a host must not exceed eight. The maximum number of host bus adapter (HBA) ports must not exceed four (for example, no more than two two-port HBAs or four one-port HBAs).

In this example, your SAN environment looks similar to the following:

- 2 SAN Volume Controller nodes called A and B

- nodes A and B have 4 ports each
  1. Node A has ports A0, A1, A2, and A3
  2. Node B has ports B0, B1, B2, and B3
- 4 hosts called P, Q, R, and S
- the 4 hosts have 4 ports each called:

*Table 18. Four hosts and their ports*

P	Q	R	S
C0	D0	E0	F0
C1	D1	E1	F1
C2	D2	E2	F2
C3	D3	E3	F3

- 2 switches called X and Y
- 1 storage controller
- storage controller has 4 ports on it called I0, I1, I2, and I3

An example configuration would be the following:

1. Attach ports 1 (A0, B0, C0, D0, E0, and F0) and 2 (A1, B1, C1, D1, E1, and F1) of each node and host to switch X.
2. Attach ports 3 (A2, B2, C2, D2, E2, and F2) and 4 (A3, B3, C3, D3, E3, and F3) of each node and host to switch Y.
3. Attach ports 1 and 2 (I0 and I1) of the storage controller to switch X.
4. Attach ports 3 and 4 (I2 and I3) of the storage controller to switch Y.

On switch X we would create the following:

5. Create a host zone containing ports 1 (A0, B0, C0, D0, E0, and F0) of each node and host.
6. Create a host zone containing ports 2 (A1, B1, C1, D1, E1, and F1) of each node and host.

Similarly, we would create the following:

7. Create two host zones on switch Y containing ports 3 (A2, B2, C2, D2, E2, and F2) of each node and host.
8. Create two host zones on switch Y containing ports 4 (A3, B3, C3, D3, E3, and F3) of each node and host.

Lastly, we would create the following:

9. Create a storage zone, which would be configured on each switch.  
Each switch contains all the SAN Volume Controller storage ports on that switch.

### **Example:**

In this second example, your SAN environment looks similar to the first example, however, here we have an additional 2 hosts with 2 ports each:

- 2 SAN Volume Controller nodes called A and B
- nodes A and B have 4 ports each
  1. Node A has ports A0, A1, A2, and A3
  2. Node B has ports B0, B1, B2, and B3
- 6 hosts called P, Q, R, S, T and U
- the 4 hosts have 4 ports each and the two additional hosts have 2 ports each called:

*Table 19. Four hosts and their ports*

P	Q	R	S	T	U
C0	D0	E0	F0	G0	H0
C1	D1	E1	F1	G1	H1
C2	D2	E2	F2	—	—
C3	D3	E3	F3	—	—

- 2 switches called X and Y
- 1 storage controller
- storage controller has 4 ports on it called I0, I1, I2, and I3

An example configuration would be the following:

1. Attach ports 1 (A0, B0, C0, D0, E0, F0, G0, and H0) and 2 (A1, B1, C1, D1, E1, and F1) of each node and host to switch X.
2. Attach ports 3 (A2, B2, C2, D2, E2, F2, G1 and H1) and 4 (A3, B3, C3, D3, E3, and F3) 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 (G0 and H0) and (G1 and H1) are zoned to different SAN Volume Controller ports so that each SAN Volume Controller port is zoned to the same number of host ports.

On switch X we would create the following:

5. Create a host zone containing ports 1 (A0, B0, C0, D0, E0, F0, G0, and H0) of each node and host.
6. Create a host zone containing ports 2 (A1, B1, C1, D1, E1, F1, G1, and H1) of each node and host.

Similarly, we would create the following:

7. Create two host zones on switch Y containing ports 3 (A2, B2, C2, D2, E2, and F2) of each node and host.
8. Create two host zones on switch Y containing ports 4 (A3, B3, C3, D3, E3, and F3) of each node and host.

Lastly, we would create the following:

9. Create a storage zone, which would be configured on each switch. Each switch contains all the SAN Volume Controller storage ports on that switch.

#### Related topics:

- “Fibre channel switches” on page 44
- Chapter 26, “Overview about zoning a switch”, on page 233

---

## Zoning considerations for Remote Copy

This topic provides information about different zoning considerations for Remote Copy.

SAN configurations that use the Remote Copy feature between two clusters need additional zoning considerations. These considerations include:

- Additional zones for remote copy. For Remote Copy operations involving two clusters, these clusters must be zoned so that the nodes in each cluster can see the ports of the nodes in the other cluster.
- Use of extended fabric settings in a switched fabric.

- Use of ISL trunking in a switched fabric.
- Use of redundant fabrics.

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

When two clusters are used to participate in a Remote Copy partnership, the following two rules apply:

- The clusters must be within the same SAN fabric or within the same set of redundant SAN fabrics.
- An additional zone must be provided in each SAN fabric which includes all ports of all nodes in both clusters.

The following zones would therefore be needed in a typical intercluster Remote Copy configuration:

1. A zone in the local cluster, that contains all the ports in the SAN Volume Controller nodes in that local cluster, and the ports on the backend storage associated with that local cluster. These zones would be required whether or not Remote Copy was in use.
2. A zone in the remote cluster, that contains all the ports in the SAN Volume Controller nodes in that remote cluster, and the ports on the backend storage associated with that remote cluster. These zones would be required whether or not Remote Copy was in use.
3. A zone that contains all the ports in the SAN Volume Controller nodes in both the local and remote cluster. This zone is required for intercluster communication and is specifically required by Remote Copy.
4. Additional zones that contain ports in host HBAs and selected ports on the SAN Volume Controller nodes in a particular cluster. These are the zones that allow a host to see VDisks presented by an I/O group in a particular cluster. These zones would be required whether or not Remote Copy was in use.

**Notes:**

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

## Zoning rules and example

This topic provides detailed information about zoning rules. It also provides an example.

For intra-cluster Remote Copy relationships no additional switch zones are required. For inter cluster Remote Copy relationships, you must:

1. Form a SAN that contains both clusters which are to be used in the Remote Copy relationships. If cluster A is in SAN A originally, and cluster B is in SAN B originally, this means that there must be at least one fibre-channel connection between SAN A and SAN B. This connection will be one or more inter switch links. The fibre-channel switch ports associated with these inter switch ports should not appear in any zone.
2. A single SAN can only be formed out of combining SAN A and SAN B if the domain numbers of the switches in each SAN are different, prior to the connection of the two SANs. You should ensure that each switch has a different domain ID before connecting the two SANs.

3. Once the switches in SAN A and SAN B are connected, they should be configured to operate as a single group of switches. Each cluster should retain the same set of zones that were required to operate in the original single SAN configuration.
4. A new zone must be added that contains all the switch ports that are connected to SAN Volume Controller ports. This will contain switch ports that were originally in SAN A and in SAN B.
5. You can adjust the switch zoning so that the hosts that were originally in SAN A can see cluster B. This allows a host to examine data in both the local and remote cluster if required. This view of both clusters is purely optional and in some cases may complicate the way you operate the overall system, therefore, unless specifically needed, it should not be implemented.
6. You should verify that the switch zoning is such that cluster A cannot see any of the back-end storage owned by cluster B. Two clusters may not share the same back-end storage devices.

**Example:**

You have two clusters. Each cluster is made up of the following:

- one host, containing 2 single port HBAs
- two SAN Volume Controller nodes, each with 4 ports
- one back-end storage device with 2 ports
- one 16 port switch

We will call the clusters A and B. In the A cluster we will call the first SAN Volume Controller ports: AN11, AN12, AN13, AN14. We will call the second SAN Volume Controller ports: AN21, AN22, AN23, AN24.

Similarly we will use AF1 and AF2 for the back-end storage ports and AH1 and AH2 for the host ports.

The zones required in this simple cluster would be:

*Table 20. Cluster A set up*

Zone 1	AH1, AN11, AN21	host zone, contains one host HBA and one port per SAN Volume Controller node
Zone 2	AH2, AN12, AN22	host zone, contains one host HBA and one port per SAN Volume Controller node
Zone 3	AF1, AF2, AN11, AN12, AN13, AN14, AN21, AN22, AN23, AN24	back-end zone, contains all SAN Volume Controller nodes and the ports for the back-end storage

Here, we are naming the switch ports in each zone by the device port that they are connected to. The same zones, notated by the actual switch port numbers would be:

*Table 21.*

Zone 1	1, 3, 7	host zone, contains one host HBA and one port per SAN Volume Controller node
--------	---------	--

*Table 21. (continued)*

Zone 2	2, 4, 8	host zone, contains one host HBA and one port per SAN Volume Controller node
Zone 3	11, 12, 3, 4, 5, 6, 7, 8, 9, 10	back-end zone, contains all SAN Volume Controller nodes and the ports for the back-end storage

Lets assume that the B cluster has a similar set up.

See Table 20 on page 237 for cluster A set up. Cluster B looks like the following:

*Table 22. Cluster B set up*

Zone 4	BH1, BN11, BN21	host zone, contains one host HBA and one port per SAN Volume Controller node
Zone 5	BH2, BN12, BN22	host zone, contains one host HBA and one port per SAN Volume Controller node
Zone 6	BF1, BF2, BN11, BN12, BN13, BN14, BN21, BN22, BN23, BN24	back-end zone, contains all SAN Volume Controller nodes and the ports for the back-end storage

For the SAN Volume Controller Remote Copy zoning, see the following:

*Table 23. SAN Volume Controller Remote Copy zoning*

Zone 7	AN11, AN12, AN13, AN14 AN21, AN22, AN23, AN24 BN11, BN12, BN13, BN14 BN21, BN22, BN23, BN24	This zone contains all the switch ports on both switches that are connected to all SAN Volume Controller ports.
--------	---	---

Again, we are naming the switch ports in each zone by the device port that they are connected to. Since we now have two switches in the connected SAN, each switch must have a different domain numbers, (we will assume switch A is domain 1 and switch B is domain 2). The zones can then be expressed in switch port numbers provided the domain is included in each port description. For example, zone 7 above would be:

*Table 24. SAN Volume Controller Remote Copy zoning*

Zone 7	1,3 1,4 1,5 1,6 1,7 1,8 1 9 1,10 2,3 2,4 2,5 2,6 2,7 2,8 2,9 2,10	This zone contains all the switch ports on both switches that are connected to all SAN Volume Controller ports. Here X,X is domain X, port X.
--------	--	---

---

## Switch operations over long distances

This topic provides information about 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 Remote Copy performance. The two most significant features are ISL trunking and extended fabric.

### **ISL trunking**

Trunking enables the switch to use two links in parallel and still maintain frame ordering. It does this by routing all traffic for a given destination over the same route even when there may be more than one route available. Often trunking is limited to certain ports or port groups within a switch. For example, in the IBM 2109-F16 switch, trunking can only be enabled between ports in the same quad (for example, same group of four ports).

Some switch types may impose limitations on concurrent use of trunking and extended fabric operation. For example, with the IBM 2109-F16 switch, it is not possible to enable extended fabric for two ports in the same quad. Thus, extended fabric and trunking are effectively mutually exclusive. (Although it is possible, to enable extended fabric operation one link of a trunked pair this does not offer any performance advantages and adds complexity to the configuration setup. This mixed mode of operation is therefore not recommended.)

### **Extended fabric**

Extended fabric operation allocates extra buffer credits to a port. This is important over long links usually found in inter-cluster remote copy operation because, due to the time it takes for a frame to traverse the link, it is possible to have more frames in transmission at any instant in time than would be possible over a short link. The additional buffering is required to allow for the extra frames.

For example, the default license for the IBM 2109-F16 switch has two extended fabric options, Normal and Extended Normal.

- Normal is suitable for short links and Extended Normal is suitable for links up to 10km long. (With the additional Extended fabric license the user gets two extra options, Medium, up to 10-50km and Long, 50-100km.)
- The Extended Normal setting gives significantly better performance for the links up to 10 km long. Medium and Long settings are not recommended for use in the inter-cluster remote copy links currently supported.



---

## Appendix A. Installing the IBM TotalStorage SAN Volume Controller Console for Windows

This chapter includes an overview of the installation process and instructions for installing and configuring the IBM TotalStorage SAN Volume Controller Console on a Windows 2000 Server operating system.

**Note:** Installing the SAN Volume Controller Console on your host system is optional. The SAN Volume Controller comes preinstalled on the master console.

---

### Installation overview for the SAN Volume Controller Console

This section provides an overview of the installation and configuration of the IBM TotalStorage SAN Volume Controller Console on a Windows 2000 Server operating system. You should have some knowledge of how to administer a Windows 2000 Server operating system before you install the IBM TotalStorage SAN Volume Controller Console. You should also become familiar with the command that you use during installation and configuration of the IBM TotalStorage SAN Volume Controller Console.

You must be aware of the following list of installation and configuration tasks *before* you install the SAN Volume Controller Console:

1. Before you install the IBM TotalStorage SAN Volume Controller Console, you should check the hardware and software requirements.
2. If the SSH client software called PuTTY is not yet installed on your system, you must install the SSH client software. You can get more information about PuTTY from the PuTTY Web site home page:

<http://www.chiark.greenend.org.uk/~sgtatham/putty/>

and download PuTTY from the following Web site download page:

<http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>

For your convenience the PuTTY installation program (putty-0.53b-installer.exe) is on the SAN Volume Controller Console installation CD-ROM in the SSHClient/PuTTY directory.

3. You can choose to install the IBM TotalStorage SAN Volume Controller Console either in graphical mode with the help of an installation wizard or in unattended mode (also known as silent mode), which involves customizing a response file and issuing a command.
4. Verify the Windows services associated with the SAN Volume Controller Console.
5. Get started using the SAN Volume Controller Console. Use a Web browser to access the SAN Volume Controller Console. You will identify the clusters to be managed to the SAN Volume Controller Console as well as complete the creation (initialization) of the SAN Volume Controller clusters.
6. Remove the IBM TotalStorage SAN Volume Controller Console. You only need to perform this optional task if you get errors during installation verification.

**Related topics:**

- “SAN Volume Controller Console hardware installation requirements”
- “SAN Volume Controller Console workstation space requirements”
- “SAN Volume Controller Console software installation requirements” on page 243
- “Installing the SAN Volume Controller Console in unattended (silent) mode” on page 253
- “Verifying the Windows services associated with the SAN Volume Controller Console” on page 256
- “Generating an SSH key pair using the SSH client called PuTTY” on page 60
- “Post Installation Tasks - Getting started using the SAN Volume Controller Console” on page 257
- “Removing the SAN Volume Controller Console” on page 259

---

## SAN Volume Controller Console hardware installation requirements

Ensure that your system satisfies the following hardware installation prerequisites for installing the IBM TotalStorage SAN Volume Controller Console on a Windows 2000 Server operating system before starting the installation.

**Hardware prerequisites:**

The following hardware is required:

- Any Intel®-based PC running Windows 2000 Server SP 3
- Intel Pentium® processor at 1 GHz, or faster
- Support for a communications adapter
- CD-ROM drive
- Minimum 1 GB RAM recommended

**Related topics:**

- “SAN Volume Controller Console workstation space requirements”
- “SAN Volume Controller Console software installation requirements” on page 243

---

## SAN Volume Controller Console workstation space requirements

Ensure that your system satisfies the following workstation space prerequisites for installing the IBM TotalStorage SAN Volume Controller Console on a Windows 2000 Server operating system before starting the installation.

**Workstation space:**

The following space on your workstation is required:

- 350 MB of disk space

**Note:** You might need to increase the total available disk space on your hard drives if the IBM TotalStorage SAN Volume Controller Console and other associated products are split between more than one logical drive. Also, the IBM TotalStorage SAN Volume Controller Console might require additional memory to operate if you configure it to manage many devices or devices with large configurations.

- Up to 65 MB of temporary disk space for installation purposes

**Related topics:**

- “SAN Volume Controller Console hardware installation requirements” on page 242
- “SAN Volume Controller Console software installation requirements”

---

## SAN Volume Controller Console software installation requirements

Ensure that your system satisfies the following software installation prerequisites for installing the IBM TotalStorage SAN Volume Controller Console on a Windows 2000 Server operating system before starting the installation.

**Software:**

The following software is required:

- Operating systems:
  - Windows 2000 Server SP3
- If the SSH client software called PuTTY is not yet installed on your system, you must install the SSH client software. You can get more information about PuTTY from the PuTTY Web site home page:

<http://www.chiark.greenend.org.uk/~sgtatham/putty/>

and download PuTTY from the following Web site download page:

<http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>

For your convenience the PuTTY installation program (putty-o.53b-installer.exe) is on the SAN Volume Controller Console installation CD-ROM in the SSHClient/PuTTY directory.

- IBM TotalStorage SAN Volume Controller Console. This is on the IBM TotalStorage SAN Volume Controller Console CD.
- Transmission Control Protocol/Internet Protocol (TCP/IP)
- Adobe Acrobat Reader version 4.0 or later (optional)

You need the Adobe Acrobat Reader to read License Agreement and product information from the SAN Volume Controller Console LaunchPad. You can download the Adobe Acrobat Reader from the following Web site:

– <http://www.adobe.com/support/downloads/main.html>

**Related topics:**

- “SAN Volume Controller Console hardware installation requirements” on page 242
- “SAN Volume Controller Console workstation space requirements” on page 242

---

## Installing the SAN Volume Controller Console in graphical mode

This section includes the steps to install the IBM TotalStorage SAN Volume Controller Console in your Windows system. If you choose to install the IBM TotalStorage SAN Volume Controller Console in unattended mode skip this section and follow the instructions in “Installing the SAN Volume Controller Console in unattended (silent) mode” on page 253. You must satisfy all prerequisites before starting the installation.

### Steps:

Perform the following steps to install the IBM TotalStorage SAN Volume Controller Console:

1. Log onto your system as a local system administrator.
2. Insert the IBM TotalStorage SAN Volume Controller Console CD into the CD drive.

The IBM TotalStorage SAN Volume Controller Console program should start within 15 - 30 seconds if you have **autorun** mode set on your system. If the LaunchPad panel does not open, perform either one of the following steps:

- a. Use a Command Prompt to change to the W2K directory on the CD. Type:  
LaunchPad
- b. Using Windows Explorer, (**Start->Programs->Accessories->Windows Explorer**), go to the W2K directory located on the CD drive. Then double-click on the **LaunchPad.bat** file.

**Note:** If you are viewing the folder using the Windows Explorer with the option selected to *Hide file extensions for known file types*, find the LaunchPad file with the file type of MS-DOS Batch File.

3. The following options are displayed when the LaunchPad panel opens:

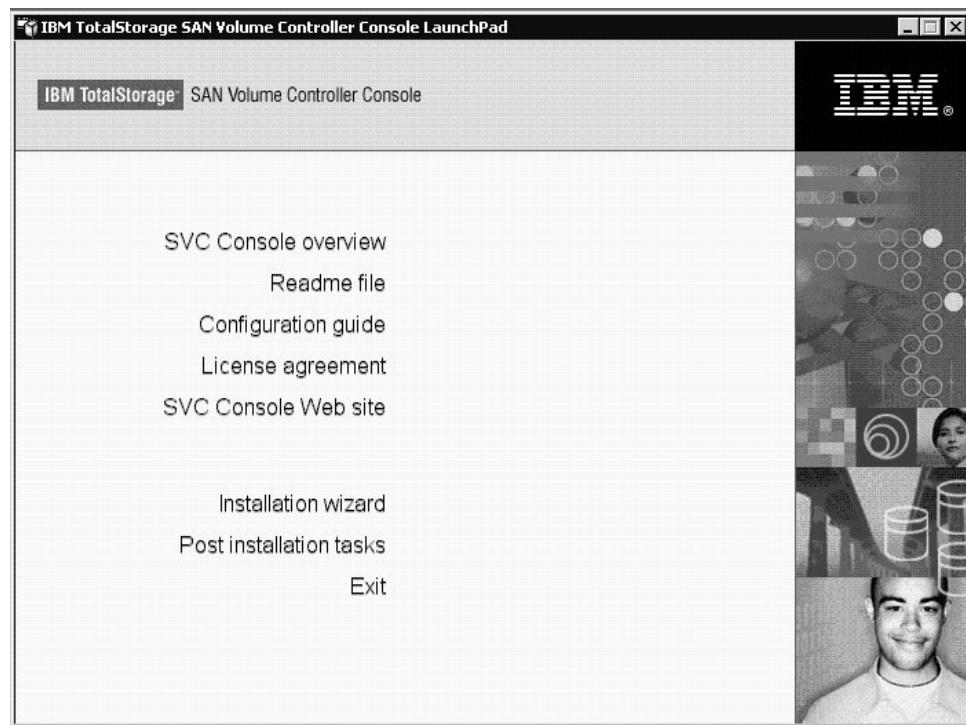


Figure 42. LaunchPad panel

**SVC Console overview**

Offers information about the IBM TotalStorage SAN Volume Controller Console.

**Readme file**

Offers any last minute product information that did not make it into sections concerning the installation of the IBM TotalStorage SAN Volume Controller Console.

**Configuration guide**

Appendix A, "Installing the IBM TotalStorage SAN Volume Controller Console for Windows", on page 241 has instructions about how to install the IBM TotalStorage SAN Volume Controller Console (a softcopy of this document).

**License agreement**

Offers information about the license for the IBM TotalStorage SAN Volume Controller Console.

**SVC Console Web site**

Offers information from the product Web site.

**Installation wizard**

Starts the IBM TotalStorage 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 cluster to the SAN Volume Controller Console management facility.

**Exit**

Exits the IBM TotalStorage SAN Volume Controller Console LaunchPad program.

- Click **Readme file** from the LaunchPad panel or from the **README.txt** file located in the doc or W2K directory on the IBM TotalStorage SAN Volume Controller Console CD to check for information that might supersede the information in this guide.
- Click **Installation wizard** from the LaunchPad panel to start the installation.

**Note:** The LaunchPad panel remains open behind the installation wizard so that you can access product information during the installation process. Click **Exit** if you want to close the LaunchPad.

- There might be a slight delay while the software loads on your system. After the software loads a DOS prompt window opens to display the following message:

```
Initializing InstallShield Wizard...
Searching for Java (tm) Virtual Machine .....
Searching for Java 1.3.1 by IBM Corporation.....
Verifying Java 1.3.1. by IBM Corporation.....
```

- The Welcome panel opens suggesting what documentation you should review prior to installation. Click **Next** to continue, or click **Cancel** to exit the installation.

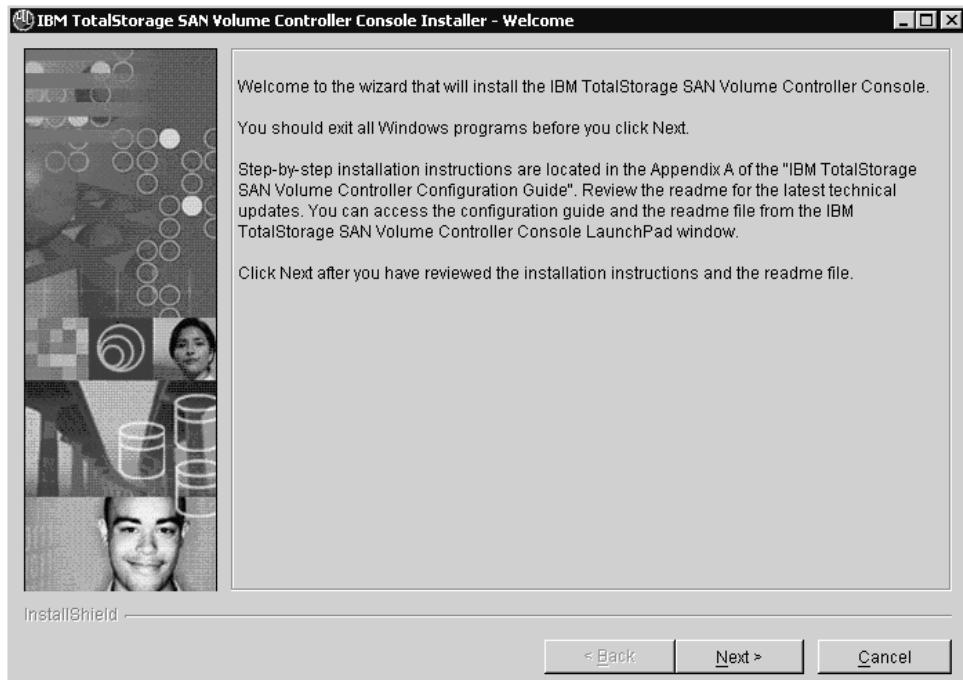


Figure 43. Welcome panel

- The License Agreement panel opens. Read the license agreement information. Select **I accept the terms of the license agreement**, then click **Next** to accept the license agreement. Otherwise, keep the selection **I do not accept the terms of the license agreement** (it is the default) and click **Cancel** to exit the installation.

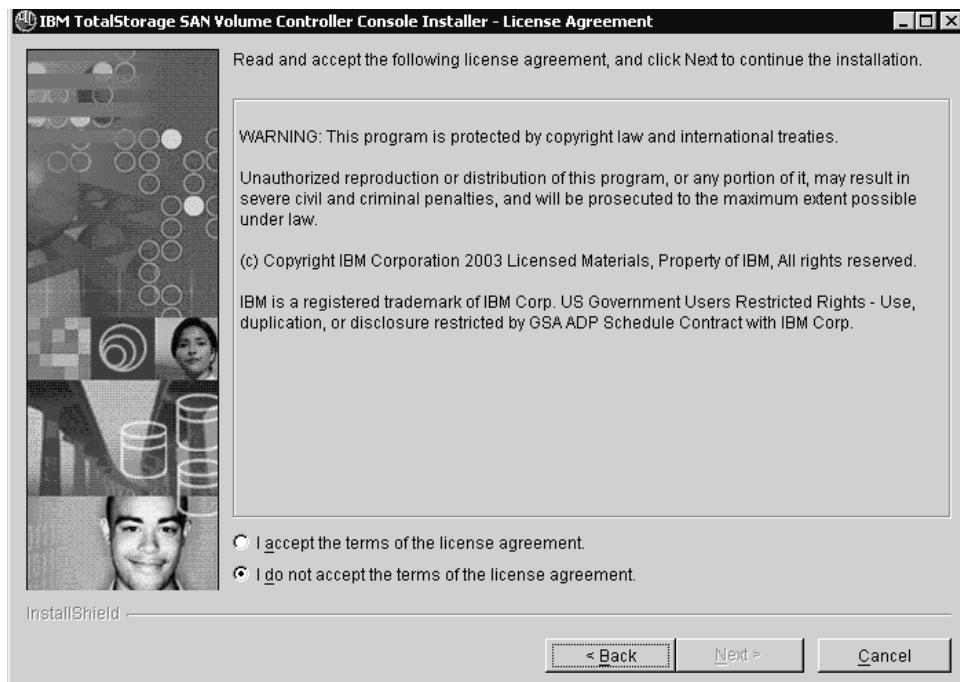


Figure 44. License Agreement panel

9. The installation wizard verifies that your machine meets the installation requirements.
  - If you have a Service Location Protocol (SLP) service that is different from the SLP that the IBM TotalStorage 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.
  - The installation wizard checks if the PuTTY SSH client is installed on your machine.
  - The installation wizard checks if a version of the IBM TotalStorage SAN Volume Controller Console is already installed. If the IBM TotalStorage SAN Volume Controller Console is *already installed*, the program checks if the Service Location Protocol (SLP), the IBM CIM Object Manager (CIMOM) service, and WebSphere Application Server V5 - SVC are started. If any of these services are started, the program asks if you want to continue the installation process by clicking **Next**. If you want to exit the installation program click **Cancel**. If you choose to continue, you must stop all the applications that use these services. (You can save the old configuration by selecting the Preserve Configuration check box that you see on the panel. If you chose to preserve the old configuration, the installation program will go directly to the Installation Confirmation panel below.)
10. The Destination Directory panel opens. Select one of the following options:



Figure 45. Destination Directory panel

- a. Click **Next** to accept the default directory.
- b. Click **Browse** to select a different directory for installation and then click **Next** to continue the installation process.
- c. Click **Cancel** to exit the installation process.

**Note:** If the program detects insufficient space for the IBM TotalStorage 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 go back by clicking **Back**, and choosing another destination directory for the product.

11. The Checking space panel is displayed.



Figure 46. Checking space panel

When the checking of available space completes, the PuTTY configuration panel opens.

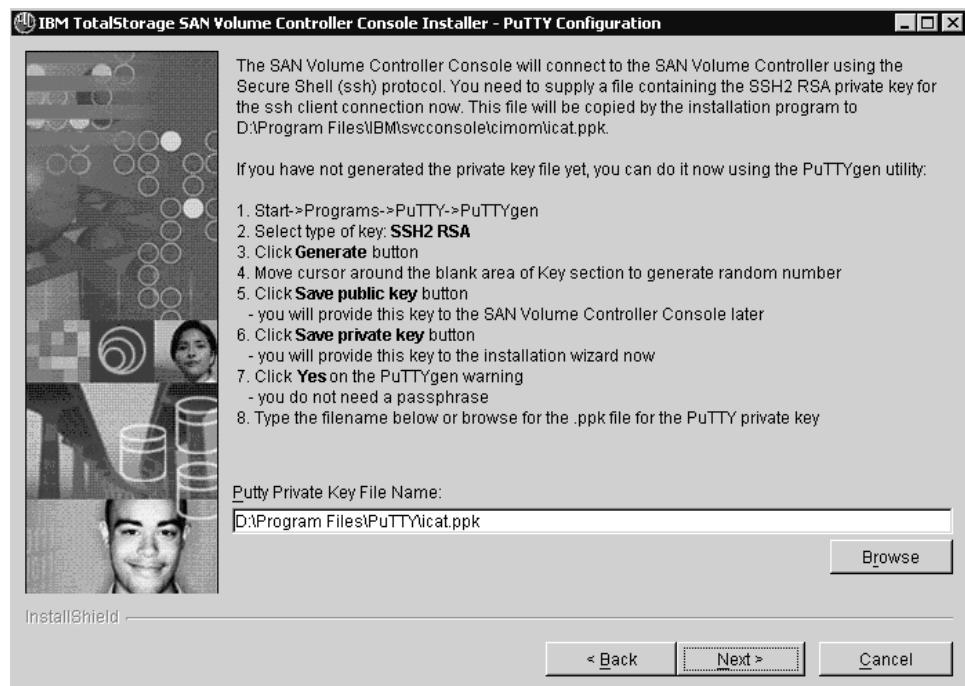


Figure 47. PuTTY Configuration panel

Enter the name and location on your system of your PuTTY SSH2 RSA private key file or click **Browse** to select the key file. If you have not prepared a PuTTY private key file yet, the steps on this panel tell you how to generate the PuTTY private and public key. Click **Next** to continue.

12. The Updating Ports panel is displayed.

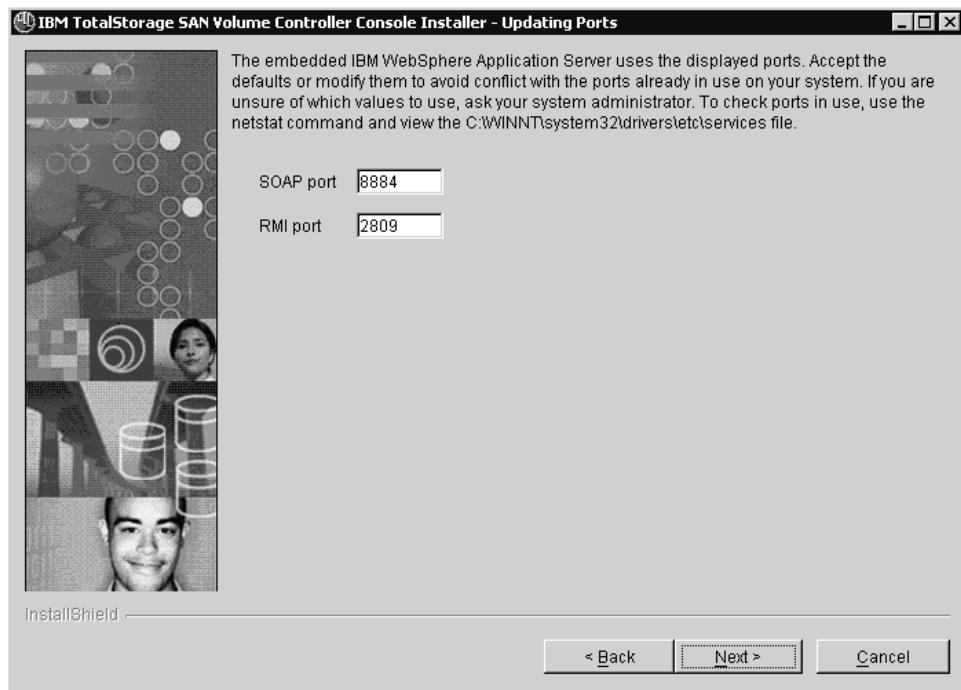


Figure 48. Updating Ports panel

Update the default ports assignments by typing unique port numbers for the products that have been registered on your system. To check ports in use, use the **netstat** command and view the C:\WINNT\system32\drivers\etc\services file. Click **Next** to continue.

13. The Installation Confirmation panel opens.

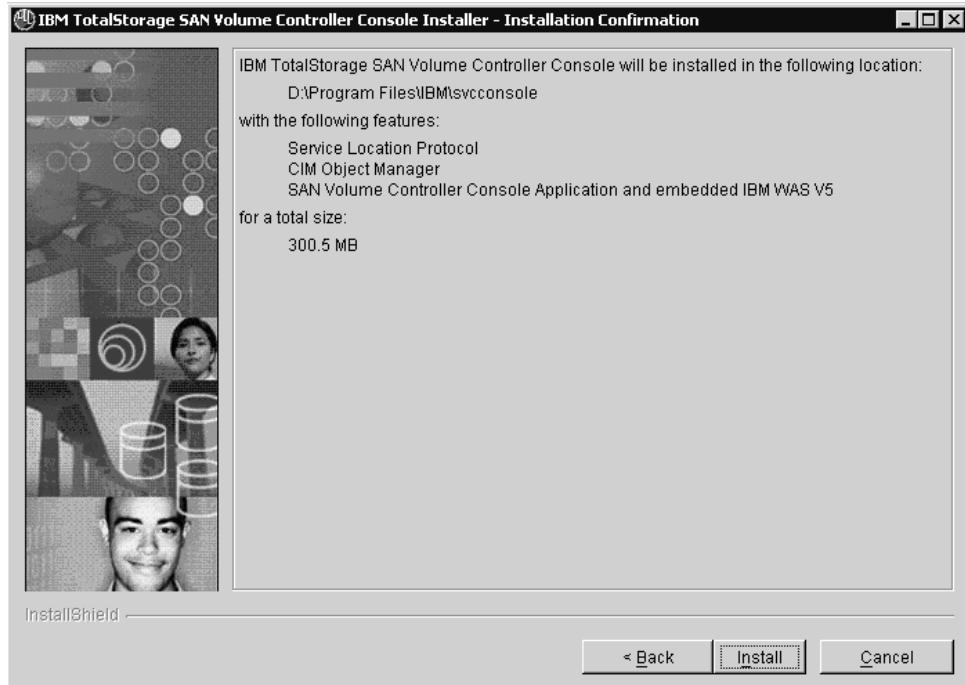


Figure 49. Installation Confirmation panel

Click **Install** to confirm the installation location and file size and to start the final installation. Click **Cancel** to exit the installation wizard or click **Back** to go to the previous panel.

14. The Installation Progress panel opens indicating how much of the installation has been completed.

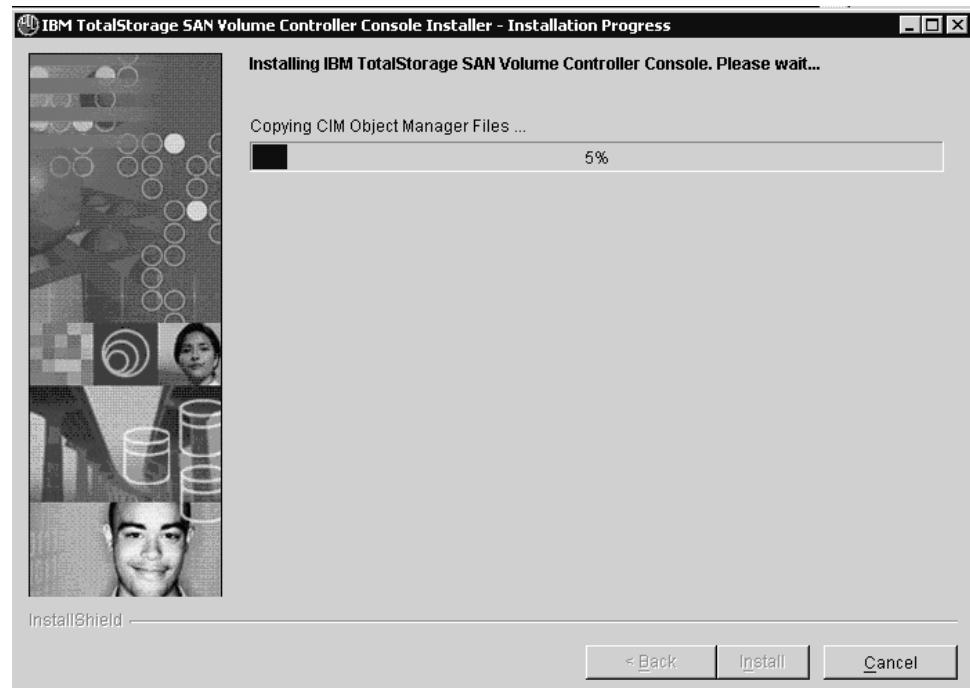


Figure 50. Installation Progress panel

Installation usually takes 3 - 10 minutes depending on the configuration of your machine.

**Note:** If you click **Cancel** a popup panel opens asking you to confirm the cancellation of the installation wizard: "Cancel the current operation? ". You may confirm the cancellation by clicking **Yes** or continue the installation by selecting **No**. If you confirm the cancellation, the information you entered or selected in the previous panel is not saved. You must start the installation again from the beginning.

After the completion of the successful installation of the IBM TotalStorage SAN Volume Controller Console, the installer attempts to start the following services:

- Service Location Protocol
- The IBM CIM Object Manager
- The IBM WebSphere Application Server V5 - SVC

Refer to "Verifying the Windows services associated with the SAN Volume Controller Console" on page 256 for further information.

15. When the Installation Progress panel closes, the Finish panel opens.

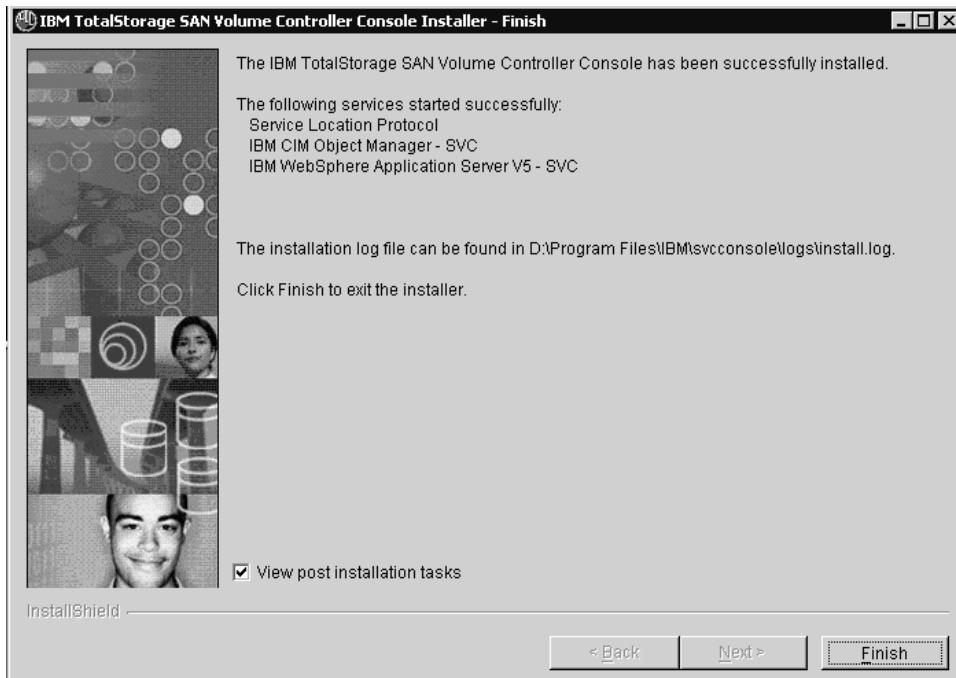


Figure 51. Finish panel

Before proceeding, you might want to review the log file for any possible error messages. The log file is located in `xxx\logs\install.log`, where `xxx` is the destination directory where the IBM TotalStorage SAN Volume Controller Console for Windows was installed. The `install.log` contains a trace of the installation actions.

**Note:** At the bottom of the Finish panel is a checkbox labeled **View post installation tasks**. If you check this box and then click **Finish**, the wizard will exit and the post installation tasks text file is displayed. This is the same information which is available in the section of this

manual entitled “Post Installation Tasks - Getting started using the SAN Volume Controller Console” on page 257. The LaunchPad panel Post Installation Tasks link also displays this same text file. You can avoid the display of the text file by unchecking the **View post installation tasks** box before you click the **Finish** button.

16. Click **Finish** to exit the installation wizard.

**Note:** Ordinarily, you do not need to restart your system during or after the installation of the IBM TotalStorage SAN Volume Controller Console. However, the installation wizard might determine that a restart is necessary. Restart your system if required. After you restart the system, the installation wizard continues with the installation.

17. If you have not yet reviewed the Post installation tasks from the installation Finish panel, review the Post installation tasks from the LaunchPad program.
  - a. Click **Post installation tasks** on the LaunchPad panel which opens the same file available from the installation Finish panel.
  - b. Continue with the post installation tasks for the SAN Volume Controller by following the instructions in this file.
18. Exit the LaunchPad program by clicking **Exit** on the LaunchPad panel.

#### Related topics:

- “Installation overview for the SAN Volume Controller Console” on page 241
- “Verifying the Windows services associated with the SAN Volume Controller Console” on page 256

---

## Installing the SAN Volume Controller Console in unattended (silent) mode

The unattended (silent) mode install option allows you to run installation unattended. Use this method of installation to customize a response file and issue a command from a command prompt window. The response file is a template on the IBM TotalStorage SAN Volume Controller Console CD. You can also create a standard response file to ensure that the product is installed consistently on multiple systems. You must satisfy all prerequisites before starting the installation.

#### Steps:

Perform the following steps to install the IBM TotalStorage SAN Volume Controller Console in your Windows environment using the unattended mode:

1. Log on to the system as a local system administrator.
2. Insert the IBM TotalStorage SAN Volume Controller Console CD into the CD drive.
3. If you have autorun mode set on your system, the IBM TotalStorage SAN Volume Controller Console program will start within 15-30 seconds. Click **Exit** from the LaunchPad.
4. Locate the response file (named *responsefile*) on your IBM TotalStorage SAN Volume Controller Console CD in the W2K directory.
5. Using Windows Explore or a command prompt, copy the response file to your hard drive.
6. Using a text editor modify the default options in the response file with the values you want:

- 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 ("").
  - 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.
  - The <-G *startUpgrade*> option enables the installation of CIM Agent over a previous installation of CIM Agent having the same version (reinstall) or lower version (upgrade). If you want to disable 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 whether or not to automatically stop SLP, CIMOM, and Embedded WAS 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 service 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 whether or not to save the configuration files when reinstalling or upgrading the product. If you do not want to save the configuration files when reinstalling or upgrading, remove the # character from the corresponding line and change the value of the option to no.
7. 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 used for a particular WebSphere service, remove the # character from the beginning of the line containing the option's value and set it to the value you desire. The following are the embedded WebSphere ports options:
- <-W ports.portSOAP="8884">
  - <-W ports.portRMI="2809">
- 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 cluster(s). Remove the # character from the corresponding line and add the fully qualified location of the PuTTY private key file. Save the responsefile *without* a file extension such as .txt.
8. From a command prompt window, type the following command:
- ```
<CD drive path>\W2K\install -options <response file path>\responsefile
```
- where <CD *drive path*> is the path of your CD drive. <*response file path*> is the path of the responsefile file that you copied in step 5 on page 253 and customized in step 6 on page 253.
9. During the installation you will see dotted lines scrolling across the screen. When the installation program ends you will see the control return to the Command Prompt.

10. Check for installation errors in the install.log file. After all the prerequisites checks have been performed, the log file is copied to the <dest-path>\logs directory. This file can be found in the <dest-path>\logs\ directory. This file is initially created in the system temporary file under the subdirectory cimagent. The following is an example of an install.log file:

```
(May 15, 2003 9:36:06 AM), This summary log is an overview of the
sequence of the installation of the IBM TotalStorage SAN Volume
Controller Console 1.0.0.12
(May 15, 2003 9:38:22 AM), IBM TotalStorage SAN Volume Controller
Console installation process started with the following install
parameters:
Target Directory: C:\Program Files\IBM\svcconsole
SOAP port: 8884
RMI port: 2809
(May 15, 2003 9:38:28 AM), Copying Service Location Protocol Files ...
(May 15, 2003 9:38:29 AM), Service Location Protocol successfully installed
(May 15, 2003 9:38:29 AM), Copying CIM Object Manager Files ...
(May 15, 2003 9:39:26 AM), The PUTTY private key successfully copied
into file C:\Program Files\IBM\svcconsole\cimom\icat.ppk
(May 15, 2003 9:39:51 AM), The file setupCmdLine.bat successfully updated.
(May 15, 2003 9:39:51 AM), Compile MOF files started ...
(May 15, 2003 9:40:06 AM), MOF files successfully compiled.
(May 15, 2003 9:40:06 AM), Generate a certificate store started ...
(May 15, 2003 9:40:19 AM), Certificate store called truststore
successfully generated.
(May 15, 2003 9:40:20 AM), IBM CIM Object Manager successfully installed
(May 15, 2003 9:40:20 AM), Installing embedded version of IBM WebSphere
Application Server ...
(May 15, 2003 9:41:42 AM), Websphere Application Server - SVC
successfully installed.
(May 15, 2003 9:43:20 AM), Copying SAN Volume Controller Console Ear Files...
(May 15, 2003 9:46:11 AM), The ICAConsole application successfully installed.
(May 15, 2003 9:47:24 AM), The SVCConsole application successfully installed.
(May 15, 2003 9:48:06 AM), The help application successfully installed.
(May 15, 2003 9:48:27 AM), The "C:\Program Files\IBM\svcconsole\console\
embeddedWAS\bin\expressPorts\UpdateExpressMultiPorts.bat" -soap 8884
-boot 2809 -remove" command updated successfully embedded WAS ports
in configuration files.
(May 15, 2003 9:48:27 AM), Command to be executed : net start cimomsrv
(May 15, 2003 9:48:49 AM), Command to be executed : net start
"IBMWAS5Service - SVC"
(May 15, 2003 9:50:15 AM), The following services started successfully:
Service Location Protocol
IBM CIM Object Manager
IBM WebSphere Application Server V5 - SVC
(May 15, 2003 9:50:15 AM), INSTSUCC: The IBM TotalStorage SAN Volume
Controller Console has been successfully installed.
```

11. Close the command prompt window by entering a command, for example **exit**.
12. After the completion of the successful installation of the IBM TotalStorage SAN Volume Controller Console, the installer attempts to start the following services:
  - Service Location Protocol
  - The IBM CIM Object Manager
  - IBM WebSphere Application Server V5 - SVC

"Verifying the Windows services associated with the SAN Volume Controller Console" on page 256
13. Continue with the post installation tasks for the IBM TotalStorage SAN Volume Controller Console using the instructions in the following section. You can also view the post installation tasks using the following option:
  - a. From a Command Prompt, change directory into the W2K directory on the CD drive. Open the LaunchPad by typing:

## LaunchPad

- b. Click **Post installation tasks** on the LaunchPad window. Continue with the post installation tasks for the IBM TotalStorage SAN Volume Controller Console by following the instructions in this file.

### Related topics:

- “Installation overview for the SAN Volume Controller Console” on page 241
- “Verifying the Windows services associated with the SAN Volume Controller Console”

---

## Verifying the Windows services associated with the SAN Volume Controller Console

This task verifies that the Windows services associated with your IBM TotalStorage SAN Volume Controller Console are correctly installed and started.

### Steps:

Perform the following steps to verify your Service Location Protocol (SLP), IBM CIM Object Manager (CIMOM), and IBM WebSphere Application Server V5 - SVC services were correctly installed:

1. Verify the installation of the Service Location Protocol (SLP).
  - a. Verify that the Service Location Protocol is started. Select **Start -> Settings -> Control Panel**. Double-click the **Administrative Tools** icon. Double-click the **Services** icon.
  - b. Find **Service Location Protocol** in the Services window list. For this component, the Status column should be marked **Started**.
  - c. If the Service Location Protocol is not started, right-click on Service Location Protocol and select **Start** from the pop-up menu. Wait for the Status column to be changed to **Started**.
  - d. Do not close the Services window because you will also use it to verify the The CIM Object Manager (CIMOM) service.
2. Verify the installation of the SAN Volume Controller Console.
  - a. Find the **IBM CIM Object Manager** in the Services window list. For this component, the Status column should be marked **Started**.
  - b. If the IBM CIM Object Manager is not started, right click on the **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 IBM WebSphere Application Server V5 - SVC service.
3. Verify the installation of the IBM WebSphere Application Server V5 - SVC service.
  - a. Find the **IBM WebSphere Application Server V5 - SVC** in the Services window list. For this component, the Status column should be marked **Started**.
  - b. If the IBM WebSphere Application Server V5 - SVC service is not started, right click on the **IBM WebSphere Application Server V5 - SVC** and select **Start** from the pop-up menu. Wait for the Status column to change to **Started**.
  - c. Close the Services window.
  - d. Close the Administrative Tools window.

---

## Post Installation Tasks - Getting started using the SAN Volume Controller Console

This section outlines how to get started using the SAN Volume Controller Console using your Web browser. In case you are new to using the SAN Volume Controller Console, this document can serve as an introduction to using the SAN Volume Controller Console.

Once you have installed the IBM TotalStorage SAN Volume Controller Console 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 Console as well as configuring SAN Volume Controller clusters.

Each time you wish to add a SAN Volume Controller cluster to the collection of clusters managed by the IBM TotalStorage SAN Volume Controller Console, you must store the PuTTY SSH client public key which is 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.

### Steps:

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 onto 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 typing the following:

`http://localhost:9080/ica`

2. Log onto 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 `passw0rd`. 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. Accessing 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 will open in 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 will open 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. See Chapter 5, “Create cluster from the front panel”, on page 51 for more information. You will be given a special password by the customer engineer (CE) to be used in later steps of initializing the SAN Volume Controller console.

After you create the SAN Volume Controller cluster using the front panel of cluster, you will need to complete the creation of the cluster by using the SAN Volume Controller Console Web pages. See Chapter 9, “Overview of creating a cluster using the SAN Volume Controller Console”, on page 81 for more information.

Enter the IP address of the cluster and check the **Create (Initialize) Cluster** box. When you click the **OK** button, the create cluster wizard will take over and present you with the panels you need to complete initializing the cluster.

The browser will then prompt you to enter the network password. Enter the user name `admin` and the password provided to you by the customer engineer (CE) during the cluster front panel creation phase.

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. See “Adding subsequent SSH public keys to the SAN Volume Controller” on page 129 for more information. 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, you simply click the **Add SAN Volume Controller Cluster** button and then add the cluster IP address but *do not* check the **Create (Initialize) Cluster** box, which is above the **OK** button. When you click the **OK** button, you will be taken to the Web page to provide the PuTTY SSH client public key to upload to the cluster. Step 5 below continues with the SSH key input description.

The browser will then prompt you to enter the network password. Enter the user name `admin` and the password which is configured for the cluster. Then Click **OK**.

5. Store the SAN Volume Controller console system SSH public key on the SAN Volume Controller Console. This PuTTY client SSH public key is the other key in the key pair you provided to the SAN Volume Controller Console during the installation program. Each key is associated with an ID string that you define that can consist of up to 30 characters. Up to 100 keys can be stored on a

cluster. You can add keys to provide either *administrator* access or *service* access. Perform the following steps to store the SSH public key on the cluster:

- a. Enter the SSH public key name and directory location on your local browser system in the field labeled **Public Key (file upload)** or click **Browse** to identify the key on the local system. Alternatively, you can paste the SSH key into the **Public Key (direct input)** field.
  - b. Enter an ID string in the field labeled **ID**. This is a unique ID to distinguish the key and is not related to a user name.
  - c. Select the *administrator* **Access Level** radio button.
  - d. Click **Add Key** to store this SSH public key on the cluster.
6. Launch the secondary Web browser window to manage your specific cluster. Once you have identified the SAN Volume Controller clusters to the SAN Volume Controller Console you can see a summary of all clusters. From this point, you can select the specific cluster in which you are interested and then launch the browser window specifically for the cluster. Perform the following steps to launch the browser window:
    - a. Click **Clusters** in the portfolio section of your browser window in the left-hand frame. A new view will be displayed in the work area.
    - b. Check the small box in the Select column left of the cluster in which you are interested to select that cluster. Select **Launch the SAN Volume Controller application** in the drop down list box of the work area and click **Go**. A secondary browser window opens to the SAN Volume Controller (SVC) 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.

#### **Result:**

This completes the verification of the connection to the SAN Volume Controller.

---

## **Removing the SAN Volume Controller Console**

This optional task provides the instructions for removing the IBM TotalStorage SAN Volume Controller Console from your Windows system.

#### **Steps:**

Perform the following steps to remove the IBM TotalStorage SAN Volume Controller Console:

1. Log onto the system where the IBM TotalStorage SAN Volume Controller Console is installed as a local system administrator.

2. Stop the IBM CIM Object Manager (CIMOM), IBM WebSphere Application Server V5 - SVC and the Service Location Protocol (SLP) services if they are started.
  - a. Click **Start -> Settings -> Control Panel**. In the Control Panel window, double-click on the **Administrative Tools** icon and then double-click the **Services** icon. The Services window opens.
  - b. Stop the IBM CIM Object Manager (CIMOM) service:
    - 1) In the Services window, scroll to **IBM CIM Object Manager**. Click on the service to select it.
    - 2) If the Status column shows Started, right-click the service, then click **Stop** on the menu.
  - c. Stop the IBM WebSphere Application Server V5 - SVC service:
    - 1) In the Services window, scroll to **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.
  - d. Stop the Service Location Protocol (SLP) service:

**Note:** You must be careful if you have other applications that use the Service Location Protocol (SLP) service. In this case, you must stop these applications before stopping Service Location Protocol (SLP) service, because during the removal process the Service Location Protocol (SLP) service will be deleted. You must also stop the configuration utilities for the IBM TotalStorage SAN Volume Controller Console, if they are running.

- 1) In the Services window, scroll to **Service Location Protocol**. Click on this service to select it.
- 2) If it is running (the Status column shows Started), right-click the service, then click **Stop** on the menu.  
(If you did not stop the IBM CIM Object Manager service, the system now asks if you want to stop the IBM CIM Object Manager service. Because the IBM CIM Object Manager service is dependent on the Service Location Protocol service which you just stopped, you must click **Yes** to stop the IBM CIM Object Manager service.)
- 3) Wait for the services to stop.
- 4) Close the Services window.
- 5) Close the Administrative Tools window.
3. Use the Windows Add/Remove Programs facility to remove the IBM TotalStorage SAN Volume Controller Console and the Service Location Protocol components.
  - a. From the Windows® menu bar, click **Start -> Settings -> Control Panel**. Double-click **Add/Remove Programs**.
  - b. Click **IBM TotalStorage SAN Volume Controller Console** from the list of currently installed programs and click **Remove** to remove the product.
4. The Welcome panel for the Uninstaller opens. Click **Next** to continue or click **Cancel** to stop the removal of the IBM TotalStorage SAN Volume Controller Console.
5. The program detects whether the Service Location Protocol, IBM CIM Object Manager, and the IBM 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 IBM TotalStorage SAN Volume Controller Console are dependent on the services. You can either:
    - Click **Next** to have the program stop the services for you.
    - Click **Cancel** to exit the removal process if you wish to manually stop the services and any dependent applications. Instructions for stopping the services are described in step 2 on page 260. You must then restart the removal process from the Windows Add/Remove facility.
6. The Confirmation panel opens. Click **Remove** to continue or click **Cancel** to stop the removal of the IBM TotalStorage SAN Volume Controller Console. Click **Back** to return to the previous panel.
  7. The Uninstallation Progress panel opens. Wait for the program to remove the IBM TotalStorage SAN Volume Controller product.
  8. The Finish panel for the Uninstaller opens. This panel indicates the result of the removal process (successful or failed). Click **Finish** to complete the removal process and exit the wizard.  
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 chose to either restart your computer now or restart your computer at a later time. Then click **Finish** to complete the removal process and exit the wizard.
  9. Close the Add/Remove Programs window.

#### **Post-processing requirements:**

Perform the following steps to complete the removal process:

1. If the system has not been restarted since IBM TotalStorage SAN Volume Controller Console was removed, do so now.
2. Log onto the system as a local system administrator.
3. The removal process saves files uniquely related to the configuration in a backup directory under the destination path where you installed the IBM TotalStorage 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\svcconsole.
4. Perform other cleanup tasks:
  - Empty your Windows Recycle Bin to reclaim the disk space that was made available during the removal process.



---

## Appendix B. Valid combinations of FlashCopy and Remote Copy functions

The following table outlines the combinations of FlashCopy and Remote Copy functions that are valid for a single virtual disk (VDisk).

*Table 25. Valid combinations of FlashCopy and Remote Copy interactions*

| FlashCopy        | Remote Copy Primary | Remote Copy Secondary |
|------------------|---------------------|-----------------------|
| FlashCopy source | Supported           | Supported             |
| FlashCopy target | Not supported       | Not supported         |

**Related topics:**

- “FlashCopy” on page 27
- “Remote Copy” on page 37



---

## Appendix C. Setting up SNMP traps

This topic provides overview information about setting up SNMP traps if the master console has been installed on a separate machine.

### Prerequisites:

There are two steps required to enable the Call-Home process:

1. Set up the SAN Volume Controller SNMP Trap destination, a specific machine (IP Address)
2. Set up IBM Director to send a correctly formatted e-mail

### Overview:

To set up the SAN Volume Controller SNMP trap destination, the destination is normally set up as part of the SAN Volume Controller Installation process, but can also be done through the SAN Volume Controller Web pages, by using a browser to log on to the SAN Volume Controller cluster and selecting the option Error notification. See *IBM TotalStorage Virtualization Family SAN Volume Controller Installation and Hardware Reference* for more information.



---

## Appendix D. Configuring IBM Director overview

This task provides step-by-step instructions for configuring IBM Director for Call-Home and E-mail, if it has been installed on a separate machine or is re-installed on the master console.

### Steps:

Perform the following steps to configure the IBM Director:

1. Set up and Event Action Plan
2. Set up a correctly formatted e-mail

### Related topics:

- “Setting up an Event Action Plan”
- “Setting up an e-mail” on page 268

---

## Setting up an Event Action Plan

This task provides step-by-step instructions for setting up Event Action Plans if the IBM Director has been installed on a separate machine or is re-installed on the master console. In order for IBM Director to present the correct SAN Volume Controller information to enable an Action Plan to be configured, it has to have received a Trap from the SAN Volume Controller.

### Steps:

Perform the following steps to set up an Event Action Plan:

1. Create a SAN Volume Controller Trap by removing the AC power from one of the UPS units, supplying the cluster and replace the power after 30 seconds.
2. Click **Event Log (ALL)** from the IBM Director Console and check that the Trap from the SAN Volume Controller has been received.
3. Click **Tasks -> Event Action Plan Builder** from the IBM Director Console.
4. Right-click **Simple Event Filter**.
5. Click **New**.
6. Click the **Event type** tab from the Simple Event Filter Builder window.
7. Uncheck the box labeled Any.
8. In the list, select in sequence the following:
  - a. SNMP
  - b. 1 (iso)
  - c. 2 (org)
  - d. 6 (dod)
  - e. 1 (internet)
  - f. 4 (private)
  - g. 1 (enterprise)
  - h. 2 (ibm)
  - i. 6 (ibmprod)
  - j. 190

- k. 1
9. Click the **Category** tab.
  10. Uncheck the box labeled Any.
  11. Click the **Alert** option.
  12. On the menu bar, click **File** and save the file with the name **2145 Error**.
  13. From the Event Filter List, select the newly created **2145 Error** filter and drag and drop it on to the Log All Events icon in the Event Action Plan column. This action causes the **2145 Error** filter to be called upon when any Event is logged.
  14. Perform steps 4 through 11 again (do not do step 8k). On the menu bar, click **File** and save the file with the name **2145 Event**.
  15. From the Event Filter List, select the newly created **2145 Event** filter and drag and drop it on to the Log All Events icon in the Event Action Plan column. This action causes the **2145 Event** filter to be called upon when any Event is logged.

**Related topics:**

- “Setting up an e-mail”

---

## Setting up an e-mail

This task provides step-by-step instructions for setting up the e-mails if the IBM Director has been installed on a separate machine or is re-installed on the master console.

**Steps:**

Perform the following steps to set up e-mail for Call-Home:

1. From the IBM Director Console menu bar, select **Tasks** → **Event Action Plan Builder**.
2. In the **Actions** column, right-click on **Send an Internet (SMTP) E-mail** and select **Customize**.
3. In the resulting **Customize Action: Send an Internet (SMTP) E-mail** panel fill-in:

**Internet E-mail Address**

- Enter the IBM Retain E-mail address
  - CALLHOME1@de.ibm.com for USA customer’s
  - CALLHOME0@de.ibm.com for customer’s outside of the USA.

**Reply to**

- Enter the E-mail address that you require any replies to be directed

**SMPT E-mail Server**

- Enter the address of your E-mail server

**SMPT Port**

- Change this, if required to your SMTP Server port number

**Subject of E-mail Message**

- Fill in **2145 Error Notification**.

### **Body of the E-mail Message**

- Fill in the following information:
  - Contact name.....not required in the E-mail to Admin
  - Contact phone number.....not required in the E-mail to Admin
  - Offshift phone number.....not required in the E-mail to Admin
  - Machine location
  - Record Type = 1

iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.1  
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.2  
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.3  
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.4  
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.5  
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.6  
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.7  
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.8  
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.9  
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.10  
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.11  
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.12

4. Click **Save** to save the information, using the name **2145CallHome**.
5. From the **Send an Internet (SMTP) E-mail** list select the newly created **2145CallHome** E-mail and Drag and Drop it on to the **2145 Error** action plan icon in the **Event Action Plan** column. This action causes the **2145CallHome** to be call when the **2145 Error** filter is satisfied.

---

## **Setting up an e-mail user notification**

This task provides step-by-step instructions for setting up the e-mails if the IBM Director has been installed on a separate machine or is re-installed on the master console.

### **Steps:**

Perform the following steps to set up e-mail for user notification:

1. From the IBM Director Console menu bar, select **Tasks -> Event Action Plan Builder**.
2. In the **Actions** column, right-click on **Send an Internet (SMTP) E-mail** and select **Customize**.
3. In the resulting **Customize Action: Send an Internet (SMTP) E-mail** panel fill-in :

#### **Internet E-mail Address**

- Enter the E-mail address you require for notification

#### **Reply to**

- Enter the E-mail address that you require any replies to be directed

#### **SMPT E-mail Server**

- Enter the address of your E-mail server

#### **SMPT Port**

- Change this, if required to your SMTP Server port number

### **Subject of E-mail Message**

- Fill in 2145 Error Notification.

### **Body of the E-mail Message**

- Fill in the following information:
  - # Machine location = xxxx

iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.1  
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.2  
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.3  
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.4  
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.5  
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.6  
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.7  
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.8  
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.9  
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.10  
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.11  
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.12

Where xxxx is information relevant to your organization.

4. Click **Save** to save the information, using the name **2145ErrorNot**.
5. From the **Send an Internet (SMTP) E-mail** list select the newly created **2145ErrorNot** E-mail and Drag and Drop it on to the **2145 Event** action plan icon in the **Event Action Plan** column. This action causes the **2145ErrorNot** to be call when the **2145 Event** filter is satisfied.

---

## Appendix E. Object types

This topic provides information about object types.

The following table lists the object codes and its corresponding object type.

*Table 26. Object types*

| Object code | Object type                 |
|-------------|-----------------------------|
| 0           | IC_TYPE_Unknown             |
| 1           | IC_TYPE_Vlun                |
| 2           | IC_TYPE_Vlungrp             |
| 3           | IC_TYPE_Hlun                |
| 4           | IC_TYPE_Node                |
| 5           | IC_TYPE_Host                |
| 6           | IC_TYPE_Hostgrp             |
| 7           | IC_TYPE_Hws                 |
| 8           | IC_TYPE_Fcgrp               |
| 9           | IC_TYPE_Rcgrp               |
| 10          | IC_TYPE_Fcmap               |
| 11          | IC_TYPE_Rcmap               |
| 12          | IC_TYPE_Wwpn                |
| 13          | IC_TYPE_Cluster             |
| 15          | IC_TYPE_Hba                 |
| 16          | IC_TYPE_Device              |
| 17          | IC_TYPE_SCSI_Lun            |
| 18          | IC_TYPE_Quorum              |
| 19          | IC_TYPE_TimeSeconds         |
| 20          | IC_TYPE_ExtSInst            |
| 21          | IC_TYPE_ExtInst             |
| 22          | IC_TYPE_Percentage          |
| 23          | IC_TYPE_VPD_SystemBoard     |
| 24          | IC_TYPE_VPD_Processor       |
| 25          | IC_TYPE_VPD_Processor_Cache |
| 26          | IC_TYPE_VPD_Memory_Module   |
| 27          | IC_TYPE_VPD_Fan             |
| 28          | IC_TYPE_VPD_FC_Card         |
| 29          | IC_TYPE_VPD_FC_Device       |
| 30          | IC_TYPE_VPD_Software        |
| 31          | IC_TYPE_VPD_Front_Panel     |
| 32          | IC_TYPE_VPD_UPS             |
| 33          | IC_TYPE_VPD_Port            |

*Table 26. Object types (continued)*

| Object code | Object type        |
|-------------|--------------------|
| 34          | IC_TYPE_FC_Adapter |
| 35          | IC_TYPE_Migrate    |

---

## Appendix F. Event codes

This topic provides information about information and configuration event codes.

There are two different types of event codes:

- Information event codes
- Configuration event codes

Information event codes, when generated, provide information on the status of a particular operation. Information event codes are recorded in the error log and an SNMP trap and sometimes an e-mail is generated if the corresponding management flag is set in the Preference cache.

Configuration event codes are generated when configuration parameters are set. Configuration event codes are recorded in a separate log and do not generate SNMP traps or e-mails and their error fixed flags are ignored.

### Related topics:

- “Information event codes”
- “Configuration event codes” on page 274

---

## Information event codes

This topic provides information about the information event codes.

The information event codes, when generated, provide information on the status of a particular operation. Information event codes are recorded in the error log and an SNMP trap and sometimes an e-mail is generated if the corresponding management flag is set in the Preference cache.

Information event codes generate information type (I) descriptions or warning type (W) descriptions.

*Table 27. Information event codes*

| Event code | Type | Description                                                                                  |
|------------|------|----------------------------------------------------------------------------------------------|
| 980321     | W    | The Managed Disk group is degraded or offline and consequently this Virtual Disk is offline. |
| 980350     | I    | The node is now a functional member of the cluster                                           |
| 980351     | I    | A non-critical hardware error has occurred                                                   |
| 980370     | I    | Both nodes in the I/O group are available                                                    |
| 980371     | I    | One node in the I/O group is not available                                                   |
| 980372     | W    | Both nodes in the I/O group are not available                                                |
| 980392     | I    | Cluster recovery completed.                                                                  |
| 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                                                                      |

*Table 27. Information event codes (continued)*

| Event code | Type | Description                                                                                                                                                    |
|------------|------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 981001     | W    | Cluster Fabric View has been updated by a multiphase discovery                                                                                                 |
| 981007     | W    | Preferred port is not being used for Managed Disk access                                                                                                       |
| 981014     | W    | LUN Discovery failed. Cluster has a connection to a device through this node but this node cannot discover the Managed Disks associated LUN correctly.         |
| 981020     | W    | Managed Disk error count warning threshold met.                                                                                                                |
| 982003     | W    | Insufficient Virtual Extents.                                                                                                                                  |
| 982007     | W    | Migration Stopped.                                                                                                                                             |
| 982009     | I    | Migrate 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. |
| 985001     | I    | Remote copy, background copy complete                                                                                                                          |
| 985002     | I    | Remote copy 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.                                                                                                       |

#### Related topics:

- Appendix F, “Event codes”, on page 273
- “Configuration event codes”

---

## Configuration event codes

This topic provides information about the configuration event codes.

Configuration event codes are generated when configuration parameters are set. Configuration event codes are recorded in a separate log and do not generate SNMP traps or e-mails and their error fixed flags are ignored.

*Table 28. Configuration event codes*

| Event code | Description                                                                                           |
|------------|-------------------------------------------------------------------------------------------------------|
| 990101     | Modify cluster (attributes in the <b>svctask chcluster</b> command)                                   |
| 990105     | Delete node from cluster (attributes in the <b>svctask rmnode</b> command)                            |
| 990106     | Create host (attributes in the <b>svctask mkhost</b> command)                                         |
| 990112     | Cluster config dumped to file (attributes in the <b>svctask dumpconfig</b> command)                   |
| 990117     | Create cluster (attributes in the <b>svctask mkcluster</b> command)                                   |
| 990118     | Modify node (attributes in the <b>svctask chnode</b> command)                                         |
| 990120     | Shutdown node (attributes in the <b>svctask stopcluster</b> command)                                  |
| 990128     | Modify host (attributes in the <b>svctask chhost</b> command)                                         |
| 990129     | Delete node (attributes in the <b>svctask rmnode</b> command)                                         |
| 990138     | Virtual Disk Modify (attributes in the <b>svctask chvdisk</b> command)                                |
| 990140     | Virtual Disk Delete (attributes in the <b>svctask rmvdisk</b> command)                                |
| 990144     | Modify Managed Disk group (attributes in the <b>svctask chmdiskgrp</b> command)                       |
| 990145     | Delete Managed Disk group (attributes in the <b>svctask rmdiskgrp</b> command)                        |
| 990148     | Create Managed Disk group (attributes in the <b>svctask mkmdiskgrp</b> command)                       |
| 990149     | Modify Managed Disk (attributes in the <b>svctask chmdisk</b> command)                                |
| 990158     | VLUN included                                                                                         |
| 990159     | Quorum created                                                                                        |
| 990160     | Quorum Destroy                                                                                        |
| 990168     | Modify the HWS an Virtual Disk is assigned to                                                         |
| 990169     | Create a new Virtual Disk (attributes in the <b>svctask mkvdisk</b> command)                          |
| 990173     | Add a Managed Disk to Managed Disk group (attributes in the <b>svctask addmdisk</b> command)          |
| 990174     | Delete a Managed Disk from Managed Disk group (attributes in the <b>svctask rmmdisk</b> command)      |
| 990178     | Add a port to a Host (attributes in the <b>svctask addhostport</b> command)                           |
| 990179     | Delete a port from a Host (attributes in the <b>svctask rmhostport</b> command)                       |
| 990182     | Create an Virtual Disk to Host SCSI mapping (attributes in the <b>svctask mkvdiskhostmap</b> command) |
| 990183     | Delete an Virtual Disk to Host SCSI mapping (attributes in the <b>svctask rmdiskhostmap</b> command)  |
| 990184     | Create a FlashCopy mapping (attributes in the <b>svctask mkfcmap</b> command)                         |
| 990185     | Modify a FlashCopy mapping (attributes in the <b>svctask chfcmap</b> command)                         |
| 990186     | Delete a FlashCopy mapping (attributes in the <b>svctask rmfcmap</b> command)                         |

*Table 28. Configuration event codes (continued)*

| Event code | Description                                                                                           |
|------------|-------------------------------------------------------------------------------------------------------|
| 990187     | Prepare a FlashCopy mapping (attributes in the <b>svctask prestartfcmap</b> command)                  |
| 990188     | Prepare a FlashCopy consistency group (attributes in the <b>svctask prestartfcconsistgrp</b> command) |
| 990189     | Trigger a FlashCopy mapping (attributes in the <b>svctask startfcmap</b> command)                     |
| 990190     | Trigger a FlashCopy consistency group (attributes in the <b>svctask startfcconsistgrp</b> command)    |
| 990191     | Stop a FlashCopy mapping (attributes in the <b>svctask stopfcmap</b> command)                         |
| 990192     | Stop a FlashCopy consistency group (attributes in the <b>svctask stopfcconsistgrp</b> command)        |
| 990193     | FlashCopy set name                                                                                    |
| 990194     | Delete a list of ports from a Host (attributes in the <b>svctask rmhostport</b> command)              |
| 990196     | Shrink a virtual disk.                                                                                |
| 990197     | Expand a Virtual Disk (attributes in the <b>svctask expandvdisksize</b> command)                      |
| 990198     | Expand single extent a Virtual Disk                                                                   |
| 990199     | Modify govern a Virtual Disk                                                                          |
| 990203     | Initiate manual Managed Disk discovery (attributes in the <b>svctask detectmdisk</b> command)         |
| 990204     | Create FlashCopy consistency group (attributes in the <b>svctask mkfcconsistgrp</b> command)          |
| 990205     | Modify FlashCopy consistency group (attributes in the <b>svctask chfcconsistgrp</b> command)          |
| 990206     | Delete FlashCopy consistency group (attributes in the <b>svctask rmfcconsistgrp</b> command)          |
| 990207     | Delete a list of Hosts (attributes in the <b>svctask rmhost</b> command)                              |
| 990213     | Change the HWS a node belongs to (attributes in the <b>svctask chiogrp</b> command)                   |
| 990216     | Apply software upgrade (attributes in the <b>svcservicetask applysoftware</b> command)                |
| 990219     | Analyze error log (attributes in the <b>svctask finderr</b> command)                                  |
| 990220     | Dump error log (attributes in the <b>svctask dumperlog</b> command)                                   |
| 990221     | Clear error log (attributes in the <b>svctask clearerrlog</b> command)                                |
| 990222     | Fix error log entry (attributes in the <b>svctask cherrstate</b> command)                             |
| 990223     | Migrate a single extent (attributes in the <b>svctask migrateexts</b> command)                        |
| 990224     | Migrate a number of extents                                                                           |
| 990225     | Create Remote copy relationship (attributes in the <b>svctask mkrcrelationship</b> command)           |
| 990226     | Modify Remote copy relationship (attributes in the <b>svctask chrcrelationship</b> command)           |

*Table 28. Configuration event codes (continued)*

| Event code | Description                                                                                          |
|------------|------------------------------------------------------------------------------------------------------|
| 990227     | Delete Remote copy relationship (attributes in the <b>svctask rmrcrelationship</b> command)          |
| 990229     | Start Remote copy relationship (attributes in the <b>svctask startrcrelationship</b> command)        |
| 990230     | Stop Remote copy relationship (attributes in the <b>svctask stoprcrelationship</b> command)          |
| 990231     | Switch a Remote copy relationship (attributes in the <b>svctask switchrcrelationship</b> command)    |
| 990232     | Start Remote copy consistency group (attributes in the <b>svctask startrcconsistgrp</b> command)     |
| 990233     | Stop Remote copy consistency group (attributes in the <b>svctask stoprcconsistgrp</b> command)       |
| 990234     | Switch a Remote copy consistency group (attributes in the <b>svctask switchrcconsistgrp</b> command) |
| 990237     | Create partnership with remote cluster (attributes in the <b>svctask mkpartnership</b> command)      |
| 990238     | Modify partnership with remote cluster (attributes in the <b>svctask chpartnership</b> command)      |
| 990239     | Delete partnership with remote cluster (attributes in the <b>svctask rmpartnership</b> command)      |
| 990240     | Create Remote copy consistency group (attributes in the <b>svctask mkrcconsistgrp</b> command)       |
| 990241     | Modify Remote copy consistency group (attributes in svctask chrcconsistgrp)                          |
| 990242     | Delete Remote copy consistency group (attributes in the <b>svctask rmrcconsistgrp</b> command)       |
| 990245     | Node pend                                                                                            |
| 990246     | Node remove                                                                                          |
| 990247     | Node unpend                                                                                          |
| 990380     | Time zone changed (attributes in the <b>svctask settimezone</b> command)                             |
| 990383     | Change cluster time (attributes in the <b>svctask setclustertime</b> command)                        |
| 990385     | System time changed                                                                                  |
| 990386     | SSH key added (attributes in the <b>svctask addsshkey</b> command)                                   |
| 990387     | SSH key removed (attributes in the <b>svctask rmsshkey</b> command)                                  |
| 990388     | All SSH keys removed (attributes in the <b>svctask rmallsshkeys</b> command)                         |
| 990390     | The add node cluster command executed. Check error data.                                             |
| 990395     | Node Shutdown/Reset command failed                                                                   |
| 990410     | Software Install started                                                                             |
| 990415     | Software Install completed                                                                           |
| 990420     | Software Install failed                                                                              |
| 990430     | Planar Serial Number changed                                                                         |
| 990501     | The featurization has changed. See feature log for details.                                          |

*Table 28. Configuration event codes (continued)*

| Event code | Description                                                       |
|------------|-------------------------------------------------------------------|
| 991024     | IO tracing has finished, trigger occurred for given Managed Disk. |

**Related topics:**

- Appendix F, “Event codes”, on page 273
- “Information event codes” on page 273

---

## Appendix G. Accessibility

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

### Features:

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

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

### Navigating by keyboard:

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

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

### Accessing the publications:

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

[www.ibm.com/storage/support/2145/](http://www.ibm.com/storage/support/2145/)

### Related topics:

- “Related information” on page 285



---

## Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.*

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

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATIONS "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been

estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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

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

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

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

**Related topics:**

- “Trademarks” on page 283

---

## Trademarks

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

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

Microsoft, Windows, and Windows NT are trademarks of Microsoft 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.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, and service names may be trademarks or service marks of others.



---

## Related information

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

- The publications that make up the library for the IBM® TotalStorage™ Virtualization Family SAN Volume Controller (SAN Volume Controller)
- Other IBM publications that relate to the SAN Volume Controller

### SAN Volume Controller library:

Table 29 lists and describes the publications that make up the SAN Volume Controller library. Unless otherwise noted, these publications are available in Adobe portable document format (PDF) on a compact disc (CD) that comes with the SAN Volume Controller. If you need additional copies of this CD, the order number is SK2T-8811. These publications are also available as PDF files from the following Web site:

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

*Table 29. Publications in the SAN Volume Controller library*

| Title                                                                                                    | Description                                                                                                                                                                                     | Order number |
|----------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------|
| <i>IBM TotalStorage Virtualization Family SAN Volume Controller: CIM Agent Developer's Reference</i>     | This reference guide describes the objects and classes in a Common Information Model (CIM) environment.                                                                                         | SC26-7545    |
| <i>IBM TotalStorage Virtualization Family SAN Volume Controller: Command-Line Interface User's Guide</i> | This guide describes the commands that you can use from the SAN Volume Controller command-line interface (CLI).                                                                                 | SC26-7544    |
| <i>IBM TotalStorage Virtualization Family SAN Volume Controller: Configuration Guide</i>                 | This guide provides guidelines for configuring your SAN Volume Controller.                                                                                                                      | SC26-7543    |
| <i>IBM TotalStorage Virtualization Family SAN Volume Controller: Host Attachment Guide</i>               | This guide provides guidelines for attaching SAN Volume Controller to your host system.                                                                                                         | SC26-7563    |
| <i>IBM TotalStorage Virtualization Family SAN Volume Controller: Installation Guide</i>                  | This guide includes the instructions the service representative uses to install the SAN Volume Controller.                                                                                      | SC26-7541    |
| <i>IBM TotalStorage Virtualization Family SAN Volume Controller: Planning Guide</i>                      | This guide introduces the SAN Volume Controller and lists the features you can order. It also provides guidelines for planning the installation and configuration of the SAN Volume Controller. | GA22-1052    |

*Table 29. Publications in the SAN Volume Controller library (continued)*

| Title                                                                                          | Description                                                                                                                                         | Order number |
|------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|--------------|
| <i>IBM TotalStorage Virtualization Family SAN Volume Controller: Service Guide</i>             | This guide describes how to maintain the SAN Volume Controller. It also includes a parts listing.                                                   | SC26-7542    |
| <i>IBM TotalStorage Virtualization Family SAN Volume Controller: Translated Safety Notices</i> | This guide contains the danger and caution notices for the SAN Volume Controller. The notices are shown in English and in numerous other languages. | SC26-7577    |

**Other IBM publications:**

Table 30 lists and describes other IBM publications that contain additional information that is related to the SAN Volume Controller.

*Table 30. Other IBM publications*

| Title                                                         | Description                                                                                                         | Order number |
|---------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------|--------------|
| <i>IBM TotalStorage: Subsystem Device Driver User's Guide</i> | This guide describes the IBM TotalStorage Subsystem Device Driver and how to use it with the SAN Volume Controller. | SC26-7540    |

**Related topics:**

- “Ordering IBM publications” on page 287
- “How to send your comments” on page vii

---

## **Ordering IBM publications**

This topic explains how to order copies of IBM publications and how to set up a profile to receive notifications about new or changed publications.

### **The IBM publications center:**

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

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

[www.ibm.com/shop/publications/order/](http://www.ibm.com/shop/publications/order/)

### **Publications notification system:**

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

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

[www.ibm.com/shop/publications/order/](http://www.ibm.com/shop/publications/order/)

### **Related topics:**

- “Related information” on page 285



---

## Glossary

This glossary includes terms for the IBM TotalStorage Virtualization Family SAN Volume Controller. This glossary includes selected terms and definitions from:

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

The following cross-references are used in this glossary:

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

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

**See also**

Refers the reader to one or more related terms.

**Contrast with**

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

## A

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

## C

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

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

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

**consistent copy.** In a Remote Copy relationship, a copy of a secondary virtual disk that is in the identical state as the primary virtual disk from the viewpoint of a host system if a power failure occurred while I/O activity was in progress and power was later restored.

**consistent copy.** In a Remote Copy relationship, a copy of a secondary virtual disk that is in the identical state as the primary virtual disk from the viewpoint of a host system if a power failure occurred while I/O activity was in progress and power was later restored.

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

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

## D

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

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

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

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

**directed maintenance procedures.** The set of maintenance procedures that can be run for a cluster. These procedures are documented in the *IBM TotalStorage Virtualization Family SAN Volume Controller: Service Guide*.

**disconnected.** In a Remote Copy relationship, pertains to two clusters when they cannot communicate.

**disconnected.** In a Remote Copy relationship, pertains to two clusters when they cannot communicate.

**disk controller.** A device that coordinates and controls the operation of one or more disk drives and synchronizes the operation of the drives with the

operation of the system as a whole. Disk controllers provide the storage that the cluster detects as managed disks (MDisks).

**disk zone.** A zone defined in the SAN fabric in which the SAN Volume Controllers can detect and address the logical units that the disk controllers present.

## E

**error code.** A value that identifies an error condition to a user.

**ESS.** See *IBM TotalStorage Enterprise Storage Server*.

**exclude.** To remove an MDisk from a cluster because of certain error conditions.

**excluded.** In SAN Volume Controller, the status of a managed disk that the cluster has excluded 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.

**FC.** See *fibre channel*.

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

**FlashCopy mapping.** A relationship between two virtual disks.

**FlashCopy relationship.** A depreciated term for *FlashCopy mapping*.

## H

**HBA.** See *host bus adapter*.

**host bus adapter (HBA).** In SAN Volume Controller, an interface card that connects a host bus, such as a peripheral component interconnect 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 LUN mapping. For each host ID, there is a separate mapping of SCSI IDs to virtual disks.

**host zone.** A zone defined in the SAN fabric in which the hosts can and address the SAN Volume Controllers.

## I

**IBM Subsystem Device Driver (SDD).** An IBM pseudo device driver designed to support the multipath configuration environments in IBM products.

**IBM TotalStorage Enterprise Storage Server (ESS).** A storage server product.

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

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

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

**image VDisk.** A virtual disk in which there is a direct block-for-block translation from the managed disk to the virtual disk.

**inconsistent.** In a Remote Copy relationship, pertains to a secondary virtual disk that is being synchronized with the primary.

**inconsistent.** In a Remote Copy relationship, pertains to a secondary virtual disk that is being synchronized with the primary.

**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 to 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 and node relationships that present a common interface to host systems.

**I/O throttling rate.** The maximum rate at which I/Os are accepted for this virtual disk.

**IP.** See *Internet Protocol*.

## L

**LBA.** See *logical block address*.

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

**local/remote fabric interconnect.** The 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 SCSI commands are addressed, for example, a virtual disk or managed disk.

**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 SCSI logical unit that a RAID controller provides and the cluster manages. The managed disk is not visible to host systems on the SAN.

**managed disk group.** A collection of managed disks that together contain all the data for a specified set of virtual disks.

**mapping.** See *FlashCopy mapping*.

**master virtual disk.** In most cases, the virtual disk that contains a production copy of the data and that an application accesses. See also *auxiliary virtual disk*.

**MDisk.** See *managed disk*.

**migration.** See *data migration*.

## N

**node.** One SAN Volume Controller. Each node provides virtualization, cache, and Copy Services to the 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.

## 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

**partnership.** In Remote Copy, the relationship between two clusters. In a cluster partnership, one cluster is defined as the local cluster and the other cluster as the remote cluster.

**partnership.** In Remote Copy, the relationship between two clusters. In a cluster partnership, one cluster is defined as the local cluster and the other cluster as the remote cluster.

**paused.** In SAN Volume Controller, 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.

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

**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.

**PPCR.** See *Peer-to-Peer Remote Copy*.

**primary virtual disk.** In a Remote Copy relationship, the target of write operations issued by the host application.

**primary virtual disk.** In a Remote Copy relationship, the target of write operations issued by the host application.

## Q

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

## R

**RAID.** See *redundant array of independent disks*.

**RAID 1.** A form of storage array in which two or more identical copies of data are maintained on separate media. (S)

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

**RAID 5.** 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)

**RAID 10.** A type of RAID that optimizes high performance while maintaining fault tolerance for up to two failed disk drives by striping volume data across several disk drives and mirroring the first set of disk drives on an identical set.

**redundant SAN.** A SAN configuration in which any one single component might fail, but connectivity between the devices within the SAN is maintained, possibly with degraded performance. This configuration is normally achieved by splitting the SAN into two independent, counterpart SANs. See also *counterpart SAN*.

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

**relationship.** In Remote Copy, the association between a master virtual disk and an auxiliary virtual disk. These virtual disks also have the attributes of a primary or secondary virtual disk. See also *auxiliary virtual disk*, *master virtual disk*, *primary virtual disk*, and *secondary virtual disk*.

**relationship.** In Remote Copy, the association between a master virtual disk and an auxiliary virtual disk. These virtual disks also have the attributes of a primary or secondary virtual disk. See also *auxiliary virtual disk*, *master virtual disk*, *primary virtual disk*, and *secondary virtual disk*.

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

## S

**SAN.** See *storage area network*.

**SAN Volume Controller fibre-channel port fan in.** The number of hosts that can see any one SAN Volume Controller port.

**SCSI.** See *Small Computer Systems Interface*.

**sequential VDisk.** A virtual disk that uses extents from a single managed disk.

**Small Computer System Interface (SCSI).** A standard hardware interface that enables a variety of peripheral devices to communicate with one another.

**secondary virtual disk.** In Remote Copy, the virtual disk in a relationship that contains a copy of data written by the host application to the primary virtual disk.

**secondary virtual disk.** In Remote Copy, the virtual disk in a relationship that contains a copy of data written by the host application to the primary virtual disk.

**Simple Network Management Protocol.** An internet standard for the management of devices.

**SNMP.** See *Simple Network Management Protocol*.

**stand-alone relationship.** In FlashCopy and Remote Copy, relationships that do not belong to a consistency group and that have a null consistency group attribute.

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

**stopped.** The status of a pair of virtual disks 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)

**superuser authority.** The level of access required to add users.

**suspended.** The status of a pair of virtual disks that have a copy relationship that has been temporarily broken because of a problem.

**symmetric virtualization.** A virtualization technique in which the physical storage (RAID arrays) is split into smaller chunks of storage known as extents. These extents are then concatenated together, using various policies, to make virtual disks.

**synchronized.** In Remote Copy, the status condition that exists when both virtual disks of a pair that has a copy relationship contain the same data.

**synchronized.** In Remote Copy, the status condition that exists when both virtual disks of a pair that has a copy relationship contain the same data.

## T

**trigger.** To initiate or reinitiate copying between a pair of virtual disks that have a copy relationship.

## U

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

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

## 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 SAN recognize as a SCSI disk.

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

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

**vital product data (VPD).** Information that uniquely defines system, hardware, software, and microcode elements of a processing system.

## W

**worldwide node name (WWNN).** An identifier for an object that is globally unique. WWNNs are used by fibre channel and other standards.

**WWNN.** See *worldwide node name*.

**WWPN.** See *worldwide port name*.

**worldwide port name (WWPN).** A unique 64-bit identifier associated with a fibre-channel adapter port. It is assigned in an implementation-and protocol-independent manner.



---

# Index

## A

about this guide vii  
accessibility 279  
  keyboard 279  
  shortcut keys 279  
adding  
  storage controllers  
    using the CLI (command-line interface) 180  
    using the SAN Volume Controller  
      Console 118  
advanced functions  
  overview  
    using CLI (command-line interface) 169, 190  
    using SAN Volume Controller  
      Console 111, 126  
Remote Copy  
  using the CLI (command-line interface) 190  
  using the SAN Volume Controller  
      Console 125  
analyzing error logs 135, 195  
audience vii

## B

book  
  about this vii

## C

call-home feature  
  enabling 69  
  overview 69  
CLI (command-line interface)  
  advanced functions 169  
  examples 149  
  getting started 141  
  issuing commands from a PuTTY SSH client system 143  
  preparing SSH client systems 142  
  scenarios 149  
  upgrading software 199  
  using to set cluster features 147  
clusters  
  configuring  
    using the CLI (command-line interface) 190  
  using the SAN Volume Controller  
      Console 126  
creating  
  from the front panel 51  
maintaining 128  
setting  
  features 127, 147  
  time 92, 147  
shutting down 136, 196  
codes  
  configuration events 274

codes (*continued*)  
  events 273  
  information events 273  
command-line interface (CLI)  
  advanced functions 169  
  examples 149  
  getting started 141  
  issuing commands from a PuTTY SSH client system 143  
  preparing SSH clients 142  
  scenarios 149  
  upgrading software 199  
  using to set cluster features 147  
  using to set cluster time 147  
commands  
  CLI (command-line interface)  
    available during the software upgrade process 217  
communications  
  determining between hosts and virtual disks 170  
configuration  
  event codes 274  
configuration rules 41  
  HBAs 43  
  nodes 44  
  power 44  
  switches 44  
configuring  
  clusters 82, 146  
    using the CLI (command-line interface) 190  
    using the SAN Volume Controller  
      Console 126  
  disk controllers 223  
Enterprise Storage Server 224, 230  
FASST Storage Manager 224, 230  
FASST Storage Server 224, 228  
master console 57  
PuTTY 64  
remote support 66  
secure shell (SSH) 59  
SSH (secure shell) 59  
  Tivoli Storage Manager 64  
consistency group, Remote Copy 38  
consistency groups, FlashCopy  
  deleting 110  
  starting 110  
  stopping 110  
console  
  master  
    upgrading software 72  
SAN Volume Controller  
  banner area 78  
  layout 78  
  portfolio 79  
  starting 77  
  task bar 78  
  work area 79

controllers  
  adding  
    using the CLI (command-line interface) 180  
    using the SAN Volume Controller  
      Console 118  
  removing  
    using the CLI (command-line interface) 181  
    using the SAN Volume Controller  
      Console 119  
conventions  
  numbering vii  
counts  
  errors 207  
creating  
  clusters  
    from the front panel 51  
    from the SAN Volume Controller  
      Console 81  
FlashCopy  
  consistency groups 106  
  mappings 106, 165, 166  
hosts 104  
managed disk groups 101  
SSH keys 60  
virtual disk-to-host mappings 105, 163

## D

deleting  
  FlashCopy  
    mappings 109  
    nodes 126, 190  
determining  
  communications between hosts and virtual disks 170  
disability 279  
discovering  
  managed disks 155, 159  
disks  
  migrating 123, 187  
disruptive software upgrade  
  using the CLI (command-line interface) 201

## E

e-mail  
  setting up 71, 268, 269  
enabling  
  call-home feature 69  
  cluster maintenance procedure 128  
errors  
  counts 207  
events  
  codes 273  
    configuration 274  
    information 273

events (*continued*)  
  setting up an action plan for 267  
examples  
  using the CLI (command-line interface) 149  
  using the SAN Volume Controller Console 95  
extents  
  migrating  
    using the CLI (command-line interface) 185

## F

fabric 3  
FC. See Fibre Channel (FC) 3  
features  
  setting  
    using the CLI (command-line interface) 147  
    using the SAN Volume Controller Console 127  
Fibre Channel (FC) 3  
FlashCopy  
  mappings 164

## G

gateways 3  
general cluster properties  
  viewing 93, 148  
getting started  
  using SAN Volume Controller 77, 241  
  using the CLI (command-line interface) 141  
  using the command-line interface (CLI) 141  
glossary 289  
guide  
  about this vii  
  who should read vii

## H

hubs 3

## I

IBM Director  
  configuring 267  
  overview 67, 267  
  starting 68  
image-mode virtual disks  
  converting to managed mode  
    using CLI (command-line interface) 189  
  using SAN Volume Controller Console 124  
information  
  center 285  
  event codes 273  
installing  
  SAN Volume Controller 244, 253  
  verification 256

IP addresses  
  modifying 133, 193  
issuing  
  CLI commands 143

## K

keyboard 279  
  shortcut keys 279

## L

language 134, 194  
link 3  
listing  
  dump files 134, 193  
  log files 134, 193

## M

maintaining  
  passwords 129, 148, 192  
  SSH keys 192  
maintenance procedures  
  clusters 128  
managed mode virtual disks  
  converting from image mode  
    using SAN Volume Controller Console 124  
    using the CLI (command-line interface) 189  
mappings, FlashCopy  
  deleting 109  
  starting 109  
  stopping 109  
master console  
  configuring 57  
  overview 57  
  upgrading software 72  
measurements vii  
migrating  
  extents  
    using the CLI (command-line interface) 185  
migration 227  
monitoring  
  software upgrades, automatic 209, 211

## N

node 3  
nodes  
  adding 96, 150  
  deleting 126  
  status 9  
  viewing  
    general details 100, 154  
notices 281

## O

object descriptions 8  
operating over long distances 238  
ordering publications 287

overview  
  advanced functions  
    using the CLI (command-line interface) 169, 190  
    using the SAN Volume Controller Console 111, 126  
call-home feature 69  
creating a cluster 81  
IBM Director 67  
SSH (secure shell) 213  
zoning 67, 233

## P

plink utility  
  running 144  
preinstalled software  
  recovering from installation failures 219  
preparing  
  SSH client system  
    overview 142  
    to issue CLI commands 142  
problems  
  counting 207  
public SSH keys  
  storing 129  
publications  
  ordering 287  
PuTTY 64  
  configuring 64  
  issuing CLI commands from 143  
  running the plink utility 144

## R

related information 285  
relationships, Remote Copy  
  overview 38  
Remote Copy  
  overview 37, 38  
    using the CLI (command-line interface) 190  
    using the SAN Volume Controller Console 125  
  partnerships 37  
  zoning considerations 235  
remote support  
  configuring 66  
removing  
  storage controllers  
    using the CLI (command-line interface) 181  
    using the SAN Volume Controller Console 119  
requirements 241, 242, 243  
routers 3  
running  
  PuTTY plink utility 144

## S

SAN Volume Controller  
  advanced functions 111  
Console  
  banner area 78

SAN Volume Controller (*continued*)  
Console (*continued*)  
examples 95  
layout 78  
portfolio 79  
post installation tasks 257  
scenarios 95  
starting 77  
task bar 78  
using to create a cluster 81  
using to set cluster features 127  
using to set cluster time 92  
work area 79  
removing 259  
SAN. See storage area network (SAN) 3  
scenarios  
    using the CLI (command-line interface) 149  
    using the SAN Volume Controller Console 95  
secure shell (SSH) 58  
    client system  
        issuing CLI commands from 143  
        overview 142  
        preparing to issue CLI commands 142  
    configuring 59  
    creating keys 60  
    keys  
        generating 60  
        storing 129  
    overview 213  
security  
    overview 58  
setting  
    action plan for events 267  
cluster features  
    using the CLI (command-line interface) 147  
    using the SAN Volume Controller Console 127  
cluster time  
    using the CLI (command-line interface) 147  
    using the SAN Volume Controller Console 92  
e-mail account 71, 268, 269  
features  
    using the CLI (command-line interface) 147  
    using the SAN Volume Controller Console 127  
time  
    using the CLI (command-line interface) 147  
    using the SAN Volume Controller console 92  
    traps 265  
settings  
    error notification 131, 193  
shortcut keys 279  
SNMP  
    setting up traps 265  
software  
    upgrading 203, 215

software, upgrading  
    available CLI (command-line interface) commands 217  
disruptive  
    using the CLI (command-line interface) 201  
    master console 72  
    using the CLI (command-line interface) 199  
SSH (secure shell) 58  
    client system  
        issuing CLI commands from 143  
        overview 142  
        preparing to issue CLI commands 142  
    configuring 59  
    creating 60  
    keys  
        generating 60  
        storing 129  
    overview 213  
starting  
    FlashCopy  
        consistency groups 110  
        mappings 109  
    IBM Director 68  
    Tivoli Storage Manager 65  
stopping  
    FlashCopy  
        mappings 109  
    Remote Copy  
        consistency groups 110  
storage area network (SAN) 3  
storage controllers  
    adding  
        using the CLI (command-line interface) 180  
    using the SAN Volume Controller Console 118  
    removing  
        using the CLI (command-line interface) 181  
    using the SAN Volume Controller Console 119  
storing  
    public SSH keys 129  
strategy  
    software upgrade  
        using the CLI (command-line interface) 199  
support  
    configuring remote 66  
switches 3  
    operating over long distances 238  
    zoning 233  
synchronous copy  
    overview 37

Tivoli SAN Manager  
    configuring 64  
    starting 65  
trademarks 283

## U

uninterruptible power supply  
    introducing 15  
upgrading software  
    command-line interface  
        commands available 217  
    disruptive  
        using the CLI (command-line interface) 201  
    master console 72  
    strategy  
        using the CLI (command-line interface) 199  
using  
    object classes and instances 271

## V

VDisks (virtual disks)  
    converting  
        from image mode to managed mode 124, 189  
    creating 103, 159  
    migrating 188  
viewing  
    clusters  
        feature logs 135, 195  
virtual disks (VDisks)  
    converting  
        from image mode to managed mode 124, 189  
virtualization  
    overview 4

## W

who should read this guide vii

## Z

zoning  
    considerations for Remote Copy 235  
    overview 67, 233  
    switches 233

## T

time  
    setting  
        using the CLI (command-line interface) 147  
    using the SAN Volume Controller Console 92



---

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

IBM® TotalStorage® Virtualization Family  
SAN Volume Controller™  
Configuration Guide  
Version 1 Release 1

Publication No. SC26-7543-00

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

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

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

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

**Please tell us how we can improve this book:**

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

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

---

Name

---

Address

---

Company or Organization

---

Phone No.

**Readers' Comments — We'd Like to Hear from You**  
SC26-7543-00



Cut or Fold  
Along Line

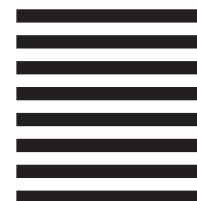
Fold and Tape

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  
RCF Processing Department  
M86/050-3  
5600 Cottle Road  
San Jose, CA 95193-0001



Fold and Tape

Please do not staple

Fold and Tape

SC26-7543-00

Cut or Fold  
Along Line





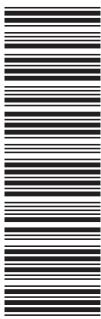
Part Number: 64P7905

Printed in U.S.A.

SC26-7543-00



(1P) P/N: 64P7905



Spine information:

IBM® TotalStorage® Virtualization  
Family  
SAN Volume Controller™



Version 1  
Release 1

SAN Volume Controller Configuration Guide