

IBM Storwize V7000 Unified

Problem Determination Guide



Note

Before using this information and the product it supports, read the following information:

- The general information in “Notices” on page 439
- The information in the “Safety and environmental notices” on page xi
- The information in the *IBM Environmental Notices and User Guide* (provided on a DVD)

This edition applies to IBM Storwize V7000 Unified and to all subsequent releases and modifications until otherwise indicated in new editions.

This edition replaces GA32-1057-12.

© **Copyright IBM Corporation 2011, 2017.**

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

Contents

Figures	vii
----------------	------------

Tables	ix
---------------	-----------

Safety and environmental notices xi

Safety notices and labels	xi
Caution notices for the Storwize V7000 Unified	xii
Danger notices for Storwize V7000 Unified	xv
Special caution and safety notices	xviii
General safety	xviii
Handling static-sensitive devices	xix
Sound pressure	xix
Environmental notices	xix

About this guide xxi

Who should use this guide	xxi
Storwize V7000 Unified library and related publications	xxi
How to order IBM publications	xxiv
Related websites	xxiv
Sending your comments	xxiv
How to get information, help, and technical assistance	xxv
What's new	xxvii

Chapter 1. Storwize V7000 Unified hardware components 1

Components in the front of the 2073-700 file module	1
Components in the front of the 2073-720 file module	2
Components in the rear of the 2073-700 file module	3
Components in the rear of the 2073-720 file module	3
Components in the front of the enclosure	4
Drives for control enclosures	4
Drive indicators for control enclosures	5
Enclosure end cap indicators	7
Components in the rear of the enclosure	8
Power supply units for control enclosures	9
Power supply units for expansion enclosures	13
Storwize V7000 2076-524 node canister ports and indicators	15
Node canister ports and indicators	25
Expansion canister ports and indicators	32
Storwize V7000 Gen2 expansion canister SAS ports and indicators	32
Storwize V7000 Gen1 expansion canister SAS ports and indicators	34
Storwize V7000 Gen2 expansion canister LEDs	34
Expansion canister LEDs	35

Chapter 2. Best practices for troubleshooting 37

Record the access information	37
Follow proper power management procedures	38

Follow proper Storwize V7000 Gen2 power management procedures	39
Follow proper power management procedures	39
Set up event notifications	39
Set up inventory reporting	40
Back up your data	40
Manage your spare and failed drives	40
Resolve alerts in a timely manner	40
Keep your software up to date	41
Keep your records up to date	41
Keep your Storwize V7000 Gen2 records up to date	42
Keep your Storwize V7000 Unified Gen1 records up to date	42
Subscribe to support notifications	42
Know your IBM warranty and maintenance agreement details	43
How to get information, help, and technical assistance	43

Chapter 3. Getting started troubleshooting 47

Installation troubleshooting	48
Problems with initial setup	48
Installation error codes	51
Problems reported by the CLI commands during software configuration	59
Management GUI wizard failure	60
GUI access issues	60
Health status and recovery	62
Connectivity issues for the 2073-720	63
Host to file modules connectivity	63
Ethernet connectivity between file modules	65
Ethernet connectivity from file modules to the control enclosure	68
Fibre Channel connectivity between file modules and control enclosure	72
Understanding LED hardware indicators	77
File module hardware indicators for 2073-720	77
Enclosure hardware indicators	87
Management GUI interface	87
When to use the management GUI	88
Accessing the Storwize V7000 Unified management GUI	89
Diagnosing and resolving problems with fix procedures	90

Chapter 4. File module 91

General file module procedures	91
Rebooting a file module	91
Removing a file module to perform a maintenance action	91
Removing and replacing file module components	94
Resolving hard disk drive problems	96
Monitoring memory usage on a file module	117

Errors and messages	117	Understanding the medium errors and bad blocks	235
Understanding error codes	117	Resolving a problem	236
Understanding event IDs	120	Start here: Use the management GUI	
File module hardware problems	121	recommended actions	236
Removing and replacing parts for the 2073-720	122	Problem: Management IP address unknown	237
How to reset/reboot server IMM interface.	182	Problem: Unable to connect to the management	
File module software problems	182	GUI	237
Logical devices and physical port locations for a		Problem: Unable to log on to the management	
2073-720 file module	182	GUI	240
Management node role failover procedures	183	Problem: Cannot initialize or create a clustered	
Checking CTDB health	187	system	241
Checking the GPFS file system mount on each file		Problem: Node canister service IP address	
module	189	unknown.	242
Identifying created and mounted file system		Problem: Cannot connect to the service assistant	245
mounts	189	Problem: Management GUI or service assistant	
Resolving problems with missing mounted file		does not display correctly	246
systems	190	Problem: A node canister has a location node	
Resolving stale NFS file systems	191	error	246
Checking user and server authentication issues	191	Problem: SAS cabling not valid	246
Resolving the “Missing SRV record in DNS”		Problem: New expansion enclosure not detected	248
error	192	Problem: Control enclosure not detected	249
If “netgroup” functionality with NIS or LDAP is		Problem: Mirrored volume copies no longer	
not working.	192	identical	249
Possible misconfiguration on the Storwize V7000		Problem: Command file not processed from USB	
Unified system	192	flash drive	249
Trouble accessing exports when authentication		Procedure: Resetting superuser password	250
server and client Storwize V7000 Unified		Procedure: Identifying which enclosure or	
configurations are correct	193	canister to service	252
Resolving access failures on an Storwize V7000		Procedure: Checking the status of your system	254
Unified system with a subordinate ID map role	193	Procedure: Getting node canister and system	
Checking client access	194	information using the service assistant	254
Checking network interface availability.	195	Procedure: Getting node canister and system	
Recovering a GPFS file system	196	information using a USB flash drive.	255
Resolving an ANS1267E error	197	Procedure: Understanding the system status	
Resolving issues reported by lsheal th	197	using the LEDs.	255
Error for “MGMTNODE_REPL_STATE ERROR		Procedure: Finding the status of Ethernet	
DATABASE_REPLICATION_FAILED”	197	connections	267
Resolving network errors	198	Procedure: Finding the status of Storwize V7000	
Resolving full condition for GPFS file system.	200	Gen2 SAS connections	269
Analyzing GPFS logs.	200	Procedure: Removing system data from a node	
Synchronizing time on the file modules	201	canister	270
Chapter 5. Control enclosure.	203	Procedure: Deleting a system completely	270
Storwize V7000 system interfaces.	203	Procedure: Fixing node errors	271
Service assistant interface	203	Procedure: Changing the service IP address of a	
Storage system command-line interface.	205	node canister	271
Service command-line interface	206	Procedure: Initializing a clustered system with a	
USB flash drive and Initialization tool interface	206	USB flash drive without using the initialization	
Starting statistics collection	215	tool.	272
Event reporting.	226	Procedure: Initializing a clustered system using	
Understanding events	226	the service assistant	273
Event notifications.	228	Procedure: Accessing the service assistant from	
Power-on self-test	229	the technician port	275
Understanding event codes.	229	Procedure: Accessing a Storwize V7000 Gen1	
Understanding the error codes	229	canister using a directly attached Ethernet cable.	276
Viewing logs and traces	230	Problem: Reseating a node canister	277
Battery operation for the control enclosure	230	Procedure: Removing a Storwize V7000 Gen2	
Battery operation for Storwize V7000 Gen2		node canister	279
control enclosures	230	Procedure: Powering off your system	280
Battery operation for Storwize V7000 Unified		Procedure: Powering on the Storwize V7000	
Gen1 control enclosures	232	Gen2 system.	281

Procedure: Powering off a Storwize V7000 Gen2 control enclosure	282
Procedure: Powering off a Storwize V7000 Gen2 node canister	283
Procedure: Collecting information for support	283
Procedure: Rescuing node canister software from another node (node rescue)	284
Procedure: FCoE host-linking	285
Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister	286
Procedure: Understanding Storwize V7000 Gen2 volume dependencies	287
Storwize V7000 replaceable units	288
Storwize V7000 2076-524 Gen2 replaceable units	288
Storwize V7000 2076-1xx and 2076-3xx Gen1 replaceable units	292
Replacing parts.	295
Preparing to remove and replace parts	295
Replacing a node canister	296
Replacing a fan module	298
Replacing an expansion canister	299
Replacing an SFP transceiver	303
Replacing a power supply unit for a control enclosure.	306
Replacing a power supply unit for an expansion enclosure	311
Replacing the battery in a node canister	316
Replacing a battery in a power supply unit	317
Releasing the cable retention bracket	321
Replacing a 3.5 inch drive assembly or blank carrier	321
Replacing a 2.5 inch drive assembly or blank carrier	325
Replacing enclosure end caps	329
Replacing a SAS cable to an expansion enclosure.	330
Replacing a control enclosure chassis	333
Replacing an expansion enclosure chassis	344
Replacing a Storwize V7000 Gen2 enclosure midplane.	351
Replacing the support rails.	363
Replacing node canister memory modules.	370
Replacing a host interface adapter	372
Replacing a CMOS battery	375
General storage system procedures	376
Procedure: SAN problem determination	376
iSCSI performance analysis and tuning.	377
Fibre Channel link failures	379
Ethernet iSCSI host-link problems	380
Recover system procedure	381
When to run the recover system procedure	383
Fix hardware errors	383
Removing system information for node canisters with error code 550 or error code 578 using the service assistant	385
Running system recovery using the service assistant	385
Recovering from offline volumes using the CLI	387
What to check after running the system recovery	388
Backing up and restoring the system configuration	390

Backing up the system configuration using the CLI.	391
Restoring the system configuration	393
Deleting backup configuration files using the CLI.	395

Chapter 6. Call home and remote support 397

Testing a call home connection	397
Establishing an AOS connection	397
Establishing a lights-out AOS connection	397
Establishing a lights-on AOS connection	398

Chapter 7. Recovery procedures . . . 401

User ID and system access	401
Accessing a file module as root	401
Recovering from losing the root password.	401
Resetting the NAS ssh key for configuration communications	402
Working with NFS clients that fail to mount NFS shares after a client IP change	403
Working with file modules that report a stale NFS file handle.	403
File module-related issues	404
Restoring System x firmware (BIOS) settings	404
Recovering from file systems that are offline after the volumes came back online	406
Recovering from a multipath event	406
Recovering from an NFSD service error	407
Recovering from an SCM error	407
Recovering from an httpd service error.	408
Recovering from an sshd_data service error	408
Recovering from an sshd_int service error.	409
Recovering from an sshd_mgmt service error	409
Recovering from an sshd_service service error	409
Recovering from the 1710 bus error due to the /var directory being full.	410
Control enclosure-related issues	410
Recovering when file volumes come back online	410
Recovering when a file volume does not come back online	411
Recovering from offline compressed volumes	411
Recovering from a 1001 error code	412
Restoring data	415
Restoring asynchronous data	415
Restoring IBM Spectrum Protect data	415
Upgrade recovery	417

Chapter 8. Troubleshooting compressed file systems 429

Recovery procedure: Increase capacity of the storage pool.	430
Recovery procedure: Adding additional capacity for offline compressed file systems	431
Monitoring file system compression	434

**Appendix. Accessibility features for
IBM Storwize V7000 Unified 437**

Notices 439

Trademarks	441
Electronic emission notices	441
Federal Communications Commission (FCC) statement.	441
Industry Canada compliance statement.	442
Australia and New Zealand Class A Statement	442
European Union Electromagnetic Compatibility Directive	442

Germany Electromagnetic Compatibility Directive	442
People's Republic of China Class A Statement	443
Taiwan Class A compliance statement	444
Taiwan Contact Information	444
Japan VCCI Council Class A statement	444
Japan Electronics and Information Technology Industries Association Statement	444
Korean Communications Commission Class A Statement	445
Russia Electromagnetic Interference Class A Statement	445

Figures

1. Front view of 2073-700 file module	1
2. Front view of 2073-720 file module	2
3. 2073-720 file module advanced operator panel	2
4. Rear view of 2073-700 file module	3
5. Rear view of 2073-720 file module	3
6. Storwize V7000 Gen1 Drives on a 12-drive enclosure.	4
7. Storwize V7000 Gen2 Large form factor horizontal drive	5
8. Storwize V7000 Gen1 Drives on a 24-drive enclosure.	5
9. Storwize V7000 Gen2 Small form factor vertical drive	5
10. LED indicators on a single 3.5 inch drive	6
11. LED indicators on a single 2.5 inch drive	6
12. Storwize V7000 Gen1 12 drives and two end caps	7
13. Storwize V7000 Gen1 Left enclosure end cap	8
14. Rear view of a model 2076-112 or a model 2076-124 control enclosure	9
15. Rear view of a model 2076-312 or a model 2076-324 control enclosure	9
16. Rear view of a model 2076-212 or a model 2076-224 expansion enclosure	9
17. Rear view of a Storwize V7000 Gen2 control enclosure	10
18. LEDs on the power supply units of the control enclosure	12
19. Rear view of a Storwize V7000 Gen2 expansion enclosure	14
20. LEDs on the power supply units of the expansion enclosure.	15
21. Storwize V7000 2076-524 node canister ports	16
22. Storwize V7000 2076-524 node canister indicators	16
23. USB ports on the Storwize V7000 Gen2 node canister	22
24. Fibre Channel ports and indicators.	23
25. Example of installed 10 Gbps Fibre Channel over Ethernet/iSCSI host interface adapters.	24
26. 10 Gbps Fibre Channel over Ethernet/iSCSI host interface adapter ports	24
27. 10 Gbps Fibre Channel over Ethernet/iSCSI host interface adapter indicator LEDs	25
28. Fibre Channel ports on the node canisters	26
29. LEDs on the Fibre Channel ports	26
30. USB ports on the node canisters.	28
31. Ethernet ports on the 2076-112 and 2076-124 node canisters.	29
32. 10 Gbps Ethernet ports on the 2076-312 and 2076-324 node canisters	29
33. SAS ports on the node canisters.	30
34. LEDs on the node canisters	31
35. SAS ports and LEDs at rear of expansion canister	33
36. SAS ports and LEDs in rear of expansion enclosure	34
37. Expansion canister LEDs	35
38. LEDs on the expansion canisters	36
39. File module Ethernet connections.	64
40. File module Ethernet connections.	66
41. Connecting the file modules to the Storwize V7000 Gen1 control enclosure using Fibre Channel cables	73
42. Connecting the file modules to a Storwize V7000 Gen2 control enclosure that has a Fibre Channel interface adapter in PCI slot 2 of each node canister	74
43. Locations of the power-supply LEDs	86
44. Selecting a file module to display node status	99
45. Displaying node status	100
46. Example that shows that mirroring is re-synchronizing	102
47. Example that shows that a drive is not synchronized.	103
48. Example that shows that the mirror is not created.	104
49. Example of a SMART error	105
50. Removing the cover	123
51. Installing the cover	124
52. Removing the bezel	126
53. Installing the bezel.	127
54. Configuration cable	130
55. Power cable connection	130
56. Hard disk drive cable connection	131
57. Removing the battery	133
58. Installing the battery	136
59. Removing the air baffle	138
60. Installing the air baffle	139
61. Removing the fan bracket	140
62. Installing the fan bracket.	141
63. Installing a PCI riser-card assembly	143
64. Removing a PCI adapter from a PCI riser-card assembly	144
65. Inserting the adapter into the PCI connector	145
66. Removing a hot-swap hard disk drive	149
67. Installing a hot-swap hard disk drive	150
68. Locations of the DIMM connectors on the system board	154
69. Removing a hot-swap ac power supply	159
70. Removing the hot-swap drive backplane	165
71. Installing the hot-swap drive backplane	166
72. Management GUI showing CTDB status for both file modules	188
73. LEDs on the power supply units of the control enclosure	263
74. LEDs on the node canisters	265
75. Removing a node canister	280
76. Power LEDs on a node canister	281
77. Node canister LEDs	281
78. Replacing the canister cover	287

79. Rear of node canisters that shows the handles.	298	96. Unlocking the 3.5 inch drive	324
80. Removing the canister from the enclosure	298	97. Removing the 3.5 inch drive	325
81. Removing and replacing the Storwize V7000 Gen2 expansion canister	301	98. Unlocking and removing a 2.5-inch drive from its slot	327
82. Rear of expansion canisters that shows the handles.	302	99. Installing and locking a 2.5-inch drive into its slot	327
83. Removing the canister from the enclosure	302	100. Unlocking the 2.5 inch drive	328
84. SFP transceiver	304	101. Removing the 2.5 inch drive	329
85. SFP transceiver	305	102. Proper orientation for SAS cable connector	332
86. Removing the power supply unit (left side of enclosure).	307	103. SAS cable	333
87. Directions for lifting the handle on the power supply unit	310	104. Bottom enclosure screws	356
88. Using the handle to remove a power supply unit.	310	105. Right-side enclosure screws	357
89. Removing the power supply unit from the left side of the expansion enclosure	312	106. Left-side enclosure screws	357
90. Directions for lifting the handle on the power supply unit	315	107. Angled midplane assembly	358
91. Using the handle to remove a power supply unit.	315	108. Removing a vertical style hard disk drive	360
92. Opening latching arms to disconnect a Storwize V7000 Gen2 node canister battery	317	109. Removing a horizontal style hard disk drive	361
93. Removing the battery from the control enclosure power-supply unit	320	110. Removing the screws of an expansion enclosure assembly	362
94. Unlocking and removing a 3.5-inch drive from its slot	323	111. Opening rear hinge bracket of mounting rail	365
95. Installing and locking a 3.5-inch drive into its slot	323	112. Compressing rail for removal from rack	366
		113. Opening rear hinge bracket of mounting rail	368
		114. Compressing rail for removal from rack	369
		115. Removing a rail assembly from a rack cabinet	370
		116. Installing a Storwize V7000 2076-524 node canister memory module.	371
		117. Removing the host interface adapter	373
		118. Installing the host interface adapter	374
		119. Replacing a CMOS Gen2 battery	376

Tables

1. IBM websites for help, services, and information	xxi
2. Storwize V7000 Unified library	xxii
3. IBM documentation and related websites	xxiv
4. IBM websites for help, services, and information	xxv
5. Drive LEDs	6
6. Storwize V7000 Gen1 LED descriptions for left enclosure end cap	8
7. Storwize V7000 Unified Gen1 model numbers	10
8. Storwize V7000 Unified Gen2 model numbers	10
9. Power supply unit LEDs in the rear of the control enclosure	12
10. Storwize V7000 Unified Gen1 model numbers	13
11. Storwize V7000 Unified Gen2 model numbers	13
12. Power supply LEDs	14
13. Power supply unit LEDs in the rear of the expansion enclosure	15
14. Storwize V7000 2076-524 SAS ports 1 and 2 LEDs	17
15. Storwize V7000 2076-524 battery status LEDs	19
16. Storwize V7000 2076-524 node canister system status LEDs	20
17. Fibre Channel host interface adapter port-state LEDs	23
18. Storwize V7000 2076-524 host interface adapter LED states and meanings	25
19. Fibre Channel port LED locations on canister 1	26
20. Fibre Channel port LED status descriptions	27
21. 1 Gbps Ethernet port LEDs	29
22. 10 Gbps Ethernet port LEDs	30
23. SAS port LEDs on the node canister	30
24. Node canister LEDs	31
25. SAS port LEDs on the expansion canister	33
26. SAS port LEDs on the expansion canister	34
27. Expansion canister LED descriptions	35
28. Expansion canister LEDs	36
29. Access information for your system	37
30. Storwize V7000 Unified Gen1 model numbers	41
31. Storwize V7000 Unified Gen2 model numbers	42
32. IBM websites for help, services, and information	43
33. Installation error code actions	51
34. Error messages and actions	52
35. CLI command problems	59
36. Ethernet connections available with the file modules	64
37. Ethernet connections available with the file modules	66
38. How to connect Fibre Channel cables from file modules to the control enclosure	75
39. Error code port location mapping	75
40. Fibre Channel cabling from the file module to the control enclosure	76
41. LED states and associated actions. For the Fibre Channel adapters on the file module check the amber LED lights next to the port	76
42. Fibre Channel connection on the node canister LED state and associated actions	77
43. LED indicators, corresponding problem causes, and corrective actions	80
44. Status of volume	100
45. State of drives	101
46. SMART ASC/ASCQ error codes and messages	106
47. Error code information	117
48. Originating role information	118
49. Originating file module and file module specific hardware code – Code 0, 2, 4	118
50. Originating file module specific software code – Code 1, 3, 5	119
51. Storage enclosure hardware code – Code 6	119
52. Error code break down	120
53. Components identified as customer replaceable units (CRUs) and field replaceable units (FRUs)	122
54. DIMM slots populated with the memory RDIMM	155
55. Default logical devices and physical port locations for a 2073-720 file module	183
56. Hostname and service IP reference	184
57. Statistics collection for individual nodes	215
58. Statistic collection for volumes for individual nodes	216
59. Statistic collection for volumes that are used in Metro Mirror and Global Mirror relationships for individual nodes	217
60. Statistic collection for node ports	217
61. Statistic collection for nodes	218
62. Cache statistics collection for volumes and volume copies	219
63. Statistic collection for volume cache per individual nodes	223
64. XML statistics for an IP Partnership port	224
65. Description of data fields for the event log	227
66. Notification levels	229
67. Storwize V7000 Unified Gen1 model numbers	230
68. Storwize V7000 Unified Gen2 model numbers	230
69. Bad block errors	235
70. Storwize V7000 Unified Gen1 model numbers	238
71. Storwize V7000 Unified Gen2 model numbers	238
72. Storwize V7000 Unified Gen1 model numbers	241
73. Storwize V7000 Unified Gen2 model numbers	241
74. Storwize V7000 Unified Gen1 model numbers	243
75. Storwize V7000 Unified Gen2 model numbers	243
76. Default service IP addresses	244
77. Default service IP addresses	245
78. Storwize V7000 Unified Gen1 model numbers	247
79. Storwize V7000 Unified Gen2 model numbers	247
80. Storwize V7000 Unified Gen1 model numbers	250

81. Storwize V7000 Unified Gen2 model numbers	251	103. Storwize V7000 Unified Gen1 model numbers	288
82. Storwize V7000 Unified Gen1 model numbers	252	104. Storwize V7000 Unified Gen2 model numbers	288
83. Storwize V7000 Unified Gen2 model numbers	252	105. Control enclosure replaceable units	289
84. Storwize V7000 Unified Gen2 model numbers	253	106. Expansion enclosure replaceable units	290
85. Storwize V7000 Unified Gen1 model numbers	256	107. Drive replaceable units	290
86. Storwize V7000 Unified Gen2 model numbers	256	108. Cable replaceable units	291
87. LED state descriptions used in the Storwize V7000 2076-524 enclosure	257	109. Replaceable units	292
88. Understanding the power supply unit LEDs	257	110. Storwize V7000 Unified Gen1 model numbers	333
89. Understanding the node canister status LEDs	258	111. Storwize V7000 Unified Gen2 model numbers	334
90. Understanding the node canister battery status LEDs	261	112. Storwize V7000 Unified Gen1 model numbers	344
91. Power-supply unit LEDs	263	113. Storwize V7000 Unified Gen2 model numbers	345
92. Power LEDs	264	114. Storwize V7000 Unified Gen1 model numbers	363
93. System status and fault LEDs	265	115. Storwize V7000 Unified Gen2 model numbers	363
94. Control enclosure battery LEDs	266	116. Replacing host interface adapters in two control enclosures concurrently	375
95. Storwize V7000 Unified Gen1 model numbers	267	117. Files created by the backup process	392
96. Storwize V7000 Unified Gen2 model numbers	267	118. Recovering from offline compressed volumes.	412
97. Storwize V7000 Unified Gen1 model numbers	273	119. Upgrade error codes from using the applysoftware command and recommended actions	417
98. Storwize V7000 Unified Gen2 model numbers	274	120. Upgrade error codes and recommended actions	420
99. Storwize V7000 Unified Gen1 model numbers	277	121. Capacity failure scenarios	429
100. Storwize V7000 Unified Gen2 model numbers	277		
101. Storwize V7000 Unified Gen1 model numbers	284		
102. Storwize V7000 Unified Gen2 model numbers	284		

Safety and environmental notices

Review the safety notices, environmental notices, and electronic emission notices for IBM® Storwize® V7000 Unified before you install and use the product.

Suitability for telecommunication environment: This product is not intended to connect directly or indirectly by any means whatsoever to interfaces of public telecommunications networks.

Here are examples of a caution and a danger notice:

CAUTION:

A caution notice indicates the presence of a hazard that has the potential of causing moderate or minor personal injury. (C001)

DANGER

A danger notice indicates the presence of a hazard that has the potential of causing death or serious personal injury. (D002)
--

To find the translated text for a caution or danger notice:

1. Look for the identification number at the end of each caution notice or each danger notice. In the preceding examples, the numbers (C001) and (D002) are the identification numbers.
2. Locate the *IBM Systems Safety Notices* with the user publications that were provided with the Storwize V7000 Unified hardware.
3. Find the matching identification number in the *IBM Systems Safety Notices*. Then, review the topics concerning the safety notices to ensure that you are in compliance.
4. Optionally, read the multilingual safety instructions on the Storwize V7000 Unified website. Go to www.ibm.com/storage/support/storwize/v7000/unified, search for Storwize V7000 Unified, and click the documentation link.

Safety notices and labels

Review the safety notices and safety information labels before using this product.

To view a PDF file, you need Adobe Acrobat Reader. You can download it at no charge from the Adobe website:

www.adobe.com/support/downloads/main.html

IBM Systems Safety Notices

This publication contains the safety notices for the IBM Systems products in English and other languages. Anyone who plans, installs, operates, or services the system must be familiar with and understand the safety notices. Read the related safety notices before you begin work.

Note: The *IBM System Safety Notices* document is organized into two sections. The danger and caution notices without labels are organized alphabetically by language

in the “Danger and caution notices by language” section. The danger and caution notices that are accompanied with a label are organized by label reference number in the “Labels” section.

Note: You can find and download the current *IBM System Safety Notices* by searching for Publication number **G229-9054** in the IBM Publications Center.

The following notices and statements are used in IBM documents. They are listed in order of decreasing severity of potential hazards.

Danger notice definition

A special note that emphasizes a situation that is potentially lethal or extremely hazardous to people.

Caution notice definition

A special note that emphasizes a situation that is potentially hazardous to people because of some existing condition, or to a potentially dangerous situation that might develop because of some unsafe practice.

Note: In addition to these notices, labels might be attached to the product to warn of potential hazards.

Finding translated notices

Each safety notice contains an identification number. You can use this identification number to check the safety notice in each language.

To find the translated text for a caution or danger notice:

1. In the product documentation, look for the identification number at the end of each caution notice or each danger notice. In the following examples, the numbers (D002) and (C001) are the identification numbers.

DANGER

A danger notice indicates the presence of a hazard that has the potential of causing death or serious personal injury. (D002)

CAUTION:

A caution notice indicates the presence of a hazard that has the potential of causing moderate or minor personal injury. (C001)

2. After you download the *IBM System Safety Notices* document, open it.
3. Under the language, find the matching identification number. Review the topics about the safety notices to ensure that you are in compliance.

Note: This product was designed, tested, and manufactured to comply with IEC 60950-1, and where required, to relevant national standards that are based on IEC 60950-1.

Caution notices for the Storwize V7000 Unified

Ensure that you understand the caution notices for Storwize V7000 Unified.

Use the reference numbers in parentheses at the end of each notice, such as (C003) for example, to find the matching translated notice in *IBM Systems Safety Notices*.

CAUTION:

The battery contains lithium. To avoid possible explosion, do not burn or charge the battery.

Do not: Throw or immerse into water, heat to more than 100°C (212°F), repair or disassemble. (C003)

CAUTION:

Electrical current from power, telephone, and communication cables can be hazardous. To avoid personal injury or equipment damage, disconnect the attached power cords, telecommunication systems, networks, and modems before you open the machine covers, unless instructed otherwise in the installation and configuration procedures. (26)

CAUTION:

- Do not install a unit in a rack where the internal rack ambient temperatures will exceed the manufacturer's recommended ambient temperature for all your rack-mounted devices.
- Do not install a unit in a rack where the air flow is compromised. Ensure that air flow is not blocked or reduced on any side, front, or back of a unit used for air flow through the unit.
- Consideration should be given to the connection of the equipment to the supply circuit so that overloading of the circuits does not compromise the supply wiring or overcurrent protection. To provide the correct power connection to a rack, refer to the rating labels located on the equipment in the rack to determine the total power requirement of the supply circuit.
- (For sliding drawers) Do not pull out or install any drawer or feature if the rack stabilizer brackets are not attached to the rack. Do not pull out more than one drawer at a time. The rack might become unstable if you pull out more than one drawer at a time.
- (For fixed drawers) This drawer is a fixed drawer and must not be moved for servicing unless specified by the manufacturer. Attempting to move the drawer partially or completely out of the rack might cause the rack to become unstable or cause the drawer to fall out of the rack.

(R001 part 2 of 2)

CAUTION:

Removing components from the upper positions in the rack cabinet improves rack stability during a relocation. Follow these general guidelines whenever you relocate a populated rack cabinet within a room or building.

- Reduce the weight of the rack cabinet by removing equipment starting at the top of the rack cabinet. When possible, restore the rack cabinet to the configuration of the rack cabinet as you received it. If this configuration is not known, you must observe the following precautions.
 - Remove all devices in the 32U position and above.
 - Ensure that the heaviest devices are installed in the bottom of the rack cabinet.
 - Ensure that there are no empty U-levels between devices installed in the rack cabinet below the 32U level.
- If the rack cabinet you are relocating is part of a suite of rack cabinets, detach the rack cabinet from the suite.
- If the rack cabinet you are relocating was supplied with removable outriggers they must be reinstalled before the cabinet is relocated.
- Inspect the route that you plan to take to eliminate potential hazards.
- Verify that the route that you choose can support the weight of the loaded rack cabinet. Refer to the documentation that comes with your rack cabinet for the weight of a loaded rack cabinet.
- Verify that all door openings are at least 760 x 230 mm (30 x 80 in.).
- Ensure that all devices, shelves, drawers, doors, and cables are secure.
- Ensure that the four leveling pads are raised to their highest position.
- Ensure that there is no stabilizer bracket installed on the rack cabinet during movement.
- Do not use a ramp inclined at more than 10 degrees.
- When the rack cabinet is in the new location, complete the following steps:
 - Lower the four leveling pads.
 - Install stabilizer brackets on the rack cabinet.
 - If you removed any devices from the rack cabinet, repopulate the rack cabinet from the lowest position to the highest position.
- If a long-distance relocation is required, restore the rack cabinet to the configuration of the rack cabinet as you received it. Pack the rack cabinet in the original packaging material, or equivalent. Also lower the leveling pads to raise the casters off the pallet and bolt the rack cabinet to the pallet.

(R002)

CAUTION:

- Rack is not intended to serve as an enclosure and does not provide any degrees of protection required of enclosures.
- It is intended that equipment installed within this rack will have its own enclosure. (R005).

CAUTION:

Tighten the stabilizer brackets until they are flush against the rack. (R006)

CAUTION:

Use safe practices when lifting. (R007)

CAUTION:

Do not place any object on top of a rack-mounted device unless that rack-mounted device is intended for use as a shelf. (R008)

CAUTION:

If the rack is designed to be coupled to another rack only the same model rack should be coupled together with another same model rack. (R009)

Danger notices for Storwize V7000 Unified

Ensure that you are familiar with the danger notices for Storwize V7000 Unified.

Use the reference numbers in parentheses at the end of each notice, such as (C003) for example, to find the matching translated notice in *IBM Systems Safety Notices*.

DANGER

When working on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard:

- If IBM supplied a power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.
- Do not open or service any power supply assembly.
- Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.
- Connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate.
- Connect any equipment that will be attached to this product to properly wired outlets.
- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.

To disconnect:

1. Turn off everything (unless instructed otherwise).
2. Remove the power cords from the outlets.
3. Remove the signal cables from the connectors.
4. Remove all cables from the devices.

To connect:

1. Turn off everything (unless instructed otherwise).
 2. Attach all cables to the devices.
 3. Attach the signal cables to the connectors.
 4. Attach the power cords to the outlets.
 5. Turn on the devices.
- Sharp edges, corners and joints might be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching. (D005)

DANGER

Heavy equipment—personal injury or equipment damage might result if mishandled. (D006)

DANGER

Observe the following precautions when working on or around your IT rack system:

- Heavy equipment—personal injury or equipment damage might result if mishandled.
- Always lower the leveling pads on the rack cabinet.
- Always install stabilizer brackets on the rack cabinet.
- To avoid hazardous conditions due to uneven mechanical loading, always install the heaviest devices in the bottom of the rack cabinet. Always install servers and optional devices starting from the bottom of the rack cabinet.
- Rack-mounted devices are not to be used as shelves or work spaces. Do not place objects on top of rack-mounted devices.



f2c00064

- Each rack cabinet might have more than one power cord. Be sure to disconnect all power cords in the rack cabinet when directed to disconnect power during servicing.
- Connect all devices installed in a rack cabinet to power devices installed in the same rack cabinet. Do not plug a power cord from a device installed in one rack cabinet into a power device installed in a different rack cabinet.
- An electrical outlet that is not correctly wired could place hazardous voltage on the metal parts of the system or the devices that attach to the system. It is the responsibility of the customer to ensure that the outlet is correctly wired and grounded to prevent an electrical shock.

(R001 part 1 of 2)

DANGER

Racks with a total weight of > 227 kg (500 lb.), Use Only Professional Movers! (R003)

DANGER


Do not transport the rack via fork truck unless it is properly packaged, secured on top of the supplied pallet. (R004)

DANGER



Main Protective Earth (Ground):

This symbol is marked on the frame of the rack.

The PROTECTIVE EARTHING CONDUCTORS should be terminated at that point. A recognized or certified closed loop connector (ring terminal) should be used and secured to the frame with a lock washer using a bolt or stud. The connector should be properly sized to be suitable for the bolt or stud, the locking washer, the rating for the conducting wire used, and the considered rating of the breaker. The intent is to ensure the frame is electrically bonded to the PROTECTIVE EARTHING CONDUCTORS. The hole that the bolt or stud goes into where the terminal conductor and the lock washer contact should be free of any non-conductive material to allow for metal to metal contact. All PROTECTIVE EARTHING CONDUCTORS should terminate at this main protective earthing terminal or at points marked with  . (R010)

Special caution and safety notices

This information describes special safety notices that apply to the Storwize V7000 Unified. These notices are in addition to the standard safety notices supplied and address specific issues relevant to the equipment provided.

General safety

When you service the Storwize V7000 Unified, follow general safety guidelines.

Use the following general rules to ensure safety to yourself and others.

- Observe good housekeeping in the area where the devices are kept during and after maintenance.
- Follow the guidelines when lifting any heavy object:
 1. Ensure that you can stand safely without slipping.
 2. Distribute the weight of the object equally between your feet.
 3. Use a slow lifting force. Never move suddenly or twist when you attempt to lift.
 4. Lift by standing or by pushing up with your leg muscles; this action removes the strain from the muscles in your back. *Do not attempt to lift any objects that weigh more than 18 kg (40 lb) or objects that you think are too heavy for you.*
- Do not perform any action that causes a hazard or makes the equipment unsafe.
- Before you start the device, ensure that other personnel are not in a hazardous position.
- Place removed covers and other parts in a safe place, away from all personnel, while you are servicing the unit.
- Keep your tool case away from walk areas so that other people cannot trip over it.
- Do not wear loose clothing that can be trapped in the moving parts of a device. Ensure that your sleeves are fastened or rolled up above your elbows. If your hair is long, fasten it.

- Insert the ends of your necktie or scarf inside clothing or fasten it with a nonconducting clip, approximately 8 cm (3 in.) from the end.
- Do not wear jewelry, chains, metal-frame eyeglasses, or metal fasteners for your clothing.

Remember: Metal objects are good electrical conductors.

- Wear safety glasses when you are hammering, drilling, soldering, cutting wire, attaching springs, using solvents, or working in any other conditions that might be hazardous to your eyes.
- After service, reinstall all safety shields, guards, labels, and ground wires. Replace any safety device that is worn or defective.
- Reinstall all covers correctly after you have finished servicing the unit.

Handling static-sensitive devices

Ensure that you understand how to handle devices that are sensitive to static electricity.

Attention: Static electricity can damage electronic devices and your system. To avoid damage, keep static-sensitive devices in their static-protective bags until you are ready to install them.

To reduce the possibility of electrostatic discharge, observe the following precautions:

- Limit your movement. Movement can cause static electricity to build up around you.
- Handle the device carefully, holding it by its edges or frame.
- Do not touch solder joints, pins, or exposed printed circuitry.
- Do not leave the device where others can handle and possibly damage the device.
- While the device is still in its antistatic bag, touch it to an unpainted metal part of the system unit for at least two seconds. (This action removes static electricity from the package and from your body.)
- Remove the device from its package and install it directly into your Storwize V7000 Unified, without putting it down. If it is necessary to put the device down, place it onto its static-protective bag. (If your device is an adapter, place it component-side up.) Do not place the device onto the cover of the Storwize V7000 Unified or onto a metal table.
- Take additional care when you handle devices during cold weather. Indoor humidity tends to decrease in cold weather, causing an increase in static electricity.

Sound pressure

Attention: Depending on local conditions, the sound pressure can exceed 85 dB(A) during service operations. In such cases, wear appropriate hearing protection.

Environmental notices

This information contains all of the required environmental notices for IBM Systems products in English and other languages.

The *IBM Systems Environmental Notices* (<http://ibm.co/1fBgwFI>) information includes statements on limitations, product information, product recycling and disposal, battery information, flat panel display, refrigeration and water-cooling systems, external power supplies, and safety data sheets.

About this guide

This publication provides information that helps you install and initialize IBM Storwize V7000 Unified systems.

Who should use this guide

This guide is intended for installers of Storwize V7000 Unified systems.

Before configuring your system, ensure that you follow the procedures as listed. Be sure to gather IP addresses that you will need before you begin the installation.

Storwize V7000 Unified library and related publications

Product manuals, other publications, and websites contain information that relates to Storwize V7000 Unified.

IBM Knowledge Center for Storwize V7000 Unified

4 The information collection in the IBM Knowledge Center contains all of the
4 information that is required to install, configure, and manage the system. The
4 information collection in the IBM Knowledge Center is updated between product
4 releases to provide the most current documentation. The information collection is
4 available at the following website:

publib.boulder.ibm.com/infocenter/storwize/unified_ic/index.jsp

Storwize V7000 Unified library

Unless otherwise noted, the publications in the library are available in Adobe portable document format (PDF) from a website.

www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss

Click **Search for publications** to find the online publications you are interested in, and then view or download the publication by clicking the appropriate item.

Table 1 lists websites where you can find help, services, and more information.

Table 1. IBM websites for help, services, and information

Website	Address
IBM home page	http://www.ibm.com
Directory of worldwide contacts	http://www.ibm.com/planetwide
Support for Storwize V7000 (2076)	www.ibm.com/storage/support/storwize/v7000
Support for Storwize V7000 Unified (2073)	www.ibm.com/storage/support/storwize/v7000/unified
Support for IBM System Storage® and IBM TotalStorage products	www.ibm.com/storage/support/

Each of the PDF publications in the Table 2 library is also available in the IBM Knowledge Center by clicking the number in the “Order number” column:

Table 2. Storwize V7000 Unified library

Title	Description	Order number
IBM Storwize V7000 Model 2073-720 Quick Start Guide	The guide provides general instructions for installing your system, and is intended for experienced developers.	
<i>Storwize V7000 Unified Quick Installation Guide</i>	The guide provides instructions for unpacking your order and installing your system. The first chapter describes verifying your order, becoming familiar with the hardware components, and meeting environmental requirements. The second chapter describes installing the hardware and attaching data cables and power cords. The last chapter describes accessing the management GUI to initially configure your system.	
<i>IBM Storwize V7000 Expansion Enclosure Installation Guide, Machine type 2076</i>	The guide provides instructions for unpacking your order and installing the 2076 expansion enclosure for the Storwize V7000 Unified system.	
<i>Adding Storwize V7000 Unified File modules to an Existing Storwize V7000 System</i>	The guide is for adding Storwize V7000 file modules to an existing Storwize V7000 system to create a Storwize V7000 Unified system.	
<i>Storwize V7000 Unified Problem Determination Guide</i>	The guide describes how to service, maintain, and troubleshoot the Storwize V7000 Unified system.	
<i>IBM Storwize V7000 Unified Safety Notices</i>	The guide contains translated caution and danger statements for the node canister documentation. Each caution and danger statement in the Storwize V7000 Unified documentation has a number that you can use to locate the corresponding statement in your language in the <i>IBM Storwize V7000 Unified Safety Notices</i> document.	

Table 2. Storwize V7000 Unified library (continued)

Title	Description	Order number
<i>Safety Information</i>	The guide contains translated caution and danger statements for the file module documentation. Each caution and danger statement in the Storwize V7000 Unified documentation has a number. Use the number to locate the corresponding statement in your language in the <i>Safety Information</i> document.	
<i>Storwize V7000 Unified Read First Flyer</i>	This document introduces the major components of the Storwize V7000 Unified system and describes how to get started with the <i>Storwize V7000 Unified Quick Installation Guide</i> .	
<i>Read First before adding file modules to an existing Storwize V7000 Unified</i>	This document introduces the major components of the Storwize V7000 Unified system and describes how to get started with <i>Adding Storwize V7000 Unified File modules to an Existing Storwize V7000 System</i> .	
<i>IBM Statement of Limited Warranty (2145 and 2076)</i>	This multilingual document provides information about the IBM warranty for machine types 2145 and 2076.	
<i>IBM Statement of Limited Warranty (2073)</i>	This multilingual document provides information about the IBM warranty for machine type 2073.	
<i>IBM License Agreement for Machine Code</i>	This multilingual guide contains the License Agreement for Machine Code for the Storwize V7000 Unified product.	
<i>Getting Started with Real-time Compression™ on IBM Storwize(r) V7000 Unified 1.4.0.1</i>	This document provides technical information and guidelines on what to consider to deploy compression in the Storwize V7000 Unified storage environment.	
<i>IBM Storwize(r) V7000 Unified Data Migration Guide: NetApp to IBM Storwize(r) V7000 Unified</i>	This document is a guide for migrating data from a NetApp system to an IBM Storwize(r) V7000 Unified NAS platform.	

IBM documentation and related websites

Table 3 lists websites that provide publications and other information about the Storwize V7000 Unified or related products or technologies. The IBM Redbooks® publications provide positioning and value guidance, installation and implementation experiences, solution scenarios, and step-by-step procedures for various products.

Table 3. IBM documentation and related websites

Website	Address
IBM Publications Center	www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss
IBM Redbooks publications	www.redbooks.ibm.com/

Related accessibility information

To view a PDF file, you need Adobe Reader, which can be downloaded from the Adobe website:

www.adobe.com/support/downloads/main.html

How to order IBM publications

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

The IBM Publications Center offers customized search functions to help you find the publications that you need. Some publications are available for you to view or download at no 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 website:

www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss

Related websites

The following websites provide information about Storwize V7000 Unified or related products or technologies.

Type of information	Website
Storwize V7000 Unified support	www.ibm.com/storage/support/storwize/v7000/unified
Technical support for IBM storage products	www.ibm.com/storage/support/
IBM Electronic Support registration	www.ibm.com/support/electronicssupport

Sending your comments

Your feedback is important in helping to provide the most accurate and highest quality information.

To submit any comments about this book or any other Storwize V7000 Unified documentation, send your comments by email to starpubs@us.ibm.com. Include the following information in your email:

- Publication title
- Publication form number
- Page, table, or illustration numbers that you are commenting on
- A detailed description of any information that should be changed

How to get information, help, and technical assistance

If you need help, service, technical assistance, or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you.

Information

IBM maintains pages on the web where you can get information about IBM products and fee services, product implementation and usage assistance, break and fix service support, and the latest technical information. For more information, refer to Table 4.

Table 4. IBM websites for help, services, and information

Website	Address
IBM home page	http://www.ibm.com
Directory of worldwide contacts	http://www.ibm.com/planetwide
Support for Storwize V7000 (2076)	www.ibm.com/storage/support/storwize/v7000
Support for Storwize V7000 Unified (2073)	www.ibm.com/storage/support/storwize/v7000/unified
Support for IBM System Storage and IBM TotalStorage products	www.ibm.com/storage/support/

Note: Available services, telephone numbers, and web links are subject to change without notice.

Help and service

Before calling for support, be sure to have your IBM Customer Number available. If you are in the US or Canada, you can call 1 (800) IBM SERV for help and service. From other parts of the world, see <http://www.ibm.com/planetwide> for the number that you can call.

When calling from the US or Canada, choose the **storage** option. The agent decides where to route your call, to either storage software or storage hardware, depending on the nature of your problem.

If you call from somewhere other than the US or Canada, you must choose the **software** or **hardware** option when calling for assistance. Choose the **software** option if you are uncertain if the problem involves the Storwize V7000 Unified software or hardware. Choose the **hardware** option only if you are certain the problem solely involves the Storwize V7000 Unified hardware. When calling IBM for service regarding the product, follow these guidelines for the **software** and **hardware** options:

Software option

Identify the Storwize V7000 Unified product as your product and supply your customer number as proof of purchase. The customer number is a 7-digit number (0000000 to 9999999) assigned by IBM when the product is purchased. Your customer number should be located on the customer information worksheet or on the invoice from your storage purchase. If asked for an operating system, use **Storage**.

Hardware option

Provide the serial number and appropriate 4-digit machine type. For Storwize V7000 Unified, the machine type is 2073.

In the US and Canada, hardware service and support can be extended to 24x7 on the same day. The base warranty is 9x5 on the next business day.

Getting help online

You can find information about products, solutions, partners, and support on the IBM website.

To find up-to-date information about products, services, and partners, visit the IBM website at www.ibm.com/storage/support/storwize/v7000/unified.

Before you call

Make sure that you have taken steps to try to solve the problem yourself before you call.

Some suggestions for resolving the problem before calling IBM Support include:

- Check all cables to make sure that they are connected.
- Check all power switches to make sure that the system and optional devices are turned on.
- Use the troubleshooting information in your system documentation. The troubleshooting section of the information center contains procedures to help you diagnose problems.
- Go to the IBM Support website at www.ibm.com/storage/support/storwize/v7000/unified to check for technical information, hints, tips, and new device drivers or to submit a request for information.

Using the documentation

Information about your IBM storage system is available in the documentation that comes with the product.

That documentation includes printed documents, online documents, readme files, and help files in addition to the information center. See the troubleshooting information for diagnostic instructions. The troubleshooting procedure might require you to download updated device drivers or software. IBM maintains pages on the web where you can get the latest technical information and download device drivers and updates. To access these pages, go to www.ibm.com/storage/support/storwize/v7000/unified and follow the instructions. Also, some documents are available through the IBM Publications Center.

Sign up for the Support Line Offering

If you have questions about how to use and configure the machine, sign up for the IBM Support Line offering to get a professional answer.

The maintenance supplied with the system provides support when there is a problem with a hardware component or a fault in the system machine code. At times, you might need expert advice about using a function provided by the system or about how to configure the system. Purchasing the IBM Support Line offering gives you access to this professional advice while deploying your system, and in the future.

Contact your local IBM sales representative or your support group for availability and purchase information.

What's new

This book describes troubleshooting a Storwize V7000 Unified system. For information about the new features and updates added in this release, see the *What's New* topic in the Storwize V7000 Unified Information Center.

Chapter 1. Storwize V7000 Unified hardware components

A Storwize V7000 Unified system consists of 1 or more machine type 2076 rack-mounted enclosures and 2 machine type 2073 rack-mounted file modules. Control enclosures contain the node canisters that manage the system operation and provide the host interfaces. Expansion enclosures provide more extra drives that can be managed by the system. Enclosures can support 2.5 inch (6.35 cm) small form factor drives or 3.5 inch (8.89 cm) large form factor drives.

There are several model types for the 2076 machine type. The main differences among the model types are the following items:

- The number of drives that an enclosure can hold. Drives are on the front of the enclosure. An enclosure can hold up to 12 3.5-inch drives or up to 24 2.5-inch drives.

- Whether the model is a control enclosure or an expansion enclosure

Control enclosures contain the main processing units that control the whole system. They are where external systems such as host application servers, other storage systems, and management workstations are connected through the Ethernet ports or Fibre Channel ports. Control enclosures can also be connected to expansion enclosures through the serial-attached SCSI (SAS) ports.

Expansion enclosures contain more storage capacity. Expansion enclosures connect either to control enclosures or to other expansion enclosures through the SAS ports.

- If the control enclosure has either 1 Gbps Ethernet capability or 10 Gbps Ethernet capability

The machine type and model for the file module is 2073-720.

Components in the front of the 2073-700 file module

This topic describes the components in the front of the file module.

Figure 1 shows the various front ports and hardware for the 2073-700 file module:

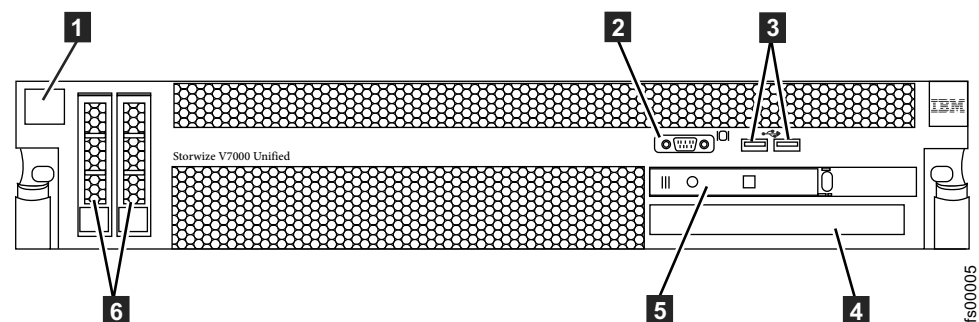


Figure 1. Front view of 2073-700 file module

- **1** File module label with MTM (machine type model) and S/N (serial number)
- **2** VGA port
- **3** USB ports
- **4** DVD drive

- **5** Control panel
- **6** 2 drives

Components in the front of the 2073-720 file module

This topic describes the components in the front of the 2073-720 file module and the functions of the advanced operator panel.

Figure 2 shows the various front ports and hardware for the 2073-720 file module.

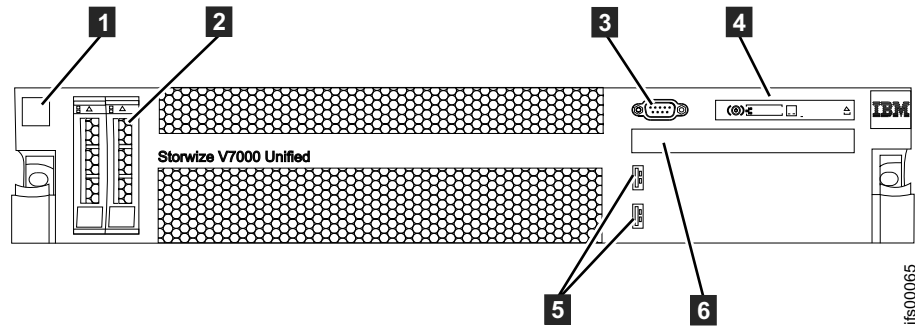


Figure 2. Front view of 2073-720 file module

- **1** File module label with MTM (machine type model) and S/N (serial number)
- **2** Boot drives
- **3** Video port
- **4** Advanced operator panel.
- **5** USB ports
- **6** DVD drive

Figure 3 displays the functions of the 2073-720 advanced operator panel.

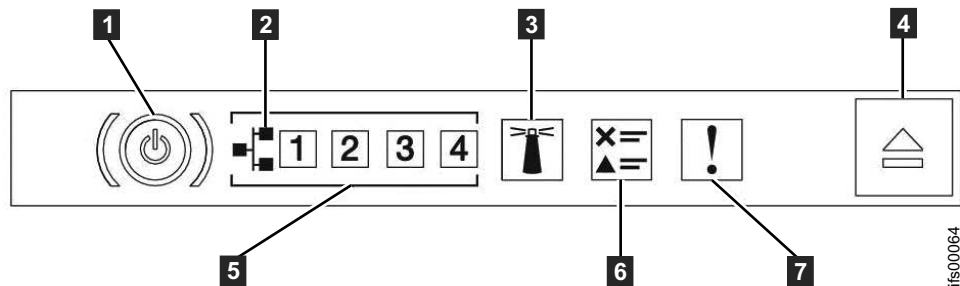


Figure 3. 2073-720 file module advanced operator panel

- **1** Power-control button and power-on LED (green)
- **2** Ethernet icon
- **3** System-locator button and LED (blue)
- **4** Release latch for the light path diagnostics panel
- **5** Ethernet activity LEDs
- **6** Check log LED
- **7** System-error LED: (yellow)

Components in the rear of the 2073-700 file module

This topic identifies the components in the rear of the 2073-700 file module.

Figure 4 identifies the various ports and hardware in the rear of the 2073-700 file module:

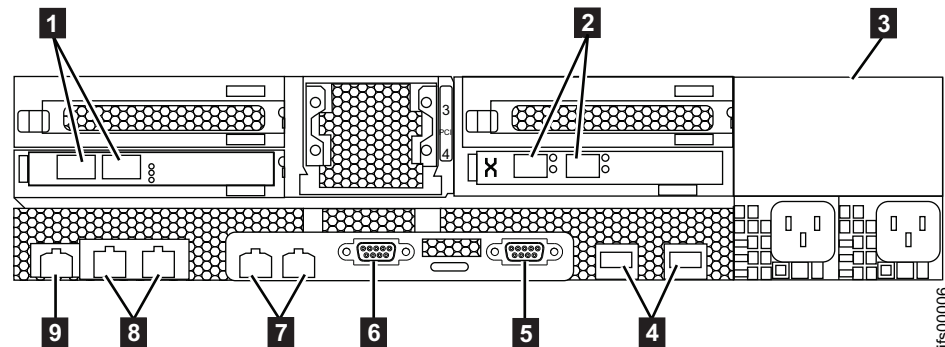


Figure 4. Rear view of 2073-700 file module

- **1** 10 Gbps Ethernet ports. Port 2 is on left, and port 1 is on right.
- **2** Fibre Channel ports. Port 2 is on left, and port 1 is on right.
- **3** Power supply unit
- **4** USB ports
- **5** Serial port
- **6** Video port
- **7** 1 Gbps Ethernet ports. Port 1 is on left, and port 2 is on right.
- **8** GbE ports. Port 3 is on left, and port 4 is on right.
- **9** GbE management port

Components in the rear of the 2073-720 file module

This topic identifies the components in the rear of the 2073-720 file module.

Figure 5 identifies the various ports and hardware in the rear of the 2073-720 file module:

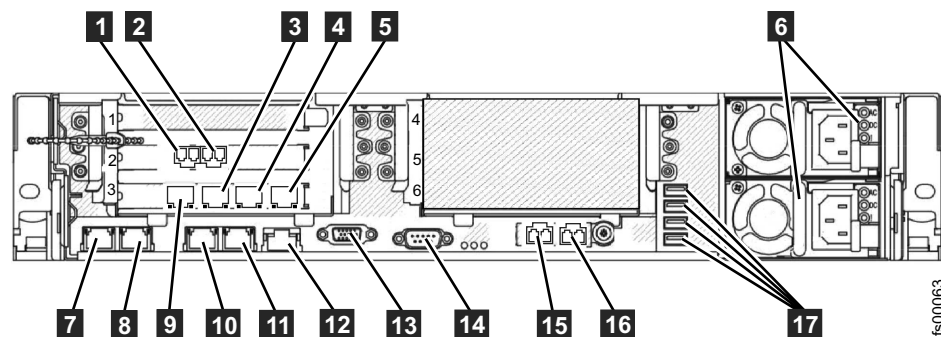


Figure 5. Rear view of 2073-720 file module

- **1** 8 Gbps Fibre Channel port 1 (connected to the control enclosure)
- **2** 8 Gbps Fibre Channel port 2 (connected to the control enclosure)

- **3** Ethernet port 8
- **4** Ethernet port 9
- **5** Ethernet port 10
- **6** Power supplies (1 is lower 2 is upper)
- **7** Ethernet port 1 (connected to the other file module)
- **8** Ethernet port 2 (connected to the other file module)
- **9** Ethernet port 7
- **10** Ethernet port 3
- **11** Ethernet port 4
- **12** Systems-management Ethernet port (NOT USED)
- **13** Video port
- **14** Serial port
- **15** Ethernet port 5 (10 Gbps)
- **16** Ethernet port 6 (10 Gbps)
- **17** USB ports

Components in the front of the enclosure

The front of each control enclosure features several different components.

Drives for control enclosures

Storwize V7000 Unified enclosures use different drives, depending on the generation of your control enclosure model. An enclosure can hold up to 12 3.5 in. (8.89 cm) drives or up to 24 2.5 in. (6.35 cm) drives.

Storwize V7000 Gen1

The drives are in the front of the enclosure. The 12 large form factor (LFF) 3.5 in. (8.89 cm) drives are mounted horizontally in four columns with three rows. The 24 small form factor (SFF) 2.5 in. (6.35 cm) drives are mounted vertically in one row.

Important: Drive slots cannot be empty. Install a drive assembly or blank carrier in each slot.

Note: Drives that are sold as Storwize V7000 Unified options are the only drives that are supported. For more information, see the Support website.

Figure 6 shows 12 LFF drives, as viewed from the front of the enclosure.



Figure 6. Storwize V7000 Gen1 Drives on a 12-drive enclosure

Figure 7 on page 5 shows a large form factor 3.5 in. (8.89 cm) drive with the latching mechanism open.

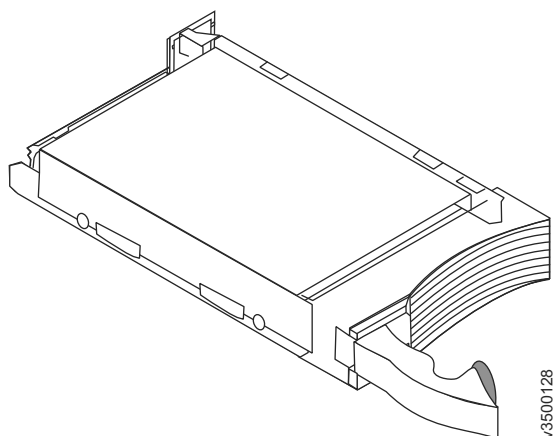


Figure 7. Storwize V7000 Gen2 Large form factor horizontal drive

Figure 8 shows 24 SFF drives, as viewed from the front of the enclosure.



Figure 8. Storwize V7000 Gen1 Drives on a 24-drive enclosure

Storwize V7000 Gen2

Figure 9 shows a small form factor 2.5 in. (6.35 cm) drive with the latching mechanism open.

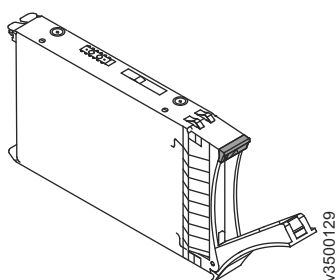


Figure 9. Storwize V7000 Gen2 Small form factor vertical drive

Drive indicators for control enclosures

Storwize V7000 Unified enclosures use different drive indicators, depending on the generation of your control enclosure model. Drives have two light-emitting diode (LED) indicators each; they have no controls or connectors.

Storwize V7000 Gen1

The LED color is the same for both drives. The LEDs for the 3.5-inch drives are placed vertically above and below each other. Figure 10 on page 6 displays the LED indicators on a single 3.5-inch drive. The LEDs for the 2.5-inch drives are placed next to each other at the bottom. Figure 11 on page 6 displays the LED

indicators on a single 2.5-inch drive.

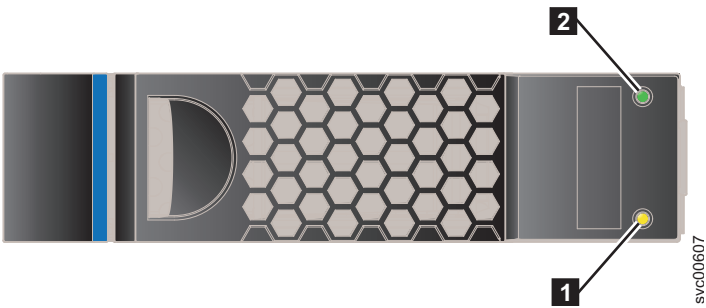


Figure 10. LED indicators on a single 3.5 inch drive



Figure 11. LED indicators on a single 2.5 inch drive

- 1 Fault LED
- 2 Activity LED

Table 5 shows the status descriptions for the two LEDs.

Table 5. Drive LEDs

Name	Description	Color
Activity	Indicates if the drive is ready or active. <ul style="list-style-type: none">• If the LED is on, the drive is ready to be used.• If the LED is off, the drive is not ready.• If the LED is flashing, the drive is ready, and there is activity.	Green
Fault	Indicates a fault or identifies a drive. <ul style="list-style-type: none">• If the LED is on, a fault exists on the drive.• If the LED is off, no known fault exists on the drive.• If the LED is flashing, the drive is being identified. A fault might or might not exist.	Amber

Enclosure end cap indicators

Storwize V7000 Unified enclosure end cap indicators vary, depending on the generation of your control enclosure model.

Storwize V7000 Gen1

Figure 12 shows where the end caps are on the front of an enclosure with 12 drives. The end caps are in the same position for an enclosure with 24 drives.

- 1** Left end cap
- 2** Drives
- 3** Right end cap

The left enclosure end caps for both enclosures are identical and contain only indicators. The left enclosure end cap contains no controls or connectors. The right enclosure end cap for both enclosures has no controls, indicators, or connectors.

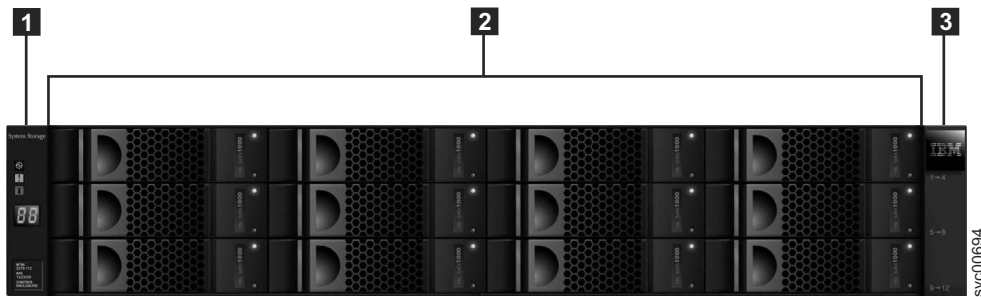


Figure 12. Storwize V7000 Gen1 12 drives and two end caps

Figure 13 on page 8 shows the indicators on the front of the enclosure end cap.

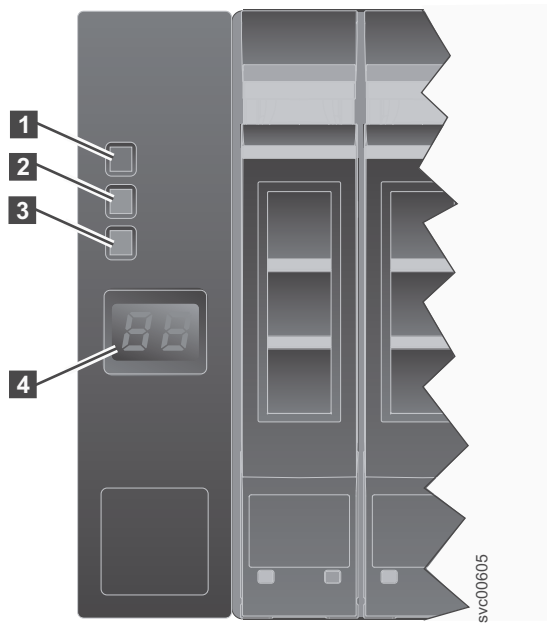


Figure 13. Storwize V7000 Gen1 Left enclosure end cap

Table 6. Storwize V7000 Gen1 LED descriptions for left enclosure end cap

Name	Description	Color	Symbol
Power	1 The power LED is the upper LED. When the green LED is lit, it indicates that the main power is available to the enclosure	Green	Ⓢ
Fault	2 The fault LED is the middle LED. When the amber LED is lit, it indicates one of the enclosure components has a hardware fault.	Amber	!
Identify	3 The identify LED is the lower LED. When the blue LED is lit, it identifies the enclosure.	Blue	Ⓜ
N/A	4 The two-character LCD display shows the enclosure ID.	N/A	N/A

Components in the rear of the enclosure

This topic describes the hardware components in the rear of the enclosure.

Two canisters are located in the middle of each enclosure. The power supply units are located on the left and right of the canisters. The left slot is power supply 1 (**1**), and the right slot is power supply 2 (**2**). Power supply 1 is top side up, and power supply 2 is inverted. The upper slot is canister 1 (**3**), and the lower slot is canister 2 (**4**). Canister 1 is top side up, and canister 2 is inverted.

Figure 14 on page 9 shows the rear view of a model 2076-112 or a model 2076-124 control enclosure. Figure 15 on page 9 shows the rear view of a model 2076-312 or a model 2076-324 control enclosure with the 10 Gbps Ethernet port (**5**). Figure 16 on page 9 shows the rear of an expansion enclosure.

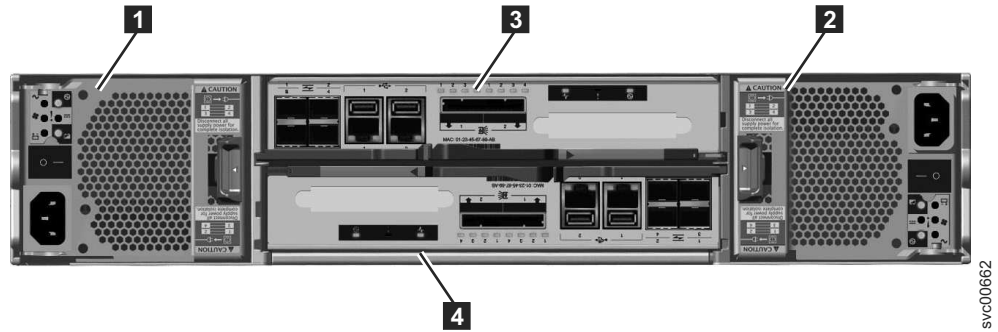


Figure 14. Rear view of a model 2076-112 or a model 2076-124 control enclosure

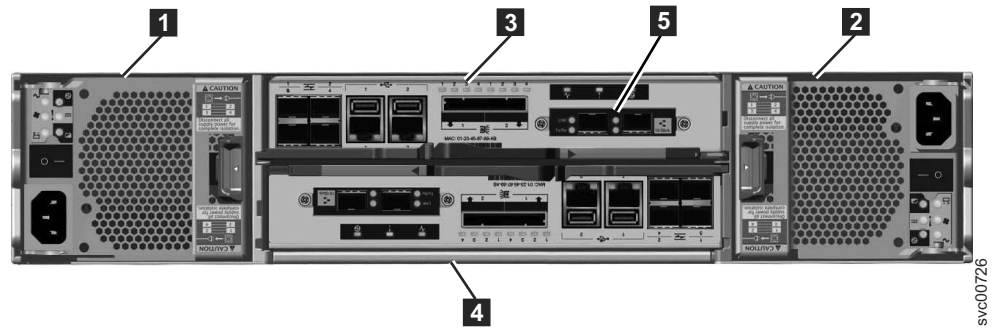


Figure 15. Rear view of a model 2076-312 or a model 2076-324 control enclosure

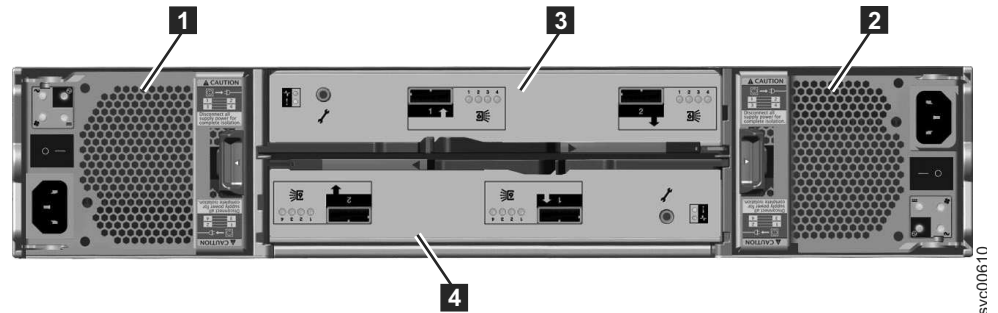


Figure 16. Rear view of a model 2076-212 or a model 2076-224 expansion enclosure

- 1** Power supply unit 1
- 2** Power supply unit 2
- 3** Canister 1
- 4** Canister 2
- 1** Node canisters
- 2** Power supply units

Power supply units for control enclosures

Storwize V7000 Unified enclosures use different power supply units, depending on the generation of your control enclosure model.

Storwize V7000 Unified *Gen1* refers to the enclosure models in the following table:

Table 7. Storwize V7000 Unified Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 Unified control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 Unified control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 Unified control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 Unified control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 Unified expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Storwize V7000 Unified Gen2 refers to the newer generation of enclosures in the following table:

Table 8. Storwize V7000 Unified Gen2 model numbers

Machine type/model	Description
2076-524	Storwize V7000 Unified control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-12F	Storwize V7000 Unified expansion enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Storwize V7000 Gen2 power supply units

Each Storwize V7000 Gen2 enclosure contains two power supply units. Each power supply unit can provide power to the whole enclosure.

Note: The power supply has no power switch. A power supply is active when a power cord is connected to the power connector and to a power source.

Figure 17 shows the rear view of a control enclosure and identifies the location of the power supply units and node canisters.

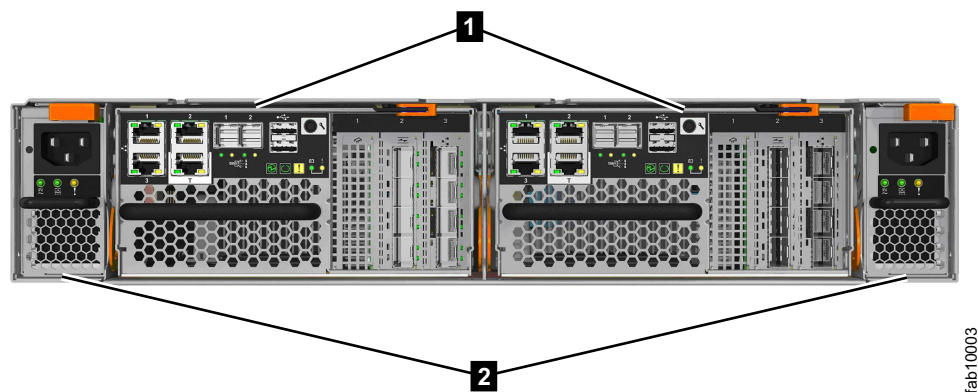


Figure 17. Rear view of a Storwize V7000 Gen2 control enclosure

- 1** Node canisters
- 2** Power supply units

Each power supply also contains a fan that cools the power supply unit itself. Cool air is drawn in and passes over each power supply. The warmed air is ejected through the rear of each power supply. For optimal cooling, do not obstruct this airflow. Also, ensure that all enclosure components or fillers are installed while the system is operational.

Storwize V7000 Gen1 Power supply unit and battery for the control enclosure

The Storwize V7000 Gen1 control enclosure contains two power supply units, each with an integrated battery. Each power supply unit can provide power to the whole enclosure.

The two power supply units in the enclosure are installed with one unit top side up and the other inverted. The power supply unit for the control enclosure has six LEDs.

Each power supply unit has a power switch. The switch must be on for the power supply unit to be operational. If the power switches are turned off, or the main power is removed, the integrated batteries temporarily continue to supply power to the node canisters. As a result, the canisters can store configuration data and cached data to their internal drives. Battery power is required only if both power supply units stop operating.

Figure 18 on page 12 shows the location of the LEDs **1** in the rear of the power supply unit.

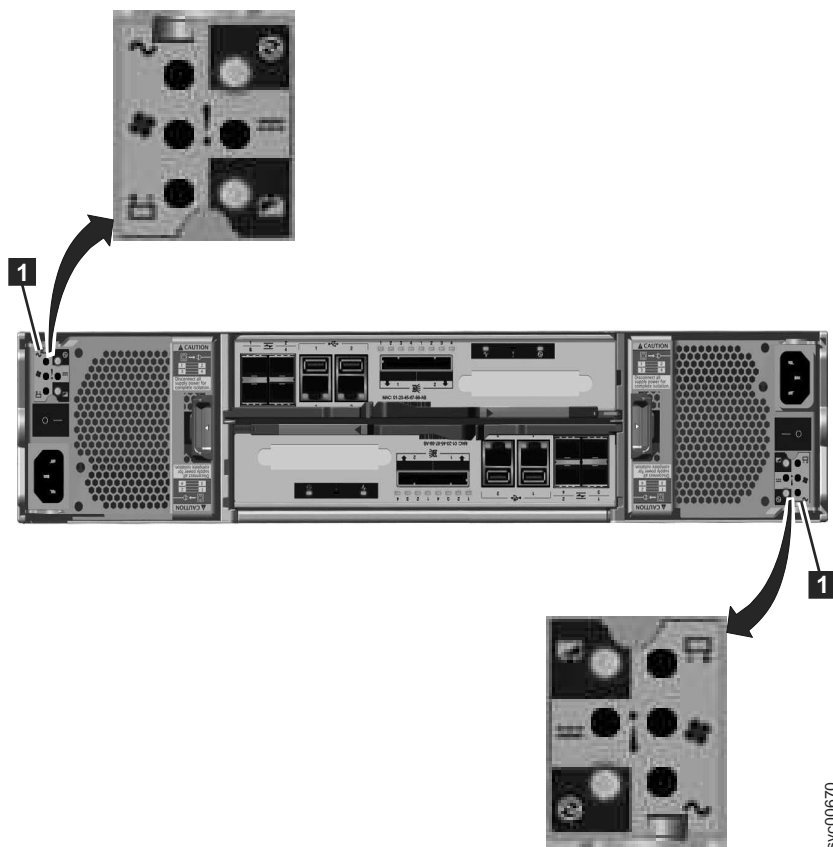


Figure 18. LEDs on the power supply units of the control enclosure

Table 9 identifies the LEDs in the rear of the control enclosure. Each power supply also contains fans that cool the enclosure. Cool air is drawn in through the front of the enclosure. The air passes over the drives, node canisters, and power supplies. The warmed air is ejected through the rear of each power supply. For optimal cooling, do not obstruct the airflow and ensure that all enclosure components or fillers are installed while the system is operational.

Table 9. Power supply unit LEDs in the rear of the control enclosure

Name	Color	Symbol
AC power failure	Amber	~
Power supply OK	Green	⏻
Fan failure	Amber	☼
DC power failure	Amber	≡
Battery failure	Amber	🔋
Battery state	Green	🔋

See “Procedure: Understanding the Storwize V7000 Gen1 system status using the LEDs” on page 262 for help in diagnosing a particular failure.

Power supply units for expansion enclosures

Storwize V7000 Unified enclosures use different power supply units, depending on the generation of your expansion enclosure model.

Storwize V7000 Unified *Gen1* refers to the enclosure models in the following table:

Table 10. Storwize V7000 Unified Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 Unified control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 Unified control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 Unified control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 Unified control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 Unified expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Storwize V7000 Unified *Gen2* refers to the newer generation of enclosures in the following table:

Table 11. Storwize V7000 Unified Gen2 model numbers

Machine type/model	Description
2076-524	Storwize V7000 Unified control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-12F	Storwize V7000 Unified expansion enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Storwize V7000 Gen2 power supply units for expansion enclosures

The Storwize V7000 Gen2 expansion enclosure contains two power supply units (PSU).

Figure 19 on page 14 shows the locations of the expansion canisters and the two power supply units in the rear of the expansion enclosure.

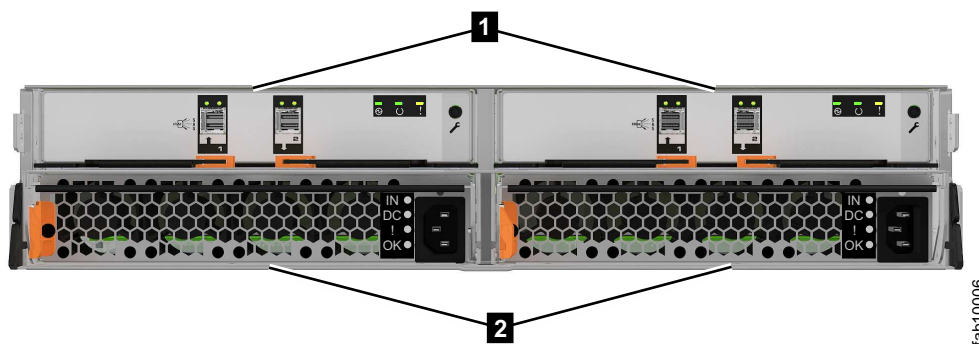


Figure 19. Rear view of a Storwize V7000 Gen2 expansion enclosure

- 1** Expansion canisters
- 2** Power supply units

Each power supply unit has four LED indicators (Table 12).

Table 12. Power supply LEDs

Name	Label	Color	Description
Input status		Green	Off No input power detected
			On Direct current input power detected
Output status		Green	Off PSU is not providing dc output power
			On PSU is providing dc output power
Fault		Amber	Off No fault detected
			On PSU fault has been detected
			BLINK PSU is being identified. A fault may have been detected.
(None)		Blue	Not used

See for help in diagnosing a particular failure.

Storwize V7000 Gen1 power supply units for expansion enclosures

The Storwize V7000 Gen1 expansion enclosure contains two power supply units (PSU).

The two power supply units in the enclosure are installed with one unit top side up and the other inverted. The power supply unit for the expansion enclosure has four LEDs, two less than the power supply for the control enclosure.

There is a power switch on each of the power supply units. The switch must be on for the power supply unit to be operational. If the power switches are turned off, the power supply units stop providing power to the system.

Table 13 on page 15 shows the locations of the LEDs **1** in the rear of the power supply unit.

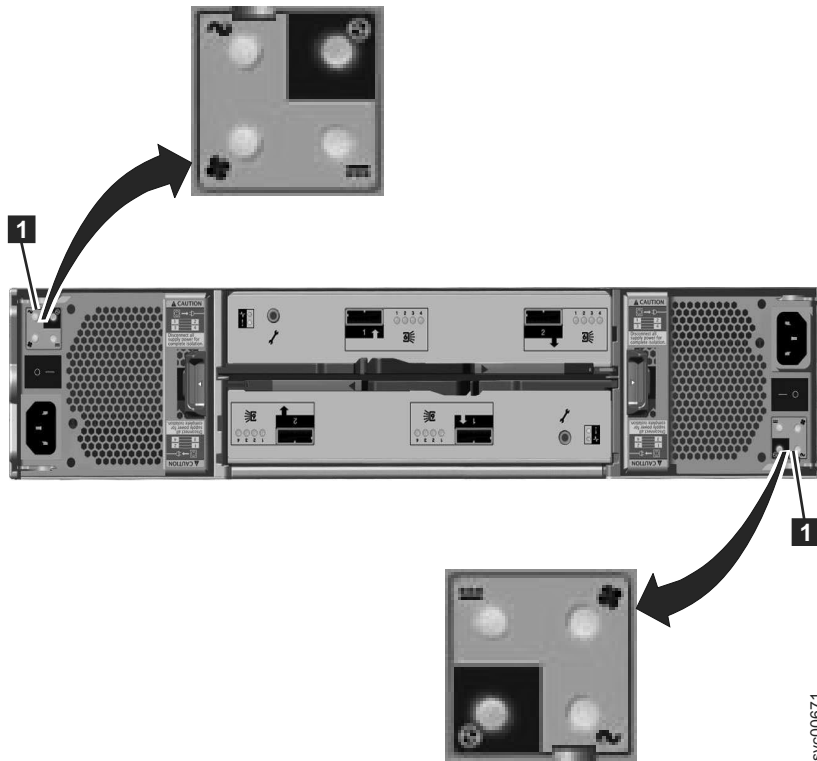


Figure 20. LEDs on the power supply units of the expansion enclosure

Table 13 identifies the LEDs in the rear of the expansion enclosure.

Table 13. Power supply unit LEDs in the rear of the expansion enclosure

Name	Color	Symbol
ac power failure	Amber	~
Power supply OK	Green	Ⓢ
Fan failure	Amber	✿
dc power failure	Amber	≡

See for help in diagnosing a particular failure.

Storwize V7000 2076-524 node canister ports and indicators

The node canister has indicators and ports but no controls.

A Fibre Channel over Ethernet (FCoE)/Internet Small Computer System Interface (iSCSI) host interface adapter may be installed in a node canister.

A node canister contains a battery that provides power to the canister as it stores cache and system data to an internal drive in the event of a power failure. This process is known as a *fire hose dump*.

Storwize V7000 Gen2 node canister ports

Each Storwize V7000 2076-524 node canister has ports for connecting Ethernet, iSCSI, and USB peripherals, and optional expansion enclosures.

Figure 21 illustrates the location of the ports.



Figure 21. Storwize V7000 2076-524 node canister ports

Storwize V7000 2076-524 node canister indicators

Each Storwize V7000 2076-524 node canister has indicator LEDs that provide status information about the canister.

Using the callout numbers in Figure 22, refer to the tables for a listing of the Storwize V7000 2076-524 node canister LEDs and a description of the meaning of the LED activity.

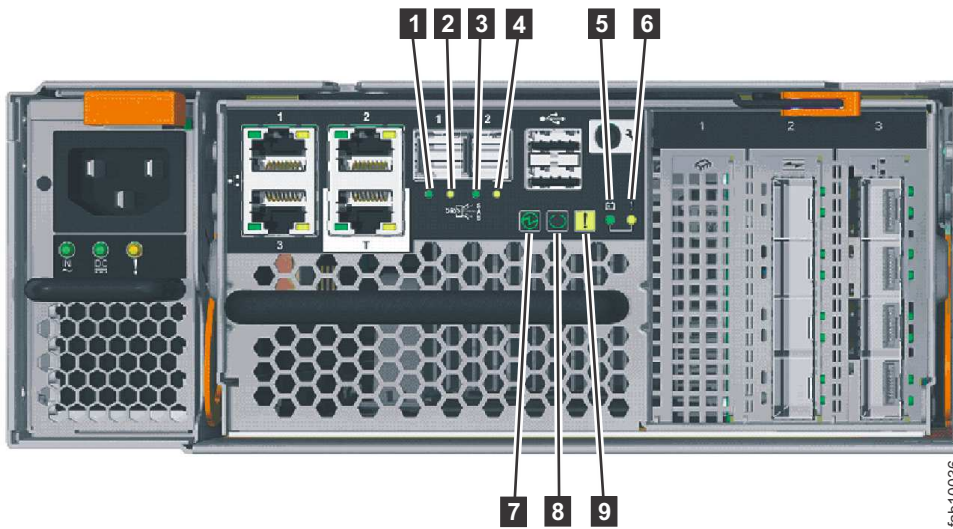


Figure 22. Storwize V7000 2076-524 node canister indicators

Storwize V7000 2076-524 node canister SAS port LEDs

Table 14 on page 17 depict the status LEDs for SAS ports 1 and 2, and their location in Figure 22.

Table 14. Storwize V7000 2076-524 SAS ports 1 and 2 LEDs. Storwize V7000 2076-524 SAS ports 1 and 2 LED legend.

Name	Call out	Symbol	Color	State	Meaning
SAS Port 1 Link	1	None	Green	OFF	No link connection on any phys (lanes). The connection is down.
				ON	The port is connected to at least one phy. At least one of the phys to that connector is up.
SAS Port 1 Fault	2	None	Amber	OFF	No fault. All four phys have a link connection.
				ON	<p>This status can indicate several different error conditions:</p> <ul style="list-style-type: none"> • One or more, but not all, of the 4 phys are connected. • Not all 4 phys are at the same speed. • One or more of the connected phys are attached to an address different from the others. • An unsupported device is plugged in to this SAS port.

Table 14. Storwize V7000 2076-524 SAS ports 1 and 2 LEDs (continued). Storwize V7000 2076-524 SAS ports 1 and 2 LED legend.

Name	Call out	Symbol	Color	State	Meaning
SAS Port 2 Link	3	None	Green	OFF	No link connection on any phys (lanes). The connection is down.
				ON	The port is connected to at least one phy. At least one of the lanes to that connector is up.
SAS Port 2 Fault	4	None	Amber	OFF	No fault. All four phys have a link connection.
				ON	<p>This status can indicate several different error conditions:</p> <ul style="list-style-type: none"> • One or more, but not all, of the 4 phys are connected. • Not all 4 phys are at the same speed. • One or more of the connected phys are attached to an address different from the others. • An unsupported device is plugged in to this SAS port.

Storwize V7000 2076-524 node canister battery status LEDs

Table 15 on page 19 show battery status LEDs and their location in Figure 22 on page 16.

Table 15. Storwize V7000 2076-524 battery status LEDs

Name	Call out	Color	State	Meaning
Battery status	5	Green	OFF	Indicates that the battery is not available for use. The battery might be missing or a battery fault was detected.
			FAST BLINK	The battery has insufficient charge to complete a “fire hose” dump.
			BLINK	The battery has sufficient charge to complete a single “fire hose” dump.
			ON	The battery has sufficient charge to complete at least two “fire hose” dumps.
Battery fault	6	Amber	OFF	No fault. An exception to this would be where a battery has insufficient charge to complete a single “fire hose” dump. Refer to the documentation for the Battery status LED.
			ON	A battery fault was detected.

Storwize V7000 2076-524 node canister system status LEDs

Table 16 on page 20 show system status LEDs and their location in Figure 22 on page 16.

Table 16. Storwize V7000 2076-524 node canister system status LEDs

Name	Call out	Color	State	Meaning
Power	7	Green	OFF	No power is available or power is coming from the battery.
			SLOW BLINK	Power is available but the main processor is not running; this state is called <i>standby mode</i> .
			FAST BLINK	In self-test.
			ON	Power is available and the system code is running.
Status	8	Green	OFF	The system code has not started. The system is off, in standby, or in self-test.
			BLINK	The canister is in candidate or service state. It is not completing I/O operations. It is safe to remove the node.
			FAST BLINK	The canister is active, able to complete I/O operations, or starting.
			ON	The canister is active, able to complete I/O operations, or starting. The node is part of a cluster.

Table 16. Storwize V7000 2076-524 node canister system status LEDs (continued)

Name	Call out	Color	State	Meaning
Canister fault	9	Amber	OFF	The canister can function as an active member of the system. If the node canister has a problem, it is not severe enough to stop the node canister from completing I/O operations.
			BLINK	The canister is being identified. There might or might not be a fault condition.
			ON	The node is in service state or an error exists that might be stopping the system code from starting. The node canister cannot become active in the system until the problem is resolved. You must determine the cause of the error before you replace the node canister. The error might be due to insufficient battery charge. To resolve this error, wait for the battery to charge.

USB ports in the Storwize V7000 Gen2 node canister

Two USB ports are located on each Storwize V7000 Gen2 node canister.

The USB ports are numbered 1 on top and 2 on the bottom as shown in Figure 23 on page 22. One port is used during installation.

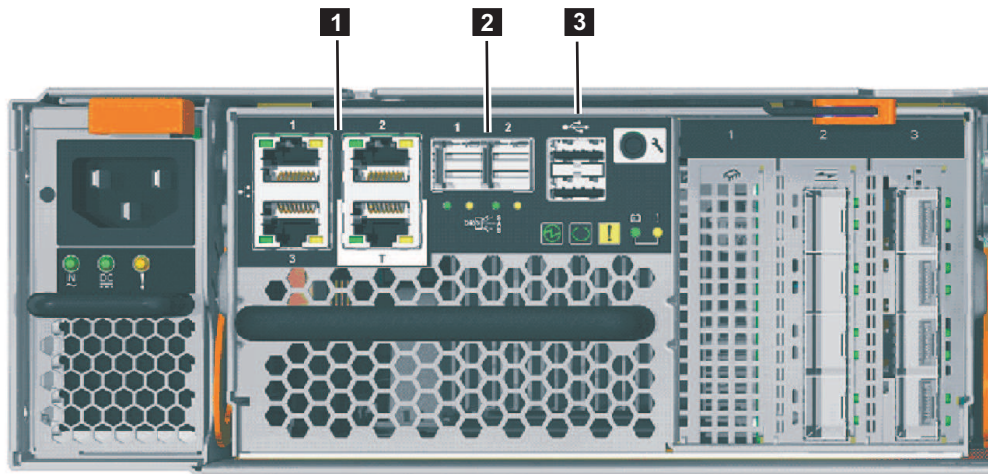


Figure 23. USB ports on the Storwize V7000 Gen2 node canister

- 1** Ethernet ports. The technician port is labeled T. It is used for troubleshooting or servicing the system.
- 2** Serial-attached SCSI (SAS) ports.
- 3** USB ports. Each canister has two USB ports.

The USB ports have no indicators.

Storwize V7000 2076-524 host interface adapter ports and indicators

An 8 Gbps or 10 Gbps host interface adapter can be installed in each node canister.

Storwize V7000 2076-524 Fibre Channel host interface adapter ports and indicators:

If you specified Fibre Channel host interface adapters for your Storwize V7000 2076-524 system, an adapter is preinstalled in each node canister.

Each 8 Gbps Fibre Channel 4-port host interface adapter (feature code ACHK) can have from two to four short wave (SW) small form-factor pluggable (SFP) transceivers installed. Cap unused ports with safety caps.

Fibre Channel host interface adapter ports

Fibre Channel ports **1** are in 1 - 4 order, starting at the top. Ports and their indicators are shown in Figure 24 on page 23.

Each port can have up to an 8 Gbps SW SFP transceiver installed. Each transceiver connects to a host or Fibre Channel switch with an LC-to-LC Fibre Channel cable.

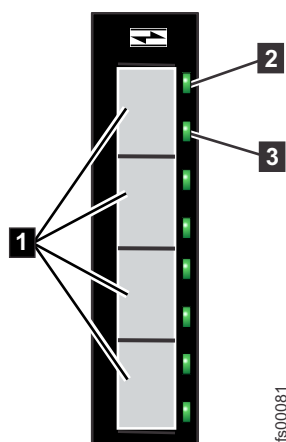


Figure 24. Fibre Channel ports and indicators

- 1** Fibre Channel 8 Gbps ports (x4)
- 2** Link-state LED (x4 - one for each port)
- 3** Speed-state LED (x4 - one for each port)

Fibre Channel host interface adapter indicators

Each Fibre Channel port has two green LED indicators. The link-state LED **2** is above the speed-state LED **3** for each port. Consider the LEDs as a pair to determine the overall link state, which is decoded in Table 17.

Table 17. Fibre Channel host interface adapter port-state LEDs

Link-state LED	Speed-state LED	Link state
OFF	OFF	Inactive
ON or FLASHING	OFF	Active low speed (2 Gbps)
ON or FLASHING	FLASHING	Active medium speed (4 Gbps)
ON or FLASHING	ON	Active high speed (8 Gbps)

One or two Fibre Channel interface adapters can be installed in each node canister. They can be installed in slots 2 and 3 of the node canister. When a single interface adapter is installed in either slot 2 or slot 3, the Fibre Channel ports on the adapter are numbered 1, 2, 3, and 4.

Storwize V7000 2076-524 10 Gbps Fibre Channel over Ethernet/iSCSI host interface adapter ports and indicators:

You have the option to install 10 Gbps Fibre Channel over Ethernet/iSCSI host interface adapters for your Storwize V7000 2076-524 system, the adapter is preinstalled in each node canister.

The 4 port FCoE/iSCSI host interface adapter is used for Fibre Channel over Ethernet (FCoE) or Internet Small Computer System Interface (iSCSI) connections to host systems or for Fibre Channel over Ethernet connections to host system or storage systems. Each port can support simultaneous FCoE and iSCSI connections. The Small Form-factor Pluggable (SFP) transceivers that are installed on the adapter support data transfer speeds of 10 Gbps.

Note: This adapter can be installed only in slots 2 and 3. Figure 25 shows two 10 Gbps Fibre Channel over Ethernet/iSCSI host interface adapters, both installed in slot 3.

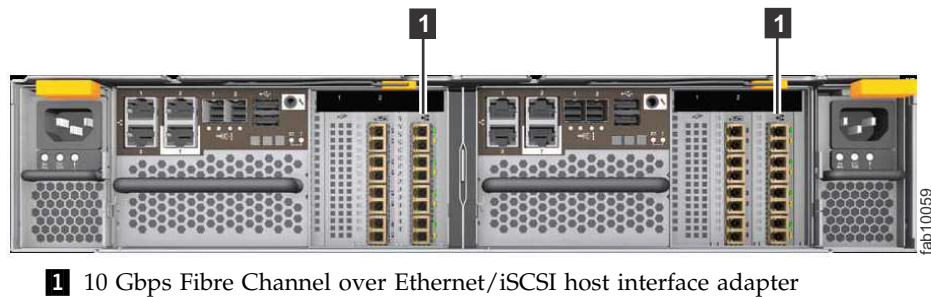


Figure 25. Example of installed 10 Gbps Fibre Channel over Ethernet/iSCSI host interface adapters

Storwize V7000 2076-524 10 Gbps Fibre Channel over Ethernet/iSCSI host interface adapter ports

The adapter has four Ethernet ports, none of which are used for system management. The ports are named 1, 2, 3 and 4 (Figure 26) when installed in a slot.



Figure 26. 10 Gbps Fibre Channel over Ethernet/iSCSI host interface adapter ports

Storwize V7000 2076-524 10 Gbps Fibre Channel over Ethernet/iSCSI host interface adapter indicators

Each port has two LED indicators, one green and one amber (see Figure 27 on page 25).

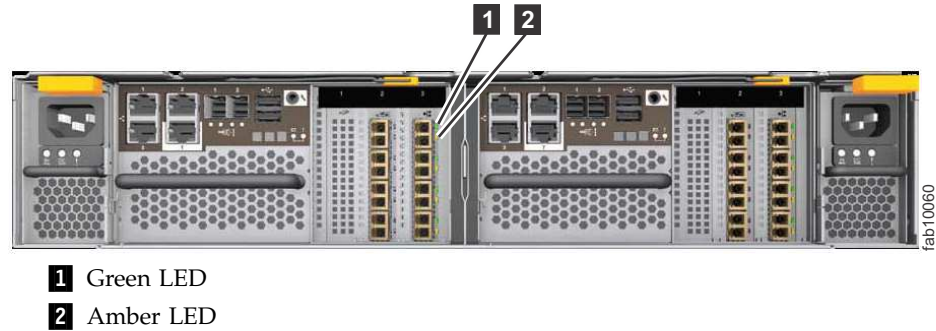


Figure 27. 10 Gbps Fibre Channel over Ethernet/iSCSI host interface adapter indicator LEDs

The LED states and their meanings are explained in Table 18.

Table 18. Storwize V7000 2076-524 host interface adapter LED states and meanings

Green LED	Amber LED	Meaning
OFF	OFF	The port is not configured in flex hardware and the port is not active in the current profile. For example, in the 2-by-16 Gbps profile, two ports are not active.
OFF	ON	The port is configured, but is not connected or the negotiation of the link failed (the link is not detected at the transport layer).
ON	OFF	The link is up and is running at the configured speed. Note: This does not indicate logical connectivity, such as the completion of FLOGI (Fabric login) or FIP (Fibre Channel over Ethernet Initialization Protocol).
ON	ON	The link is up and is running at less than the configured (degraded) speed.

Node canister ports and indicators

The node canister has indicators and ports but no controls.

Fibre Channel ports and indicators

The Fibre Channel port LEDs show the speed of the Fibre Channel ports and activity level.

Each node canister has four Fibre Channel ports located on the left side of the canister as shown in Figure 28 on page 26. The ports are in two rows of two ports. The ports are numbered 1 - 4 from left to right and top to bottom.

Note: The reference to the left and right locations applies to canister 1, which is the upper canister. The port locations are inverted for canister 2, which is the lower canister.

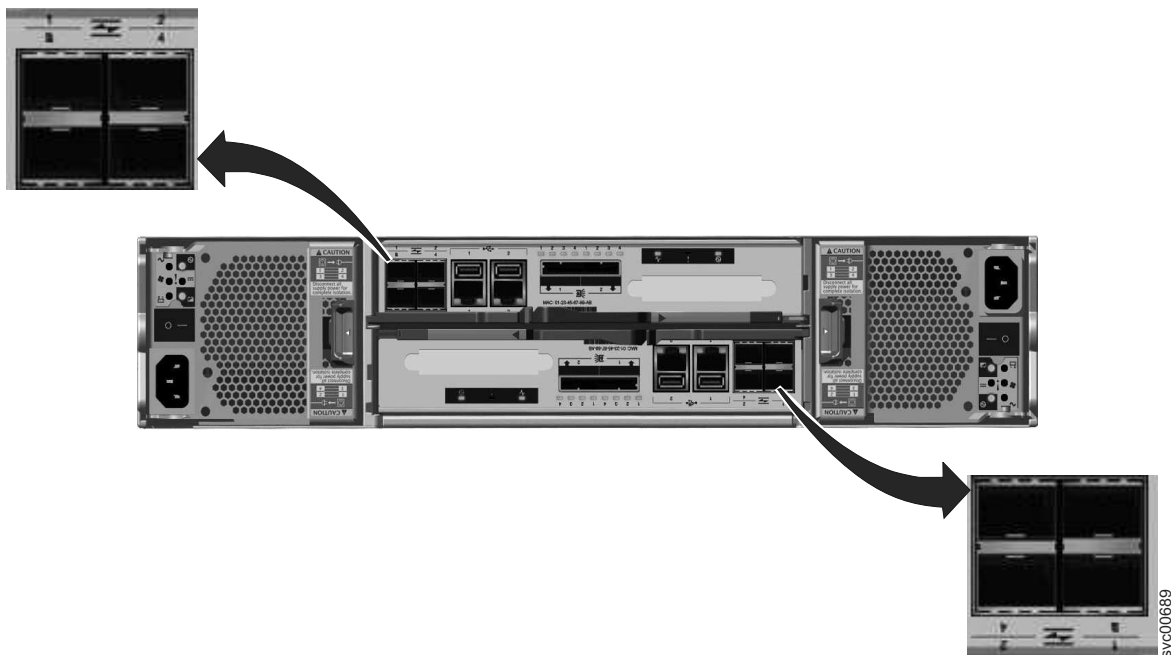


Figure 28. Fibre Channel ports on the node canisters

There are two green LEDs associated with each port: the speed LED and the link activity LED. These LEDs are in the shape of a triangle. The LEDs are located in between the two rows of the ports as shown in Figure 29. Figure 29 shows the LEDs for the Fibre Channel ports on canister 1. Each LED points to the associated port. The first and second LEDs in each set show the speed state, and the third and fourth LEDs show the link state.

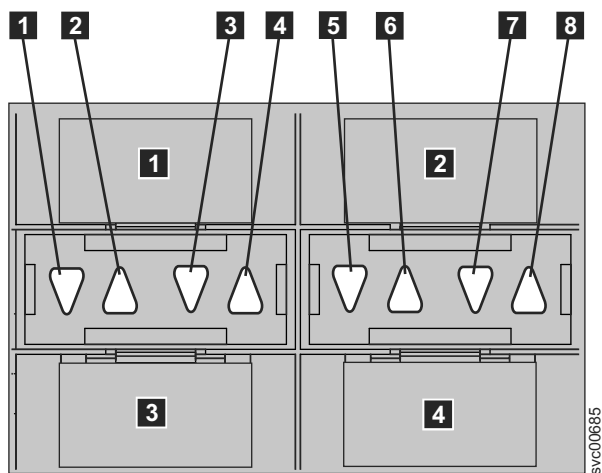


Figure 29. LEDs on the Fibre Channel ports

Table 19. Fibre Channel port LED locations on canister 1

Associated port	LED location	LED status
Port 3 3	First LED between ports 1 and 3 1	Speed
Port 1 1	Second LED between ports 1 and 3 2	Speed

Table 19. Fibre Channel port LED locations on canister 1 (continued)

Associated port	LED location	LED status
Port 3 3	Third LED between ports 1 and 3 3	Link
Port 1 1	Fourth LED between ports 1 and 3 4	Link
Port 4 4	First LED between ports 2 and 4 5	Speed
Port 2 2	Second LED between ports 2 and 4 6	Speed
Port 4 4	Third LED between ports 2 and 4 7	Link
Port 2 2	Fourth LED between ports 2 and 4 8	Link

Table 20 provides the status descriptions for the LEDs on the Fibre Channel ports.

Table 20. Fibre Channel port LED status descriptions

Speed state LED	Link state LED	Link state
Off	Off	Inactive
Off	On or flashing	Active low speed (2 Gbps)
Flashing	On or flashing	Active medium speed (4 Gbps)
On	On or flashing	Active high speed (8 Gbps)

USB ports in Storwize V7000 Gen1 node canisters

Two USB ports are located side by side on each Storwize V7000 Gen1 node canister.

The USB ports are numbered 1 on the left and 2 on the right as shown in Figure 30 on page 28. One port is used during installation.

Note: The reference to the left and right locations applies to canister 1, which is the upper canister. The port locations are inverted for canister 2, which is the lower canister.

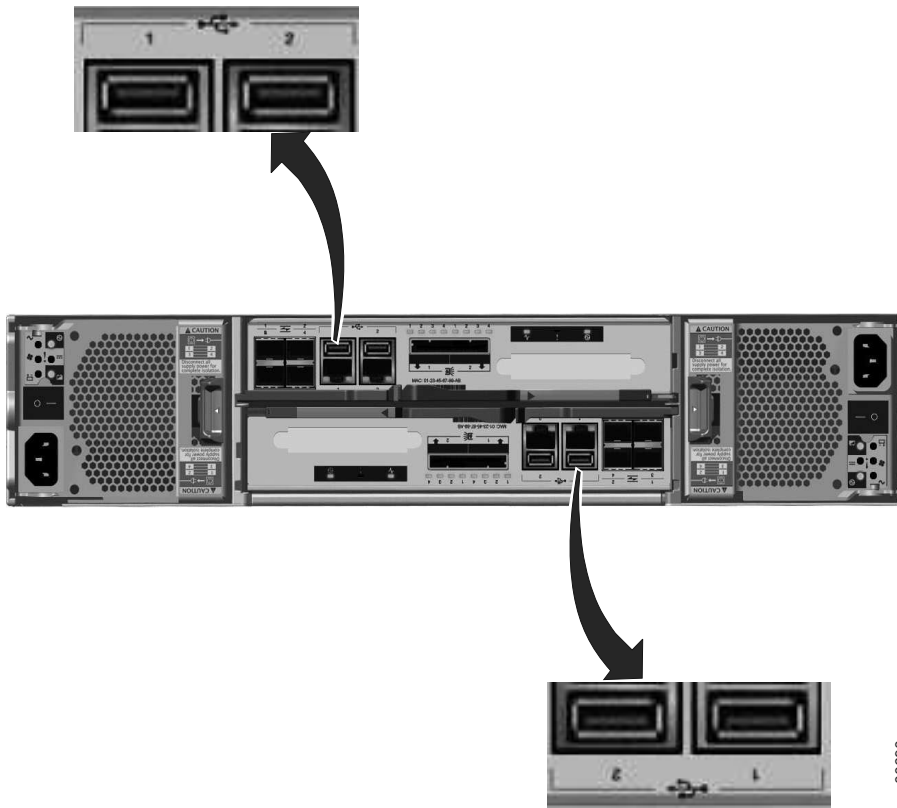


Figure 30. USB ports on the node canisters

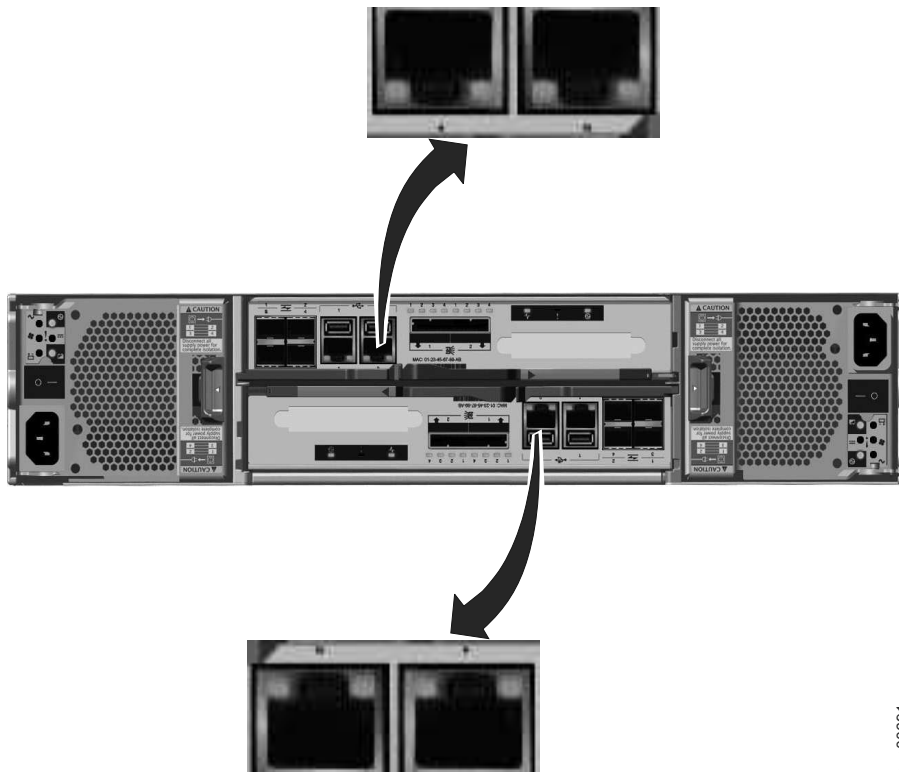
The USB ports have no indicators.

Ethernet ports and indicators

Ethernet ports are located side by side on the rear of the node canister. All control enclosure models have two 1 Gbps Ethernet ports per node canister. Model 2076-312 and model 2076-324 also have two 10 Gbps Ethernet ports per node canister.

For the 1 Gbps support, the Ethernet ports are numbered 1 on the left and 2 on the right as shown in Figure 31 on page 29. Port 1 must be connected; the use of port 2 is optional. Two LEDs are associated with each port.

Note: The reference to the left and right locations applies to canister 1, which is the upper canister. The port locations are inverted for canister 2, which is the lower canister.



svc00691

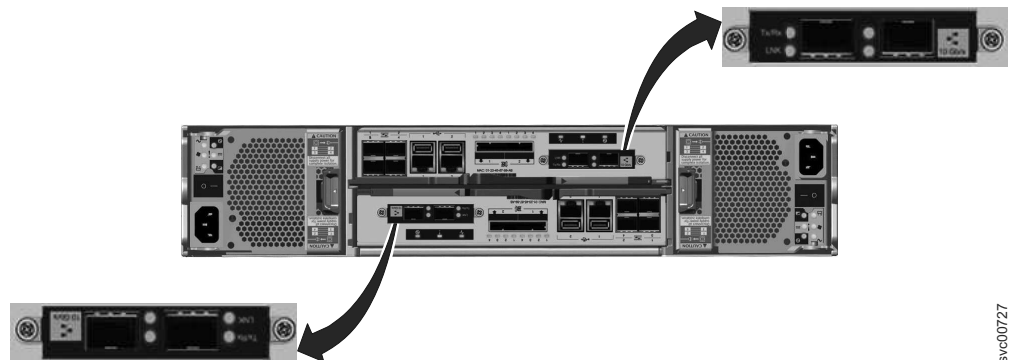
Figure 31. Ethernet ports on the 2076-112 and 2076-124 node canisters

Table 21 provides a description of the two LEDs.

Table 21. 1 Gbps Ethernet port LEDs

Name	Description	Color
Link speed (LED on right of upper canister)	The LED is on when there is a link connection; otherwise, the LED is off.	Green
Activity (LED on left of upper canister)	The LED is flashing when there is activity on the link; otherwise, the LED is off.	Yellow

Figure 32 shows the location of the 10 Gbps Ethernet ports.



svc00727

Figure 32. 10 Gbps Ethernet ports on the 2076-312 and 2076-324 node canisters

Table 22 on page 30 provides a description of the LEDs.

Table 22. 10 Gbps Ethernet port LEDs

Name	Symbol	Description	Color
Activity	Tx/Rx	The LED is flashing when there is activity on the link; otherwise, the LED is off.	Green
Link	LNK	The LED is on when there is a link connection; otherwise, the LED is off.	Amber

Node canister SAS ports and indicators

Two serial-attached SCSI (SAS) ports are located side by side in the rear of the node canister.

The SAS ports are numbered 1 on the left and 2 on the right as shown in Figure 33. Port 1 is used if you add one expansion enclosure. Port 2 is used if you add a second expansion enclosure. Each port provides four data channels.

Note: The reference to the left and right locations applies to canister 1, which is the upper canister. The port locations are inverted for canister 2, which is the lower canister.

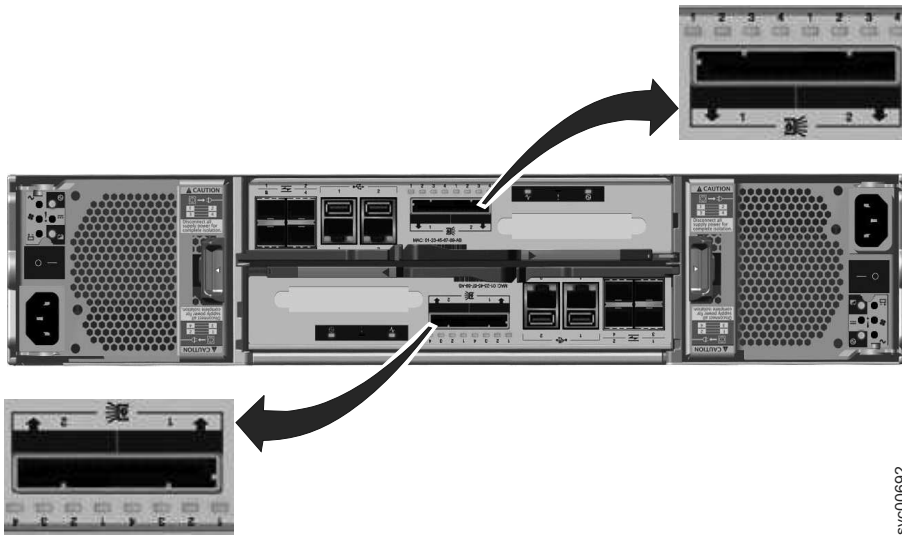


Figure 33. SAS ports on the node canisters.

SAS ports must be connected to Storwize V7000 Unified enclosures only. See “Problem: Storwize V7000 Gen1 SAS cabling not valid” on page 248 for help in attaching the SAS cables.

Four LEDs are located with each port. Each LED describes the status of one data channel within the port. The data channel number is shown with the LED.

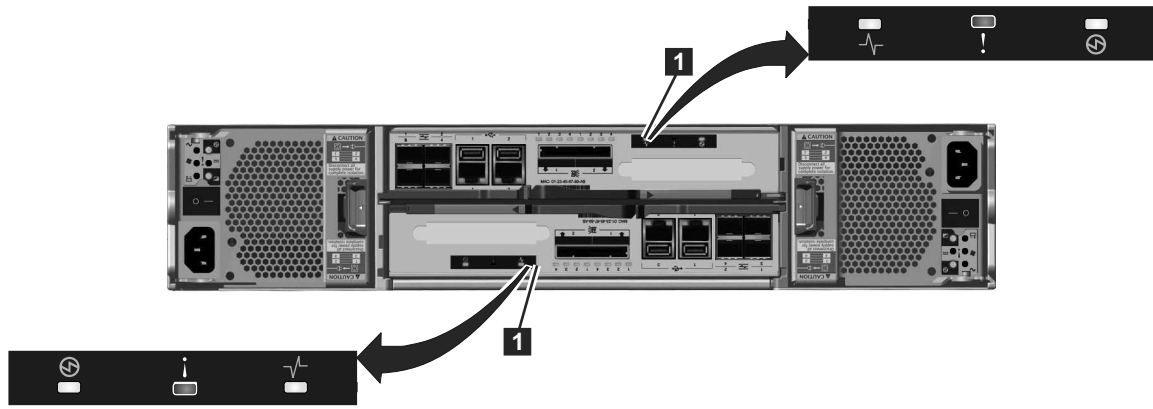
Table 23. SAS port LEDs on the node canister

LED state	Description
Off	No link is connected.
Flashing	The link is connected and has activity.
On	The link is connected.

Node canister LEDs

Each node canister has three LEDs that provide status and identification for the node canister.

The three LEDs are located in a horizontal row near the upper right of the canister **1**. Figure 34 shows the rear view of the node canister LEDs.



svc00672

Figure 34. LEDs on the node canisters

Note: The reference to the left and right locations applies to canister 1, which is the upper canister. The port locations are inverted for canister 2, which is the lower canister.

Table 24. Node canister LEDs

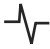
Name	Description	Color	Symbol
System status	<p>Indicates the status of the node.</p> <ul style="list-style-type: none"> The on status indicates that the node is active, that is, it is an active member of a clustered system. When the node is active, do not remove it. The off state indicates there is no power to the canister or the canister is in standby mode. These conditions can cause the off state: <ul style="list-style-type: none"> The main processor is off and only the service processor is active. A power-on self-test (POST) is running on the canister. The operating system is loading. The flashing status indicates that the node is in candidate state or service state. It is not able to perform I/O in a system. When the node is in either of these states, it can be removed. Do not remove the canister unless directed by a service procedure. 	Green	

Table 24. Node canister LEDs (continued)

Name	Description	Color	Symbol
Fault	<p>Indicates if a fault is present and identifies which canister.</p> <ul style="list-style-type: none"> The on status indicates that the node is in service state or an error exists that might be preventing the code from starting. Do not assume that this status indicates a hardware error. Further investigation is required before replacing the node canister. The off status indicates that the node is a candidate or is active. This status does not mean that there is not a hardware error on the node. Any error that was detected is not severe enough to stop the node from participating in a system. The flashing status indicates that the canister is being identified. This status might or might not be a fault. 	Amber	!
Power	<p>Indicates if power is available and the boot status of the canister.</p> <ul style="list-style-type: none"> The on status indicates that the canister is powered on and that the main processor or processors are running. The off status indicates that no power is available. The slow flashing (1 Hz) status indicates that power is available and that the canister is in standby mode. The main processor or processors are off and only the service processor is active. The fast flashing (2 Hz) indicates that the canister is running the power-on self-test (POST). 	Green	Ⓢ
<p>Notes:</p> <ol style="list-style-type: none"> If the system status LED is on and the fault LED is off, the node canister is an active member of a system. If the system status LED is on and the fault LED is on, there is a problem establishing a system. <p>For a more complete identification of the system LEDs, go to “Procedure: Understanding the Storwize V7000 Gen1 system status using the LEDs” on page 262.</p>			

Expansion canister ports and indicators

An expansion canister is one of two canisters that is located in the rear of the expansion enclosure. The expansion canister has no controls.

There is a diagnostic port on the right of the expansion canister. There are no indicators that are associated with the port. There are no defined procedures that use the port.

Storwize V7000 Gen2 expansion canister SAS ports and indicators

Two SAS ports are located in the rear of the Storwize V7000 Gen2 expansion canister.

SAS ports are numbered at the bottom of the port, with 1 on the left and 2 on the right, as shown in Figure 35. Use of port 1 is required. Use of port 2 is optional. Each port connects four data channels.

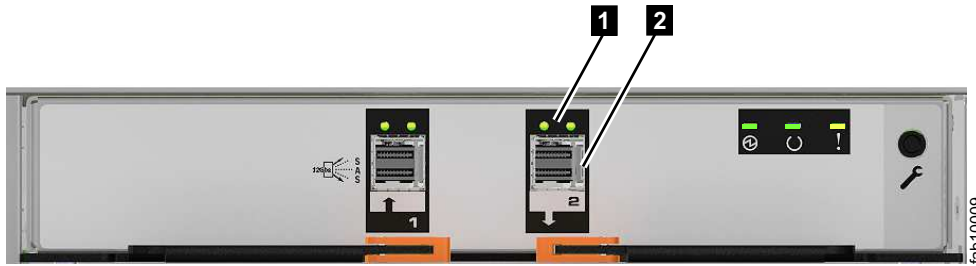


Figure 35. SAS ports and LEDs at rear of expansion canister

Figure 35 has callouts to show the location of the LEDs and the port for SAS port 2:

- 1** Port 2 LEDs
- 2** Port 2 12 Gbps SAS port

Table 25 describes LED states for each of the two LEDs per SAS port. The link LED is on the left of each set of ports.

Table 25. SAS port LEDs on the expansion canister

Name	Color	State	Meaning
SAS Port 1 Link	Green	OFF	No link connection on any phys (lanes). The connection is down.
		ON	There is a connection on at least one phy. At least one of the phys to that connector is up.
SAS Port 1 Fault	Amber	OFF	No fault. All four phys have a link connection.
		ON	This can indicate a number of different error conditions: <ul style="list-style-type: none"> One or more, but not all, of the 4 phys are connected. Not all 4 phys are at the same speed. One or more of the connected phys are attached to an address different from the others
SAS Port 2 Link	Green	OFF	No link connection on any phys (lanes). The connection is down.
		ON	There is a connection on at least one phy. At least one of the lanes to that connector is up.
SAS Port 2 Fault	Amber	OFF	No fault. All four phys have a link connection.
		ON	This can indicate a number of different error conditions: <ul style="list-style-type: none"> One or more, but not all, of the 4 phys are connected. Not all 4 phys are at the same speed. One or more of the connected phys are attached to an address different from the others

Storwize V7000 Gen1 expansion canister SAS ports and indicators

Two SAS ports are located in the rear of the Storwize V7000 Gen1 expansion canister.

SAS ports are numbered 1 on the left and 2 on the right as shown in Figure 36. Use of port 1 is required. Use of port 2 is optional. Each port connects four data channels.

Note: The reference to the left and right locations applies to canister 1, which is the upper canister. The port locations are inverted for canister 2, which is the lower canister.

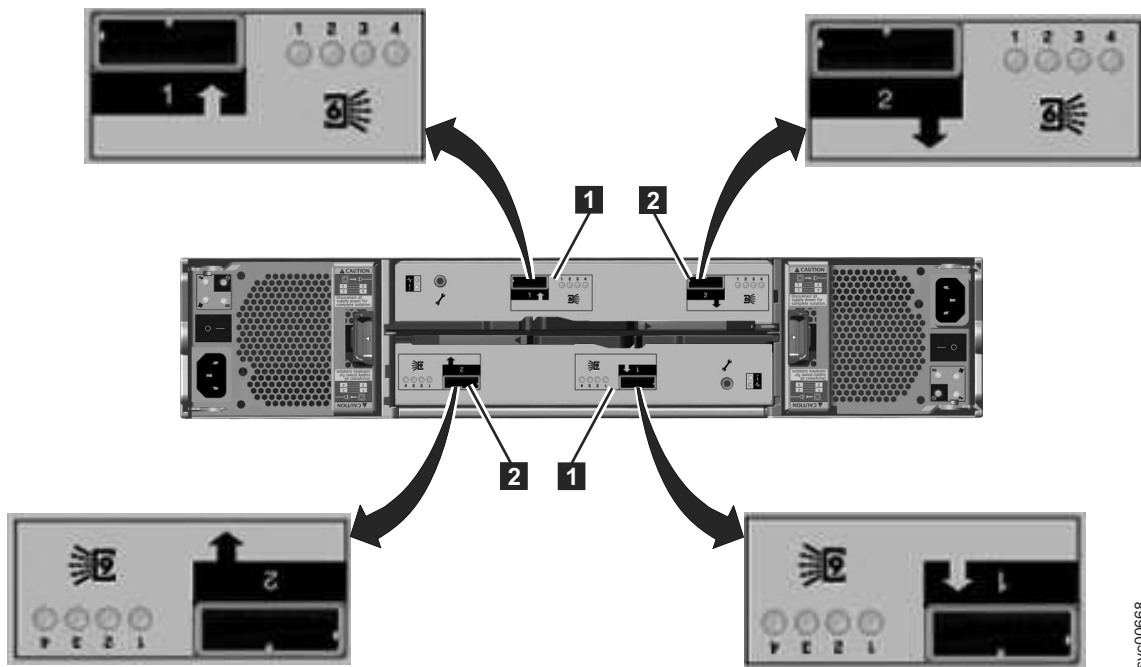


Figure 36. SAS ports and LEDs in rear of expansion enclosure

- **1** Port 1, 6 Gbps SAS port and LEDs
- **2** Port 2, 6 Gbps SAS port and LEDs

Four LEDs are located with each port. Each LED describes the status of one data channel within the port. The data channel is shown with the LED. Table 26 indicates the status of the SAS port LEDs on the expansion canister.

Table 26. SAS port LEDs on the expansion canister

LED state	Description
Off	No link is connected.
Flashing	The link is connected and has activity.
On	The link is connected.

Storwize V7000 Gen2 expansion canister LEDs

Each Storwize V7000 Gen2 expansion canister has three LEDs that provide status and identification for the expansion canister.

Three LEDs are located in a horizontal row on the right side (when viewed from the rear) of the expansion canister. Figure 37 shows the expansion canister LEDs, and Table 27 describes the LEDs.

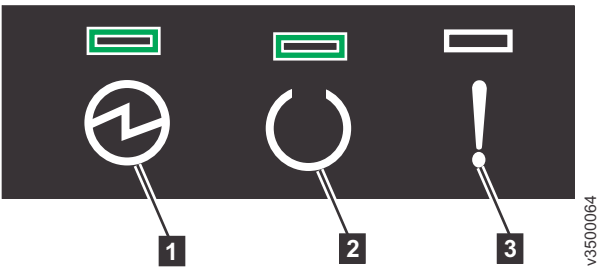


Figure 37. Expansion canister LEDs

Table 27. Expansion canister LED descriptions

Name	Description	Color	Symbol
1 Power	Indicates whether the expansion canister has power. <ul style="list-style-type: none">• If the LED is on, the canister has power.• If the LED is off, the canister does not have power.	Green	
2 Status	Indicates whether the expansion canister is active. <ul style="list-style-type: none">• If the LED is on, the canister is active.• If the LED is off, the canister is not active.• If the LED is flashing, there is a vital product data (VPD) error.	Green	
3 Fault	Indicates whether a fault is present and identifies the expansion canister. <ul style="list-style-type: none">• If the LED is on, a fault exists.• If the LED is off, no fault exists.• If the LED is flashing, the expansion canister is being identified. This status might or might not be a fault.	Amber	

Expansion canister LEDs

Each expansion canister has two LEDs that provide status and identification for the expansion canister.

The two LEDs are located in a vertical row on the left side of the canister. Figure 38 on page 36 shows the LEDs (**1**) in the rear of the expansion canister.

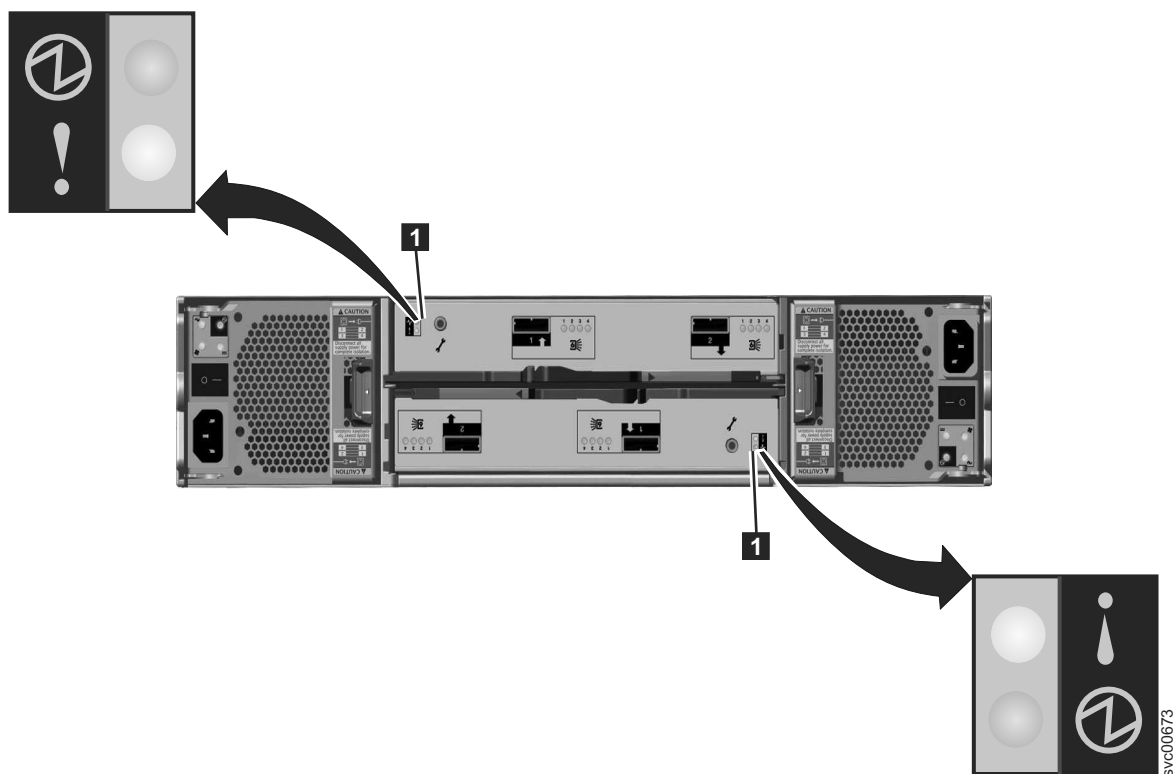


Figure 38. LEDs on the expansion canisters

Table 28 describes the expansion canister LEDs.

Table 28. Expansion canister LEDs

Name	Description	Color	Symbol
Status	<p>Indicates if the canister is active.</p> <ul style="list-style-type: none"> If the LED is on, the canister is active. If the LED is off, the canister is not active. If the LED is flashing, there is a vital product data (VPD) error. 	Green	⚡
Fault	<p>Indicates if a fault is present and identifies the canister.</p> <ul style="list-style-type: none"> If the LED is on, a fault exists. If the LED is off, no fault exists. If the LED is flashing, the canister is being identified. This status might or might not be a fault. 	Amber	!

Chapter 2. Best practices for troubleshooting

Taking advantage of certain configuration options, and ensuring vital system access information has been recorded, makes the process of troubleshooting easier.

Record the access information

It is important that anyone who has responsibility for managing the system know how to connect to and log on to the system. Give attention to those times when the normal system administrators are not available because of vacation or illness.

Record the following information in Table 29 and ensure that authorized people know how to access the information.

- The management IP addresses. This address connects to the system using the management GUI or starts a session that runs the command-line interface (CLI) commands. Record this address and any limitations regarding where it can be accessed from within your Ethernet network.
- The service IP addresses for the file module are used to access the root console on each of the file modules when needed to perform some investigation and recovery procedures.
- The root password for the file modules. The root password might be needed to perform some recovery procedures. For security reasons, the root password must be changed from its default value of `Passw0rd` using the **chrootpwd** CLI command. If you lose the root password, see “Recovering from losing the root password” on page 401.
- The control enclosure management IP address. This address is normally not needed. You might need it to access the control enclosure management GUI or the CLI during some recovery procedures. Use this address if the file modules lose their connection to the control enclosure CLI.
- The service IP addresses for the control enclosure canister. These addresses are normally not needed. You might need a service IP address to access the service assistant during some recovery procedures. Use this address if the control enclosure CLI is not working. These addresses are not set during the installation of a Storwize V7000 Unified system, but you can set these IP addresses later by using the management GUI or the **chserviceip** CLI command.

Table 29. Access information for your system

Item	Value	Notes
The management IP address for the management GUI and CLI		
The management user ID (the default is admin)		
The management user ID password (the default is admin001)		
The additional management user IDs and passwords that you create on your system		
The network gateway IP address		
File module 1 service IP address		

Table 29. Access information for your system (continued)

Item	Value	Notes
File module 2 service IP address		
The root password for the file modules (the default is Passw0rd)		
The control enclosure superuser IP address (not applicable to Storwize V7000 Gen2)		
The control enclosure superuser password (the default is passw0rd)		
Control enclosure 1 canister 1 service IP address		
Control enclosure 1 canister 2 service IP address		
Control enclosure 2 canister 1 service IP address		
Control enclosure 2 canister 2 service IP address		
Control enclosure 3 canister 1 service IP address		
Control enclosure 3 canister 2 service IP address		
Control enclosure 4 canister 1 service IP address		
Control enclosure 4 canister 2 service IP address		

Follow proper power management procedures

Access to your volume data can be lost if you incorrectly power off all or part of a system.

Use the management GUI or the CLI commands to power off a system. Using either of these methods ensures that the system fails properly in the case of powering down individual file modules and that data that is cached in the node canister memory is correctly flushed to the RAID arrays for the disk system.

The Storwize V7000 Unified system uses a pair of file modules for redundancy. Follow the appropriate power down procedures to minimize impacts to the system operations. See “Turning off the system” in the Storwize V7000 Unified information center.

Do not power off an enclosure unless instructed to do so. If you power off an expansion enclosure, you cannot read or write to the drives in that enclosure or to any other expansion enclosure that is attached to it from the SAS ports. Powering off an expansion enclosure can prevent the control enclosure from flushing all the data that it has cached to the RAID arrays.

Follow proper Storwize V7000 Gen2 power management procedures

You can lose access to volume data if you incorrectly power off all or part of a Storwize V7000 Gen2 system.

Always use the management GUI function to power off the system.

Only power off or remove a node canister if instructed to do so in a service action.

If you power off an expansion enclosure, you cannot read or write to the drives in that enclosure or to any other expansion enclosure that is attached to it from the SAS ports. Powering off an expansion enclosure can prevent the control enclosure from flushing all the data that it cached to the RAID arrays.

Follow proper power management procedures

Access to your volume data can be lost if you incorrectly power off all or part of a system.

Use the management GUI or the CLI commands to power off a system. Using either of these methods ensures that the system fails properly in the case of powering down individual file modules and that data that is cached in the node canister memory is correctly flushed to the RAID arrays for the disk system.

The Storwize V7000 Unified system uses a pair of file modules for redundancy. Follow the appropriate power down procedures to minimize impacts to the system operations. See “Turning off the system” in the Storwize V7000 Unified information center.

Do not power off an enclosure unless instructed to do so. If you power off an expansion enclosure, you cannot read or write to the drives in that enclosure or to any other expansion enclosure that is attached to it from the SAS ports. Powering off an expansion enclosure can prevent the control enclosure from flushing all the data that it has cached to the RAID arrays.

Set up event notifications

Configure your system to send notifications when a new event is reported.

Correct any issues reported by your system as soon as possible. To avoid monitoring for new events by constantly monitoring the management GUI, configure your system to send notifications when a new event is reported. Select the type of event that you want to be notified about. For example, restrict notifications to just events that require immediate action. Several event notification mechanisms exist:

- Email. An event notification can be sent to one or more email addresses. This mechanism notifies individuals of problems. Individuals can receive notifications wherever they have email access which includes mobile devices.
- Simple Network Management Protocol (SNMP). An SNMP trap report can be sent to a data-center management system, such as IBM Systems Director, that consolidates SNMP reports from multiple systems. Using this mechanism, you can monitor your data center from a single workstation.
- Syslog. A syslog report can be sent to a data-center management system that consolidates syslog reports from multiple systems. Using this mechanism, you can monitor your data center from a single workstation.

- **Call Home.** If your system is within warranty, or you have a hardware maintenance agreement, configure your system to send email events to IBM if an issue that requires hardware replacement is detected. This mechanism is called Call Home. When the event is received, IBM automatically opens a problem report, and if appropriate, contacts you to verify if replacement parts are required. If you set up Call Home to IBM, ensure that the contact details that you configure are correct and kept up to date as personnel change.

Set up inventory reporting

Inventory reporting is an extension to the Call Home email.

Rather than reporting a problem, an email is sent to IBM that describes your system hardware and critical configuration information. Object names and other information, such as IP addresses, are not sent. The inventory email is sent on a regular basis. Based on the information that is received, IBM can inform you if the hardware or software that you are using requires an update because of a known issue.

Back up your data

Back up your system configuration data, volume data, and file systems.

The file modules back up their configuration after each configuration change. Download the backup files regularly to your management workstation to protect the data.

The storage system backs up your control enclosure configuration data to a file every day. This data is replicated on each control node canister in the system. Download this file regularly to your management workstation to protect the data. This file must be used if there is a serious failure that requires you to restore your system configuration. It is important to back up this file after modifying your system configuration.

Your volume data or files in the file systems are susceptible to failures in your host application or your Storwize V7000 Unified system. Follow a backup and archive policy that is appropriate to the data that you have for storing the volume data on a different system or the files on a different system.

Manage your spare and failed drives

Your RAID arrays that are created from drives consist of drives that are active members and drives that are spares.

The spare drives are used automatically if a member drive fails. If you have sufficient spare drives, you do not have to replace them immediately when they fail. However, monitoring the number, size, and technology of your spare drives, ensures that you have sufficient drives for your requirements. Ensure that there are sufficient spare drives available so that your RAID arrays are always online.

Resolve alerts in a timely manner

Your system reports an alert when there is an issue or a potential issue that requires user attention.

The management GUI provides the capability to review these issues from the Events panel.

For file module issues, use the Storwize V7000 Unified information center to look up the events and perform the actions listed for the events.

For Storwize V7000 issues, resolve these problems through the **Recommended actions only** option from the Events panel.

Complete the recommended actions as quickly as possible after the problem is reported. Your system is designed to be resilient to most single hardware failures. However, if you operate for any period of time with a hardware failure, the possibility increases that a second hardware failure can result in some volume data that is unavailable.

If there are a number of unfixed alerts, fixing any one alert might become more difficult because of the effects of the other alerts.

Keep your software up to date

Check for new code releases and update your code on a regular basis.

This can be done using the management GUI, or by checking the IBM support website to see if new code releases are available:

www.ibm.com/storage/support/storwize/v7000/unified

The release notes provide information about new function in a release plus any issues that have been resolved. Update your code regularly if the release notes indicate a potential issue.

Keep your records up to date

Follow the proper record keeping procedures for your system. Record management procedures differ, depending on the generation of your control enclosure model.

About this task

Storwize V7000 Unified *Gen1* refers to the enclosure models in the following table:

Table 30. Storwize V7000 Unified Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 Unified control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 Unified control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 Unified control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 Unified control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 Unified expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Storwize V7000 Unified *Gen2* refers to the newer generation of enclosures in the following table:

Table 31. Storwize V7000 Unified Gen2 model numbers

Machine type/model	Description
2076-524	Storwize V7000 Unified control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-12F	Storwize V7000 Unified expansion enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Keep your Storwize V7000 Gen2 records up to date

Keep good records of the system location, names, and management addresses. Record the location information for your enclosures.

If you have only one system, it is relatively easy to identify the enclosures that make up the system. Identification becomes more difficult when you have multiple systems in your data center and multiple systems in the same rack.

Record the MT-M and serial number of each Storwize V7000 Gen2 enclosure. The information can be found on the IBM Standard Asset Tag attached to the left enclosure bezel, which includes a machine-readable data matrix to ISO/IEC 15434.

For each system, record the location of the control enclosure and the location of any expansion enclosures. It is useful to label the enclosures themselves with the system name and the management IP addresses.

Keep your Storwize V7000 Unified Gen1 records up to date

Record the location information for your enclosures and file modules.

If you have only one system, it is relatively easy to identify the enclosures that make up the system. Identification becomes more difficult when you have multiple systems in your data center and multiple systems in the same rack.

For each system, record the location of the file modules, control enclosure, and any expansion enclosures. It is useful to label the enclosures themselves with the system name and the management IP addresses.

Subscribe to support notifications

Subscribe to support notifications so that you are aware of best practices and issues that might affect your system.

Subscribe to support notifications by visiting the IBM support page on the IBM website:

www.ibm.com/storage/support/storwize/v7000/unified

By subscribing, you are informed of new and updated support site information, such as publications, hints and tips, technical notes, product flashes (alerts), and downloads.

Know your IBM warranty and maintenance agreement details

If you have a warranty or maintenance agreement with IBM, know the details that must be supplied when you call for support.

Have the phone number of the support center available. When you call support, provide the machine type and the serial number of the enclosure or file module that has the problem. The machine type is always 2076 for a control enclosure or 2073 for a file module. If the problem does not relate to a specific enclosure, provide the control enclosure serial number. The serial numbers are on the labels on the enclosures.

Support personnel also ask for your customer number, machine location, contact details, and the details of the problem.

How to get information, help, and technical assistance

If you need help, service, technical assistance, or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you.

Information

IBM maintains pages on the web where you can get information about IBM products and fee services, product implementation and usage assistance, break and fix service support, and the latest technical information. For more information, refer to Table 4 on page xxv.

Table 32. IBM websites for help, services, and information

Website	Address
IBM home page	http://www.ibm.com
Directory of worldwide contacts	http://www.ibm.com/planetwide
Support for Storwize V7000 (2076)	www.ibm.com/storage/support/storwize/v7000
Support for Storwize V7000 Unified (2073)	www.ibm.com/storage/support/storwize/v7000/unified
Support for IBM System Storage and IBM TotalStorage products	www.ibm.com/storage/support/

Note: Available services, telephone numbers, and web links are subject to change without notice.

Help and service

Before calling for support, be sure to have your IBM Customer Number available. If you are in the US or Canada, you can call 1 (800) IBM SERV for help and service. From other parts of the world, see <http://www.ibm.com/planetwide> for the number that you can call.

When calling from the US or Canada, choose the **storage** option. The agent decides where to route your call, to either storage software or storage hardware, depending on the nature of your problem.

If you call from somewhere other than the US or Canada, you must choose the **software** or **hardware** option when calling for assistance. Choose the **software** option if you are uncertain if the problem involves the Storwize V7000 Unified software or hardware. Choose the **hardware** option only if you are certain the problem solely involves the Storwize V7000 Unified hardware. When calling IBM for service regarding the product, follow these guidelines for the **software** and **hardware** options:

Software option

Identify the Storwize V7000 Unified product as your product and supply your customer number as proof of purchase. The customer number is a 7-digit number (0000000 to 9999999) assigned by IBM when the product is purchased. Your customer number should be located on the customer information worksheet or on the invoice from your storage purchase. If asked for an operating system, use **Storage**.

Hardware option

Provide the serial number and appropriate 4-digit machine type. For Storwize V7000 Unified, the machine type is 2073.

In the US and Canada, hardware service and support can be extended to 24x7 on the same day. The base warranty is 9x5 on the next business day.

Getting help online

You can find information about products, solutions, partners, and support on the IBM website.

To find up-to-date information about products, services, and partners, visit the IBM website at www.ibm.com/storage/support/storwize/v7000/unified.

Before you call

Make sure that you have taken steps to try to solve the problem yourself before you call.

Some suggestions for resolving the problem before calling IBM Support include:

- Check all cables to make sure that they are connected.
- Check all power switches to make sure that the system and optional devices are turned on.
- Use the troubleshooting information in your system documentation. The troubleshooting section of the information center contains procedures to help you diagnose problems.
- Go to the IBM Support website at www.ibm.com/storage/support/storwize/v7000/unified to check for technical information, hints, tips, and new device drivers or to submit a request for information.

Using the documentation

Information about your IBM storage system is available in the documentation that comes with the product.

That documentation includes printed documents, online documents, readme files, and help files in addition to the information center. See the troubleshooting information for diagnostic instructions. The troubleshooting procedure might require you to download updated device drivers or software. IBM maintains pages

on the web where you can get the latest technical information and download device drivers and updates. To access these pages, go to www.ibm.com/storage/support/storwize/v7000/unified and follow the instructions. Also, some documents are available through the IBM Publications Center.

Sign up for the Support Line Offering

If you have questions about how to use and configure the machine, sign up for the IBM Support Line offering to get a professional answer.

The maintenance supplied with the system provides support when there is a problem with a hardware component or a fault in the system machine code. At times, you might need expert advice about using a function provided by the system or about how to configure the system. Purchasing the IBM Support Line offering gives you access to this professional advice while deploying your system, and in the future.

Contact your local IBM sales representative or your support group for availability and purchase information.

Chapter 3. Getting started troubleshooting

This topic is an entry point to troubleshooting your system. The content provides help in correctly identifying which of the recovery procedures must be run to recover a Storwize V7000 Unified system from a problem.

About this task

Important: After you successfully fix a problem by using the instructions that follow, use the Health status and recovery procedure to set the health status back to green.

If you are here because you installed a new system and cannot initialize it by using the USB flash drive, go to “Installation troubleshooting” on page 48.

If one of the file modules does not boot up and join the GPFS™ cluster, look for a hardware problem by using the light-path diagnostics LEDs. See “File node hardware indicators for 2073-720” on page 77. If you suspect that the boot software is corrupted, call IBM support.

If any orange fault LEDs are illuminated on the control enclosure, front or rear, see “Resolving a problem” on page 236.

If you are having problems accessing the management GUI or the CLI, see “GUI access issues” on page 60. For information about accessing the management GUI, see “Accessing the Storwize V7000 Unified management GUI” on page 89.

If the health status indicator in the lower right corner of management GUI is not green, hover over the icon on the left side of the indicator to see the type of error that is causing the poor health status. Select an error type, and you are shown the critical errors in the event log. First try to fix the critical errors under the **Block** tab of the **Monitoring > Events** page before trying to fix the critical errors under the **File** tab of the **Monitoring > Events** page.

Log into the CLI interface and run the CLI command, **lslog**. Review the results for problems that may need to be resolved.

If users or applications are having trouble accessing data that is held on the Storwize V7000 Unified system, or if the management GUI is not accessible or is running slowly, the Storwize V7000 control enclosure might have a problem.

If you cannot ping the system IP address for the Storwize V7000 control enclosure, try to access the control enclosure service assistant. Use the service IP address of the node canisters in the control enclosure or the technician port to resolve any reported node errors. See “Procedure: Fixing node errors” on page 271.

Note: Use the access information that you previously recorded for the service IP address of the node canisters in the control enclosure. See “Record the access information” on page 37. If you do not know the service IP addresses for the node canisters in the control enclosure, see “Problem: node canister service IP address unknown” on page 244.

If all nodes show either node error 550 or node error 578, you might need to perform a system recovery. See “Recover system procedure” on page 381 for more details.

For more information about determining and solving block storage problems that relate to the control enclosure, see “Resolving a problem” on page 236.

Check the intrasystem connectivity by using the management GUI. Navigate to **Monitoring > System**. Use the interactive graphic to determine the connection state by hovering over each connection in the graphic.

If either of the Fibre Channel links from the file modules shows an error or degraded state, see “Fibre Channel connectivity between file modules and control enclosure” on page 72.

If mgmt0, the direct Ethernet link between the file modules, shows an error or degraded state, see “Ethernet connectivity between file modules” on page 65.

If one or both of the Fibre Channel links from the file module to the control enclosure show an error or a degraded state, see “Ethernet connectivity from file modules to the control enclosure” on page 68.

Check the core component health. Navigate to **Monitoring > System Details > Interface Nodes > *nodename* > NAS Services**. In the Status panel, check the CTDB state and the GPFS state.

If the GPFS state is Active, but the CTDB state is not Active, see “Checking CTDB health” on page 187; otherwise, see “Checking the GPFS file system mount on each file module” on page 189.

If you have lost access to the files, but there is no sign that anything is wrong with the Storwize V7000 Unified system, see “Host to file modules connectivity” on page 63.

Installation troubleshooting

This topic provides information for troubleshooting problems encountered during the installation.

Software issues are often reported through CLI commands at system configuration and through error codes. Power problems can often be solved through identifying visual symptoms.

Problems with initial setup

This topic helps you to solve initial setup problems.

About this task

If USB flash drive is missing or faulty:

- Contact the IBM Support Center.
- Install the latest InitTool.exe (or reinstall if tool is not launching). Go to <http://www-933.ibm.com/support/fixcentral/options> and select the following options to locate the tool. The options are listed under the **Select product** tab, at the bottom of the page:
 - Product Group: **Storage Systems**

- Product Family: **Disk Systems**
- Product: **IBM Storwize V7000 Unified**
- Release: **All**
- Platform: **All**

Before loading the USB flash drive verify it has a FAT32 formatted file system. Plug the USB flash drive into the laptop. Go to Start (my computer), right-click the USB drive. The general tab next to File system should say FAT32.

- If the USB flash drive is not formatted as FAT32, format it. To format, right-click it, select format, under filesystem. Select FAT32 and then click Start. Continue as prompted.

InitTool.exe is not loaded on the USB flash drive or fails to launch:

- Install the latest InitTool.exe (or reinstall if tool is not launching). Go to <http://www-933.ibm.com/support/fixcentral/options> and select the following options to locate the tool. The options are listed under the **Select product** tab, at the bottom of the page:
 - Product Group: **Storage Systems**
 - Product Family: **Disk Systems**
 - Product: **IBM Storwize V7000 Unified**
 - Release: **All**
 - Platform: **All**

Amber LED on node canister does not stop flashing during install:

Allow at least 15 minutes for the LED to stop flashing. If flashing continues beyond 15 minutes, remove the USB flash drive and insert in your laptop. Navigate to the `satask_results.html` file and scan for errors and follow the service action recommendation. Take that action and retry installation.

An error is posted in the `satask_results.html`:

Take the recommended service action given by **sainfo lsservicerecommendation** in the `satask_results.html` file, reboot the node, and restart the initial setup procedure.

If `satask_results.html` contains node error code 835 or node error code 550 then this can indicate that the node canisters were not able to communicate with each other at some time during the creation of the block cluster. This can occur because the PCIe link between the node canisters is temporarily broken when the nodes are restarted, as part of the create cluster process. This can generate node error codes 835 and 550. These are transitional errors that can be ignored if the nodes are now in active state with no errors. Follow this procedure to check that the errors are gone, using the USB flash drive:

- Save a copy of `satask.txt` and `satask_results.html`.
- Make sure that there is no `satask.txt` file on the USB flash drive before you plug it into the control enclosure. Plug the USB flash drive into the control enclosure. The orange fault light should go on for a short time only (such as a slow blink for a few seconds). Wait for the orange fault light to go out then unplug the USB flash drive and plug it into another computer so that you can look at the contents of the `satask_results.html` file on the USB flash drive. The `satask_results.html` will contain the output from a number of `sainfo` commands.
- Check the following:

- The cluster_status under **sainfo lsservicenodes** should be Active.
- The node_status should be Active for both node canisters in the cluster under **sainfo lsservicenodes**. Otherwise, follow the service action under **sainfo lsservicerecommendation**.
- There should be nothing in the error_data column against each node under **sainfo lsservicenodes**. Otherwise, follow the service action under **sainfo lsservicerecommendation**.

This is an example of what the satask_results.html can contain on a healthy storage system, with which you can compare your results:

```
Service Command Results
Thu Apr 19 08:23:42 UTC 2012
satask.txt file not found.

System Status

sainfo lsservicenodes
panel_name cluster_id cluster_name node_id node_name relation node_status
error_data
01-1 00000200A4E008BA Cluster_9.71.18.184 1 node1 local Active
01-2 00000200A4E008BA Cluster_9.71.18.184 2 node2 partner Active
sainfo lsservicestatus
panel_name 01-1
cluster_id 00000200a4e008ba
cluster_name Cluster_9.71.18.184
cluster_status Active
cluster_ip_count 2
cluster_port 1
cluster_ip 9.71.18.184
cluster_gw 9.71.18.1
cluster_mask 255.255.255.0
...
...
sainfo lsservicerecommendation
service_action
No service action required, use console to manage node.
```

Blue LED on file module, where the USB flash drive was inserted, keeps flashing (does not turn solid as stated in the instructions):

- Allow 5 minutes at least, remove the USB flash drive, insert it into your laptop. Verify that the InitTool set up information is correct, navigate to the SONAS_results.txt file, and open it. Check for errors and corrective actions. Refer to *Storwize V7000 Unified Problem Determination Guide* PDF on the CD.
- If no errors are listed, reboot the server (allow server to start), reinsert the USB flash drive, and try again.

Blue LED on the other file module (without USB flash drive) keeps flashing (does not turn solid or off as listed in instructions):

Wait for the primary file module to start flashing, remove the USB flash drive, insert it into you laptop, verify the InitTool set up information is correct , navigate to the SONAS_results.txt file and open it. Check for errors and corrective actions (refer to *Storwize V7000 Unified Problem Determination Guide* PDF on the CD). If no errors are listed, reboot both file modules, allow file modules to boot completely, reinsert the USB flash drive as originally instructed and try again.

Installed with the incorrect control enclosure or file module IP addresses:

If it is determined that the addresses were entered incorrectly, they can be changed at the command line as user **admin** with the following commands:

- For control enclosure IP changes use: **svctask chsystemip**
- For file module management node changes use: **chnwmgt**

Refer to the man pages for usage.

The file module initialization may have failed because of a duplicate IP address:

The control enclosure may have been set up with an IP address which is already in use by another machine on your network but the initial setup of the file modules has failed. Refer to Checking that IP addresses are not already in use from the Information Center, under the Installing topic.

Installation error codes

The system generates an error code that provides a recommended action if the installation fails.

Guide to using the error code table

1. Always check the entire system for any illuminated error lights first and refer to the problem systems appropriate maintenance manual. If no lights are illuminated, continue to step 2.
2. Match the error code noted in the results.txt file to the installation error codes list in Table 34 on page 52. If there are multiple errors, the first error listed is the most critical and should be addressed first.
3. Refer to Table 33 to match the code (A-I) to the recommended action. Follow the suggested action, in order, completing one before trying the next.
4. If the recommended action or actions fail, call the IBM Support Center.

.

Table actions defined

This table serves as a legend for defining the precise action to follow. The action legend defines the action that is correlated with each action key.

Table 33. Installation error code actions

Action key	Action to be taken
A	Power cycle both file modules with the power button. Wait for the file modules to come up and the flashing blue light on each to come on before proceeding, then reinsert the USB flash drive into the original file module. The installation continues from the last good checkpoint.
B	Power down both file modules, remove power from the power source (unplug it), reapply power, power up, wait for the file modules to come up and the flashing blue light on each to come on before proceeding, then reinsert the USB flash drive into the original file module. The installation continues from the last good checkpoint.
C	Verify that the cabling between file modules is correct and that the connections are seated properly. Then reinsert the USB flash drive into the original file module. The installation will continue from the last good checkpoint.
D	Verify that all IP/gateway/subnet address information is correct (InitTool) and that there are no duplicate IP's on the network. If a change is made, reinsert the USB flash drive. The installation continues from the last good checkpoint.
E	Insert the USB flash drive into the other file module and try again

Table 33. Installation error code actions (continued)

Action key	Action to be taken
F	<p>The NAS private key is probably OK if you were able to start the management GUI, but if the USB flash drive step of initial set-up failed, then do the following:</p> <p>Retrieve the NAS private key from the Storwize V7000 by doing the following:</p> <ul style="list-style-type: none"> • Create a text file with the following line: satask chnaskey -privkeyfile NAS.ppk • Save the file as satask.txt on the USB flash drive. Insert the USB flash drive into one of the top control enclosure USB ports and wait at least 20 seconds. Reinsert the USB flash drive into the original management node. The installation continues from the last good checkpoint.
G	<p>Verify that the Ethernet cabling connections are seated properly between the Storwize V7000 Unified control enclosure and the customer network, as well as the file modules cabling to the customer network. Then press the Restart button if the management GUI has already started, otherwise, reinsert the USB flash drive into the original file module. The installation will continue from last good checkpoint.</p>
H	<p>This could be caused by a number of things so look in sonas_results.txt for an error code that could have caused this, and follow the recommended action. If there is no other error code in sonas_results.txt that could have caused this then refer to "Ethernet connectivity from file modules to the control enclosure" on page 68 for help troubleshooting the file module to control enclosure management connection.</p>
I	<p>Open the Storwize V7000 management GUI in a Web browser by using the Storwize V7000 system IP address. Go to the update Software panel in the Storwize V7000 management GUI. Follow the instruction in the GUI to download and run the Storwize V7000 software update checker for the latest Storwize V7000 software level that is compatible with the current Storwize V7000 Unified software level of the file modules. Repair any errors flagged by the software update checker. Then select the retry option in the Storwize V7000 Unified management GUI if that is working or reinsert the USB flash drive into the original file module. The installation will continue from last good checkpoint.</p>

Installation error codes

Table 34 lists the error messages and keyed actions. To match the actions, see Table 33 on page 51.

Table 34. Error messages and actions

Error code	Error message	Action key
0A01	Unable to open /tmp/setup_hosts_\$\$.	A
0A02	Unable to create default users.	A
0A05	Unable to determine management IP address.	A
0A06	Unable to determine Management Mask Address.	A
0A07	Error updating /etc/hosts.	A
0A08	Unable to update VPD field.	A
0A0A	Error opening /etc/sysconfig/network.	A
0A0B	Error writing /etc/sysconfig/network.	A
0A0C	Error updating host name.	A

Table 34. Error messages and actions (continued)

Error code	Error message	Action key
0A0D	Error querying settings through ASU.	B
0A0E	Error setting ASU command.	B
0A0F	Unable to determine adapter name from VPD.	A
0A10	Unable to open the ifcfg file.	A
0A11	Unable to write to the ifcfg file.	A
0A12	Unable to bring adapter down.	A
0A13	Unable to bring adapter up.	D then C then B
0A14	Unable to determine adapter name from VPD.	A
0A15	Unable to open the ifcfg-alias file.	A
0A16	Unable to write to the ifcfg-alias file.	A
0A17	Unable to bring adapter-alias down.	A
0A18	Unable to bring adapter-alias up.	D then C then B
0A19	Unable to retrieve adapter name.	A
0A1A	Incorrect parameters.	D
0A1B	Adapter value not valid.	A
0A1C	Alias value not valid.	A
0A1D	DHCP is not valid on this adapter.	A
0A1E	DHCP is not valid on aliases.	A
0A1F	Invalid IP address.	D
0A20	Invalid netmask.	D
0A21	Invalid Gateway IP Address.	D
0A22	Gateway, netmask, and IP are incompatible.	D
0A23	Gateway is not valid on this adapter.	D
0A24	Alias is null.	A
0A25	Could not drop aliases.	A
0A26	Invalid adapter for Storwize V7000.	A
0A27	Invalid alias state argument.	A
0AA5	Invalid inputs.	A
0AA6	Called with invalid host name.	A
0AA7	Error sending password.	A
0AA8	A node name was not provided.	A
0AA9	Invalid management IP address.	A
0AAB	Invalid RSA IP address.	A
0AAC	Invalid IP for management node.	A
0AAD	The node is already a part of a cluster.	A
0AAE	Error while setting storage node peer.	A
0AAF	Unable to get node roles from VPD.	A
0AB0	Error opening /etc/sysconfig/rsyslog.	A
0AB1	Error writing to /etc/sysconfig/rsyslog.	A
0AB2	Error reading /etc/rsyslog.conf.	A

Table 34. Error messages and actions (continued)

Error code	Error message	Action key
0AB3	Unable to open /opt/IBM/sonas/etc/rsyslog_template_mgmt.conf.	A
0AB4	Unable to open /opt/IBM/sonas/etc/rsyslog_template_int.conf.	A
0AB5	Unable to open /opt/IBM/sonas/etc/rsyslog_template_strg.conf.	A
0AB6	Unknown node roles.	A
0AB7	Error writing /etc/rsyslog.conf.	A
0ABB	Unable to gather shared SSH keys.	A
0ABC	Unable to copy new private keys.	A
0ABD	Unable to copy new public keys.	A
0ABE	Unable to copy shared keys to the remote system.	A
0ABF	Unable to copy user keys on remote system.	A
0AC0	Unable to copy host keys on remote system.	A
0AC1	Unable to open local public RSA key file.	A
0AC2	Unable to parse local host's RSA public key file.	A
0AC3	Unable to open the local host public RSA key file.	A
0AC4	Unable to send local key to the remote system.	A
0AC5	Unable to access remote system after sending local key.	A
0AC6	Unable to gather remote system's public key.	A
0AC7	Unable to gather remote system's host public key.	A
0AC8	Unable to generate public/private keys.	A
0AC9	Unable to copy user SSH keys.	A
0ACA	Unable to copy host SSH keys.	A
0ACB	Unable to copy shared keys to remote host.	A
0ACC	Unable to update keys on remote host.	A
0ACD	Unable to read in shared user key.	A
0ACE	Unable to read in shared host key.	A
0ACF	Unable to open authorized keys file for reading.	A
0AD0	Unable to open temp file for writing.	A
0AD1	Error moving temporary file.	A
0AD2	Error opening known hosts file.	A
0AD3	Error opening temporary file.	A
0AD4	No host name provided to exchange keys with.	A
0AD5	Host name is invalid.	A
0AD6	Invalid parameters.	D
0AD7	Unable to open vpdnew.txt file.	A
0AD8	VPD failed to update a value.	A
0AD9	Invalid option.	D
0ADA	Error while parsing adapter ID.	B

Table 34. Error messages and actions (continued)

Error code	Error message	Action key
0ADB	Unable to open /proc/scsi/scsi.	B
0AF8	Trying to install management stack on non-management node.	A
0AF9	Invalid site ID. Curently only 'st001' is supported on physical systems.	A
0AFA	This node is already a part of a cluster. Unable to configure.	E
0AFB	Unable to generate public/private keys.	A
0AFC	Unable to copy user SSH keys.	A
0AFD	Unable to copy host SSH keys.	A
0AFE	Unable to set the system's timezone.	A
0AFF	Unable to write clock file.	A
0B00	Unable to write to /etc/ntp.conf.	A
0B01	Unable to parse internal IP range.	D
0B08	Unable to open dhcpd.conf template file.	A
0B09	Unable to open dhcpd.conf for writing.	A
0B0A	Unable to copy dhcpd.conf to /etc/.	A
0B0B	Unable to copy tftp to /etc/xinetd.d.	A
0B0E	Unable to enable the TFTP server.	A
0B12	sonas_setup_security is not present.	A
0B13	No service IP provided.	D
0B14	Unable to create RSA1 SSH keys.	A
0B15	Unable to create RSA SSH keys.	A
0B16	Unable to create DSA SSH keys.	A
0B17	Exiting on trap.	A
0B18	No controllers found in this cluster.	A
0B2F	Unable to set GPFS setting. Check logs for more details.	A
0B30	Unable to query current GPFS settings from mmlscluster.	A
0B31	There was an error while attempting to enable CTDB.	A
0B32	Unable to query current GPFS settings mmlsconfig.	A
0B33	Unable to open settings file. Check logs for more details.	A
0B34	Invalid arguments passed to the script.	A
0B4F	add_new called with improper parameters.	A
0B50	Invalid serial number.	B
0B51	Invalid forced ID.	A
0B52	Invalid site.	A
0B53	Node with serial was not found in available list.	B
0B54	Storage nodes must be added in pairs. Invalid peer serial.	A

Table 34. Error messages and actions (continued)

Error code	Error message	Action key
0B55	Storage node peer must be a different serial.	A
0B56	Peer node is not a storage node.	A
0B57	There is already a node with ID.	A
0B58	There is a node at the peer's ID.	A
0B59	No existing cluster found. Node ID must be specified.	A
0B5A	Unable to determine management IP address of this node.	A
0B5B	Unknown node type.	B
0B5C	IP address conflict detected with the management IP. There is a node that already has this IP address.	D
0B5E	IP address conflict detected with its peer management IP. There is a node that already has this IP address.	D
0B5F	Error updating node's data in newnodes.dat.	B
0B60	Error writing temporary file.	A
0B62	Node did not finish configuration before timeout.	B
0B7F	All nodes must be up before adding a new node.	A
0B80	Unable to find the peer storage node.	Check Fibre Channel cabling between the file modules and the control enclosure. Verify that the control enclosure is up. Refer to "Powering the system on and off" in the <i>IBM Storwize V7000 Unified Information Center</i> .
0B81	The host name was not set properly.	A
0B82	Unable to create temp file nodes.lst.	A
0B85	Error copying cluster configuration to node.	A
0B86	Error restoring cluster configuration on node.	A
0B87	There was an error while adding nodes to the GPFS cluster.	A
0B88	There was an error while configuring GPFS licensing.	A
0B89	There was an error while configuring GPFS quorum.	A
0B8C	There was an error in updating the configuration on the new node.	A
0B8D	Error reading checkpoint file.	A
0B8E	Error writing to checkpoint file.	A
0B8F	There was an error while installing GPFS callbacks.	A
0B92	Rsync failed between management nodes.	C
0B94	There were too many potential peer storage nodes. Storage controllers may be cabled incorrectly or UUIDs might not be set properly.	A

Table 34. Error messages and actions (continued)

Error code	Error message	Action key
0B95	Invalid parameters.	D
0B96	Failed to configure the management processes on mgmt001st001	D then A then B
0B97	IP is invalid.	D
0B98	Netmask is invalid.	D
0B99	IP, gateway, and netmask are not a valid combination.	D
0B9A	There was an internal error.	A
0B9B	Invalid NAS private key file.	F
0B9C	Unable to copy the NAS private key file.	F
0B9D	Internal error setting permissions on NAS private key file.	A
0B9E	No NAS private key file found. Verify that the Storwize V7000 configuration ran properly.	F
0B9F	Unable to find local serial number in new nodes.	B
0BA0	Unable to find node at new IP address. Check the node cabling.	C
0BA1	Remote node is at a higher code level.	E
0BA2	Management IP for node not found.	D
0BA3	The disk IP was not found in VPD.	D
0BA4	Unable to attach to Storwize V7000 system. Private key files might not match.	G then F
0BA5	Unable to add Storwize V7000 system to CLI.	A
0BA6	The addstoragesystem CLI command has failed.	G then D
0BAC	Unable to find remote serial number in newnodes.	C then D then B
0BAD	Remote node is at a higher code level.	E
0BAE	Incorrect parameters.	A
0BAF	Unable retrieve the node serial number.	A
0BCC	Unable to configure policy routing	D then C then B
0BB0	Unable to open pxeboot data file.	A
0BB1	Unable to update pxeboot data file for node.	A
0BB2	Unable to set file permissions.	A
0BB3	Unable to find node serial in pxeboot data file.	A
0BB4	Node had an internal error during configuration.	A
0BC6	Unable to configure system.	A
0BC9	Invalid arguments passed to the script.	A
0BDA	Error copying update test utility to controller.	G then I
0BDB	Error running update test utility on controller, see Storwize V7000 for more details.	I
0BD7	Yum is reporting a package error on a node.	Try running yum manually.
01B2	Unable to start performance collection daemon.	Contact IBM Remote Technical Support.

Table 34. Error messages and actions (continued)

Error code	Error message	Action key
01B3	Failed to copy update package to Storwize V7000 system.	H then G
01B4	Failed to start update on Storwize V7000 with the svctask applysoftware command.	H then G
01B5	Storwize V7000 multipaths are unhealthy.	H then G
01B6	Storwize V7000 volumes are unhealthy as indicated using the lsdisk command.	Check Fibre Channel cabling to storage and verify storage is up.
01B7	Failed to query status of update by using the lsupdate command.	H then G
01B9	Failed to check the Storwize V7000 version	H
01B8	Failed to query status of Storwize V7000 nodes using the lsnodes command.	H
01BE	Unable to distribute update callback	Check on health of the cluster using lshealth Contact IBM Remote Technical Support.
01BF	Update callback failed	Contact your customer advocate. Update callbacks are custom steps placed on a system before the start of update. Contact IBM Remote Technical Support.
01CF	Unable to configure node	Pull both power supply cables from subject node, wait 10 seconds, plug back in, after system comes up try again.
01C4	Unable to remove callbacks	Contact IBM Remote Technical Support.
01D5	Storwize V7000 stalled.	Contact IBM Remote Technical Support.
01D6	Storwize V7000 stalled_non_redundant	H
01DA	GPFS cluster is unhealthy	Refer to “Checking the GPFS file system mount on each file module” on page 189

Table 34. Error messages and actions (continued)

Error code	Error message	Action key
01DB	Failed to stop performance center	Please attempt to stop performance center using /opt/IBM/sofs/cli/cfgperfcenter --stop . If successful restart update. If you are unable to stop performance center please contact IBM Remote Technical Support.

Problems reported by the CLI commands during software configuration

Use this information when troubleshooting problems reported by the CLI commands during software configurations.

The following table contains error messages that might be displayed when running the CLI commands during software configuration.

Table 35. CLI command problems

CLI Command	Symptom/Message	Action
mkfs	SG0002C Command exception found : Disk <arrayname> might still belong to file system <filesystemname>.	<p>This message indicates that the arrays listed in the error message appear to already be part of a file system.</p> <ol style="list-style-type: none"> 1. Check the list of array names that you specified in the mkfs command. If the mkfs command has been used to create multiple file systems, you might have used the same array name in more than one file system. If this is the case, correct the list of array names. 2. If you are certain there is no data on the system, this problem might have been caused by an error during the manufacturing cleanup process before the machine was shipped. In this case, you can work around the problem by appending the --noverify parameter to the mkfs command. Never use the --noverify parameter on a system with customer data unless directed to do so by support personnel; improper use can cause unrecoverable data loss.

Management GUI wizard failure

DNS errors can cause management GUI wizard to fail with no clear error messages.

About this task

The management GUI wizard process can fail if there are issues with the DNS information entered into the system. Entering the incorrect information is a common problem, particularly in step 5. This step requires that you fill in the following fields:

- Domain name
- Domain Admin user ID
- Domain Admin user password
- DNS servers

Entering the incorrect information may result in messages such as **domain name not found** or **wrong user or password**. However, a failure can also occur when connecting or verifying the DNS server, which is the last entry in this step. In this case, an error message does not appear, but the step fails or hangs.

One known cause of this type of failure occurs when the DNS server Address Resolution Protocol (ARP) table shows the IP address entered was previously configured in the ARP table. In this case the DNS server does not allow the connection. An unused IP address needs to be entered or the address from the ARP table needs to be removed before restarting the management GUI wizard. Exit out of the management GUI wizard and restart the wizard again. You have to key in all fields for each step again. Once all steps are completed the system runs the configuration and restarts.

GUI access issues

This topic provides assistance in isolating and resolving problems with the GUI.

About this task

This section covers GUI access issues that allow you to isolate and resolve GUI problems. This section extends beyond the GUI in the case where a file module is not responding and requires a management switchover to the other file module. Accessing the GUI is critical to isolating and resolving system problems.

1. Does the GUI launch and are there problems logging into the system?

- **Yes:** Check that the user ID being used was set up to access the GUI. Refer to “Authentication basic concepts” in the *IBM Storwize V7000 Unified Information Center*.
- **No:** Proceed to next question.

2. Does the GUI launch and are there problems logging into the system?

- **Yes:** Verify that you are using a supported browser and the browser settings are correct. Refer to “Checking your web browser settings for the management GUI” in the *IBM Storwize V7000 Unified Information Center*.
- **No:** Proceed to next question.

Note: If the GUI does not load complete these steps.

3. Are you able to initiate an ssh connection to either file node and log in to either file node?

- **Yes:**
 - a. Run the CLI command **lsnode** and determine the status of the file nodes.
 - b. If the **lsnode** reports the management service is not running, refer to “Management node role failover procedures” on page 183.
 - c. If **lsnode** provides the system configuration information, check the connection status under the appropriate heading. Is the status set to **OK**:

Note: The Sample Output shown has been adjusted in regards to spacing and layout to accommodate this publication. It might not match exactly what is seen on your system.

Sample Output:

```
[admin@kq186wx.mgmt001st001 ~]# lsnode
```

Hostname	IP	Description	Role	Product version	Connection status	GPFS status	CTDB status	Last updated
mgmt001st001	172.31.8.2	active	management, interface, node	1.3.0.0-51c	OK	active	active	9/19/11 8:02 AM
mgmt002st001	172.31.8.3	passive	management, interface, node	1.3.0.0-51c	OK	active	active	9/19/11 8:02 AM

EFSSG1000I The command completed successfully.

- **Yes:** Run the CLI command **lshealth**. Reference the active management node Hostname (mgmt001st001 or mgmt002st002) obtained from the **lsnode** command. Ensure that HOST_STATE, SERVICE, and NETWORK from **lshealth** is set to OK.

Sample Output:

mgmt001st001	HOST_STATE	OK	OK
	SERVICE	OK	All services are running OK
	CTDB	OK	CTDBSTATE_STATE_ACTIVE
	GPFS	OK	ACTIVE
	SCM	OK	SCM system running as expected
	NETWORK	ERROR	Network interfaces have a degraded state
mgmt002st001	CHECKOUT	OK	Disk Subsystem have a online state
	HOST_STATE	OK	OK
	SERVICE	OK	All services are running OK
	CTDB	OK	CTDBSTATE_STATE_ACTIVE
	GPFS	OK	ACTIVE
	SCM	OK	SCM system running as expected
V7000	MGMTNODE_REPL_STATE	OK	OK
	NETWORK	ERROR	Network interfaces have a degraded state
	CLUSTER	ERROR	Alert found in component cluster
	ENCLOSURE	ERROR	Alert found in component enclosure
	IO_GRP	OK	The component io_grp is running OK
	MDISK	OK	The component mdisk is running OK
	NODE	OK	The component node is running OK
	PORT	ERROR	Alert found in component port

EFSSG1000I The command completed successfully.

- **No:** Perform network connectivity isolation procedures. Refer to “Management node role failover procedures” on page 183.

No: Perform network connectivity isolation procedures. Refer to “Management node role failover procedures” on page 183.

If none of the previous steps resolved the GUI connectivity issues, perform the following procedure.

Network port isolation for GUI:

If none of the previous steps have resolved the problem and the network connectivity and system reports nothing wrong, there might be an issue with the port configuration of your network that is not detected in any of the previous steps. The internal management services use both port 443 and port 1081. Port 443 is redirected to port 1081 that the management service listens.

1. Check to see if you can access the GUI on the default https port (no port included in the URL). If all is good with firewall and management IP, the GUI will listen on **https://<Management IP>/** and provide a login prompt.
2. Check network port settings and firewall settings. If the previous step fails, investigate the following issues:
 - The firewall is open between the administrative browser and the Primary Node Service IP but not between the administrative browser and the management IP. The firewall settings must have rules that allow port 1081 but not 443 between the administrative browser and the management IP.
 - The management IP is up but the port redirection on the switch/router is not working as expected. Check the network settings.

Health status and recovery

Use this information to review the outstanding issues that cause the **Health Status** indicator at the bottom of all management GUI panels to be red (critical errors) or yellow (warnings or degradations).

Before you begin

Use this procedure after you resolve the events from the **Monitoring > Events** page to resolve the overall system health status indicators. You can also mark events as read from the events page to display the updated health status in the health status indicator that is placed at the bottom of all panels.

About this task

Within the Storwize V7000 Unified system, the system **Health Status** is based on a set of predefined software and hardware health status sensors. The status of each component is displayed against the corresponding logical host name in the System and System Details pages.

For storage problems, resolve events and health status by running the recommended repair action or actions from the **Block** tab of **Monitoring > Events**.

For file module problems, the software and hardware sensors are different. Some of the sensors are automatic and actively reflect the current status of the system; whereas, some of the sensors, such as the hardware sensors, require a reset after the service actions are completed.

Note: For the file modules, the System Details page and sensors are separate from the events. Events that are displayed in the log might be reflected within a corresponding sensor with the **System Details > Status** indicator for the failing host name. However, be aware that resolving events and resetting the corresponding sensor changes the health status of the system but does not clear the corresponding event from the event log.

This topic instructs you where to go to view the information that is displayed, how to check the status of the various sensors, and how to manually close out sensor events. By performing these tasks, you ensure that the overall **Health Status** reflects the current system health.

To resolve the overall health status indicators, perform the following steps:

Procedure

1. Log on to the management GUI.
2. Navigate to **Monitoring > System Details**.
3. Expand any mgmt00xst001 subcomponent that shows a critical or warning event indication and select **Status**.

- a. Review the **Sensor** column and the **Level** column for **Critical Error**, **Major Warning**, or **Minor Warning** items.

If the problem that caused the **Level** item is resolved, right-click the event and select the **Mark as Read** action.

- b. Click **Yes** in the information dialog to complete the action.
- c. Review the status list for the other events that might cause the **Health Status** to be red or yellow.
- d. Perform the same steps.

As long as there is a single sensor that is marked as **Critical Error**, **Major Warning**, or **Minor Warning**, the **Health Status** is red or yellow. When you use the **Mark as Read** action against the sensor, the sensor no longer shows in the status view. If the problem is still not resolved, a new sensor update occurs that reflects the problem. An example might be if a software error event is marked as read but the system still detects the problem; then the status is properly reflected in the **Status** display.

Connectivity issues for the 2073-720

This topic provides information for troubleshooting connectivity issues. The major focus is on connectivity between the file modules and the control enclosure. Good connectivity is required to troubleshoot control enclosure problems.

Host to file modules connectivity

This procedure is used to troubleshoot Ethernet network connectivity between the host and the file modules. These network paths are used for all system requests and management operations. The paths are also needed for Ethernet network connectivity between the file module and the Storwize V7000.

About this task

Within the file modules, two internal 1 GB network ports and two 10 GbE network ports can be configured for system operations.

Figure 39 on page 64 identifies the various rear ports and hardware for the file module.

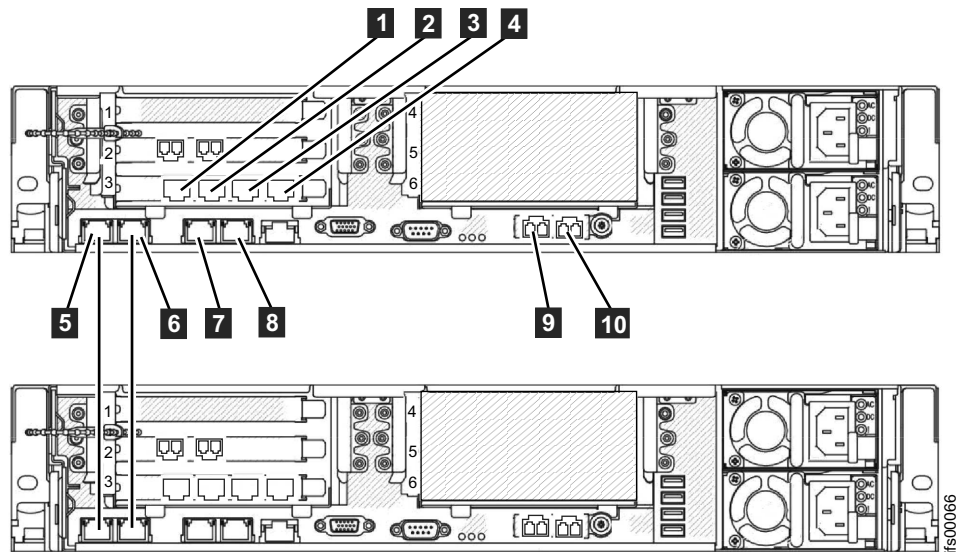


Figure 39. File module Ethernet connections.

Table 36. Ethernet connections available with the file modules

Item	Port	IP address is assigned by InitTool	Use
1	Ethernet port 7		Connect to a switch for public file access
2	Ethernet port 8		Connect to a switch for public file access
3	Ethernet port 9		Connect to a switch for public file access
4	Ethernet port 10		Connect to a switch for public file access
5	Ethernet port 1	From the internal IP address range	Connect to the other file module
6	Ethernet port 2	From the internal IP address range	Connect to the other file module
7	Ethernet port 3	File module service and system management IP address	Connect to a switch for public file access and system management
8	Ethernet port 4		Connect to a switch for public file access
9	Ethernet port 5 (10 Gbps optical)		Connect to a switch for public file access and optional system management
10	Ethernet port 6 (10 Gbps optical)		Connect to a switch for public file access

If you are looking at a problem regarding built-in Ethernet port 1 or built-in Ethernet port 2, refer to “Ethernet connectivity between file modules” on page 65.

Isolation procedures:

Ensure that the file module is powered up before you begin this procedure. The network connection being diagnosed must be connected to an active port on your Ethernet network.

- Determine the state of the Ethernet LEDs examining the LEDs of the Ethernet ports.
- The activity LED flashes when there is activity on the connection. The link state LED must be permanently on. If it is off, the link is not connected.

If your link is not connected, perform the following actions to check the port status each time until it is corrected or connected:

1. Verify that each end of the cable is securely connected.
2. Verify that the port on the Ethernet switch or hub is configured correctly. Contact your network administrator to verify the switch and network configuration information.
3. Connect the cable to a different port on your Ethernet network.
4. Replace the Ethernet cable.
5. For the 10 GbE Ethernet port, replace the small form-factor pluggable (SFP) transceiver. Refer to “Removing a PCI adapter from a PCI riser-card assembly” on page 143 and “Installing a PCI adapter in a PCI riser-card assembly” on page 145.

Ethernet connectivity between file modules

This topic covers troubleshooting Ethernet connectivity issues between the file modules. These connections are used for internal management operations between the file modules. They make use of the Internal IP address range that you provided during initializing the Storwize V7000 Unified system.

About this task

This procedure is used to troubleshoot Ethernet connectivity between the file modules. These network paths are used for all internal file system communication. Between the file module, there are two separate network paths for internal communication.

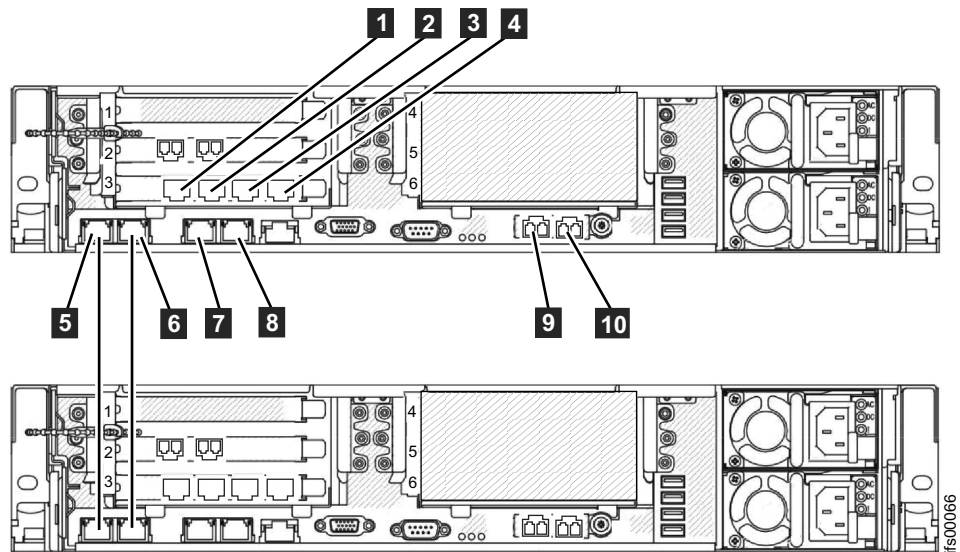


Figure 40. File module Ethernet connections.

Table 37. Ethernet connections available with the file modules

Item	Port	IP address is assigned by InitTool	Use
1	Ethernet port 7		Connect to a switch for public file access
2	Ethernet port 8		Connect to a switch for public file access
3	Ethernet port 9		Connect to a switch for public file access
4	Ethernet port 10		Connect to a switch for public file access
5	Ethernet port 1	From the internal IP address range	Connect to the other file module
6	Ethernet port 2	From the internal IP address range	Connect to the other file module
7	Ethernet port 3	File module service and system management IP address	Connect to a switch for public file access and system management
8	Ethernet port 4		Connect to a switch for public file access
9	Ethernet port 5 (10 Gbps optical)		Connect to a switch for public file access and optional system management
10	Ethernet port 6 (10 Gbps optical)		Connect to a switch for public file access

If you are looking at a problem regarding built-in Ethernet port 3, built-in Ethernet port 4, or any network connections to PCI slot 4, refer to “Host to file modules connectivity” on page 63.

Isolation procedures:

Ensure that both of the file module are powered up before you begin this procedure:

- Determine the state of the Ethernet LEDs by examining the Ethernet port LEDs.
- The activity LED flashes when there is activity on the connection. The link state LED must be permanently on. If it is off, the link is not connected.

If your link is not connected perform the following actions to check the port status until it is corrected or connected:

1. Verify that each end of the cable is securely connected.
2. Replace the Ethernet cable.
3. Replace the failed Ethernet port on the server by replacing the system planer. Refer to “Removing the system board” on page 177 and “Installing the system board” on page 179.

Duplicate IP address procedure:

If you are experiencing odd intermittent problems with communications between the file modules then it could be that some other machine on your network is using the same IP address as one of the four IP addresses used for the file modules to communicate with each other. These IP addresses were set during initial setup from the internal IP address range that you chose in the initialization tool.

It is always possible that somebody in your site could set up another machine to use one or more IP address that your Storwize V7000 Unified system is already using. Use the management GUI to check which four IP addresses the file modules are currently using to communicate with each other. See the device = mgmt0 box, in **Monitoring > System Details > Network** panel that is available under each file module interface node name.

Follow this procedure:

1. Find the system IP address of the control enclosure in the **Settings > Network > IP Report**. Log on to the storage system CLI. For example: (default password is passw0rd):

```
ssh superuser@<system IP address>
```
2. Use ping from the storage system CLI to see if any packets are returned from each of the internal IP addresses used for file modules to communicate with each other. For example:

```
IBM_2076:mssystem:superuser>ping 10.254.8.1
PING 10.254.8.1 (10.254.8.1) 56(84) bytes of data.
--- 10.254.8.1 ping statistics ---
5 packets transmitted, 0 received, 100% packet loss, time 4005ms
```
3. You should get 100% packet loss. If you do not get 100% packet loss then some other machine on your network is using this IP address.

If you cannot stop other machines on your network using these IP addresses and must change the internal IP address range used then you need to contact IBM Remote Technical Support to help you to put your file modules back to an out-of-box state so you can choose a different internal IP address range. All other IP addresses used by the system can be changed without needing to put the file modules back to an out-of-box state.

Ethernet connectivity from file modules to the control enclosure

This topic covers troubleshooting Ethernet network connectivity issues between the file modules and the attached control enclosure. These network paths are used for all management operations between the file module and the control enclosure.

About this task

This procedure is used to troubleshoot Ethernet network connectivity between the file modules and the control enclosure. These connections are used for the active management node on one of the file modules to ssh storage command-line interface (CLI) commands to the main configuration node canister in a control enclosure.

There are no direct physical Ethernet connections between the file module hardware and the control enclosure. All network connections are done through your network infrastructure. When configuring your network switches, be sure that there is an available communication path between the file module network connections and the control enclosure network connections. Ideally the file modules and control enclosure should be connected to the same 1 Gbps Ethernet switch

If you want redundant connectivity to the control enclosure from the file modules, then both 1 Gbps ports from each node canister in the control enclosure are connected to your network. If you do not want redundancy, connecting port 2 of the control enclosure node canister to your network is optional.

If you seem to have intermittent management communication problems between the file module which is the active management node and the control enclosure CLI, then it is possible that another machine on your network could be using the IP address used by the control enclosure. Refer to Problem: Another system may be using the system IP address for how to check for a duplicate IP address on your network and how to change the control enclosure IP address if necessary.

If the file modules can no longer ssh CLI commands to the storage system CLI then the first thing to do is make sure that the management IP addresses are correctly set. You may find the GUI works very slowly in this case, so access the CLI by using ssh to the log on to the management IP address as admin (default password **admin001**).

For example:

```
ssh admin@<mangementIP>
```

Use the **lsnwmgt** CLI command to show you the IP addresses used by the file modules for management. For example:

```
[kd52y0g.ibm]$ lsnwmgt
Interface Service IP Node1 Service IP Node2 Management IP Network Gateway VLAN ID
ethX0 9.71.18.160 9.71.18.161 9.71.18.210 255.255.255.0 9.71.18.1
EFSSG1000I The command completed successfully.
```

Use the **lsstoragesystem** CLI command to show you the IP address that the active management node, running on one of the file modules, will use to ssh commands to the storage system CLI. For example:

```
[kd52y0g.ibm]$ lsstoragesystem
name          primaryIP    secondaryIP id
StorwizeV7000 9.71.18.180 9.71.18.180 00000200A6002C08
EFSSG1000I The command completed successfully.
```

Check that these 5 or 6 IP addresses and sub net mask are as expected. Attempt the **lsystemip** CLI command which will probably fail when the active management node running on a file module attempts to ssh it to the storage system CLI running on a control enclosure. For example:

```
[kd52y6h.ibm]$ lssystemip
EFSSG0655C Error in communication with the storage system.
Failed to open SSH connection
```

However, if this CLI command now works then the original problem with ssh to the storage system CLI may have gone away. Otherwise, use ping to check the network connections between the storage system and the file modules. It does not work from the management CLI but it should work from the storage system CLI.

From an external computer ssh as superuser to the primary IP given by the **lsstoragesystem** CLI. The IP that the active management node will be trying to ssh commands to the storage system CLI. For example (default password is passw0rd):
ssh superuser@9.71.18.180

If you can not ssh to the storage system primary IP or secondary IP (that was given by the **lsstoragesystem** CLI command) then follow the procedure to use the USB flash drive to discover the state and settings of the Storwize V7000. Make sure that there is no satask.txt file on the USB flash drive before you plug it into the control enclosure.

Plug the USB flash drive into the control enclosure. The orange fault light should go on for a short time only. (such as a slow blink for a few seconds) . Wait for the orange fault light to go out then unplug the USB flash drive and plug it into another computer so that you can look at the contents of the satask_results.html file on the USB flash drive. The satask_results.html will contain the output from a number of sainfo commands.

Check the following;

- The cluster_id under sainfo lsservicestatus should match the id (that was given by the **lsstoragesystem** CLI command). Otherwise you may have plugged the USB flash drive into the wrong control enclosure (such as one that is not part of this Storwize V7000 unified system). The node_status should be active for each node canister in the cluster under sainfo lsservicestatus. Otherwise follow the service action under sainfo lsservice recommendation.
- The cluster_ip under sainfo lsservicestatus should match the Primary IP (that was given by the **lsstoragesystem** CLI command). Otherwise investigate which of the IP addresses is the correct one and make the other one match it. Refer to the instructions later on this page if you need to change the storage system IP address but can not log onto the storage system IP address to use the CLI.

This is an example of some of what the satask_results.html would contain on a healthy storage system for you to compare to your results:

```
Thu Apr 19 08:23:42 UTC 2012
satask.txt file not found.
System Status
sainfo lsservicenodes
panel_name cluster_id    cluster_name    node_id node_name relation node_status
error_data
01-1      00000200A4E008BA Cluster_9.71.18.184 1      node1      local      Active
```

```

01-2      00000200A4E008BA Cluster_9.71.18.184 2      node2      partner  Active
sainfo lsservicestatus
panel_name 01-1
cluster_id 00000200a4e008ba
cluster_name Cluster_9.71.18.184
cluster_status Active
cluster_ip_count 2
cluster_port 1
cluster_ip 9.71.18.184
cluster_gw 9.71.18.1
cluster_mask 255.255.255.0

```

When you can ssh to the storage system IP then use the **lssystem** CLI command on the storage system CLI to show you what it thinks that its system IP address is:

```

IBM_2076:tbcluster-ifs4:superuser>lssystemip
cluster_id  cluster_name  location port_id IP_address  subnet_mask  gateway
IP_address_6 prefix_6 gateway_6
00000200A6402C08 tbcluster-ifs4 local 1 9.71.18.180 255.255.255.0 9.71.18.1
00000200A6402C08 tbcluster-ifs4 local 2

```

Check that these IP addresses and sub net mask are as expected. Use the **chsystemip** CLI if you must change anything. Use ping to check the path back to the management IP address.

```

IBM_2076:tbcluster-ifs4:superuser>ping 9.71.18.160
PING 9.71.18.160 (9.71.18.160) 56(84) bytes of data.
64 bytes from 9.71.18.160: icmp_seq=1 ttl=64 time=0.103 ms
64 bytes from 9.71.18.160: icmp_seq=2 ttl=64 time=0.096 ms
64 bytes from 9.71.18.160: icmp_seq=3 ttl=64 time=0.082 ms
64 bytes from 9.71.18.160: icmp_seq=4 ttl=64 time=0.081 ms
64 bytes from 9.71.18.160: icmp_seq=5 ttl=64 time=0.082 ms

--- 9.71.18.160 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4001ms
rtt min/avg/max/mdev = 0.081/0.088/0.103/0.014 ms

```

If the ping from the storage system CLI back to the management IP has 100% packet loss then investigate the physical 1 Gbps Ethernet cabling and the configuration of the Ethernet switch. Also check the Ethernet port LEDs on:

- Built in Ethernet port 3 of each file module
- Ethernet port 1 on each node canister of the control enclosure

If the ping from the Storwize V7000 to each file module has 0% packet loss, then the ssh key should be reset. Follow the “Resetting the NAS ssh key for configuration communications” on page 402 procedure in the Information Center to reset the NAS key.

If you need to change the IP settings on the storage system but cannot use ssh to access the current system IP to run the **chsystemip** CLI command then refer to Problem: Unable to change the system IP address because you cannot access the CLI.

If you plan to change the system IP address and can ssh to the current system IP address, then you can run the **chsystemip** CLI command. Here is an example:

```

>ssh superuser@<system IP address>
$ chsystemip -clusterip 9.20.136.5 -gw 9.20.136.1 -mask 255.255.255.0 -port 1

```

The default password for superuser is **passw0rd**.

Update the file module's record of the control enclosure system IP:

To find the file module's current record of the control enclosure system IP address, use the Storwize V7000 Unified management CLI to issue the **lsstoragesystem** command. Here is an example:

```
>ssh admin@<management_IP>
[kd01ghf.ibm]$ lssstoragesystem
name           primaryIP      secondaryIP    id
StorwizeV7000  9.11.137.130  9.11.137.130  00000200A2601508
EFSSG1000I The command completed successfully.
```

If the primary and secondary IP address shown by the **lsstoragesystem** CLI do not match the system IP addresses shown in the output of the **lssystemip** CLI command, then it is necessary to update the record. The **chstoragesystem** command changes the file module record of the control enclosure system IP. Here is an example:

```
>[kd01ghf.ibm]$ chstoragesystem --ip1 9.71.18.136 --ip2 9.71.18.136
EFSSG1000I The command completed successfully.
```

Verify that communication from the file module to the control enclosure is now possible by running the **lssystemip** command on the Storwize V7000 Unified management CLI:

```
>ssh admin@<management IP address>
[kd01ghf.ibm]$ lssystemip
```

Changing the cluster IP of the file modules:

If the cluster IP address of the file modules is not known, or has been incorrectly set, the value can be changed by logging into the system using a console.

Connect a keyboard and monitor directly into the front of the file module which is the active management node. Login as a user with administrative access rights:

- Login: admin
- Password: <default is admin>

View the current cluster IP setting using the **lsnwmgt** command:

```
>$ lsnwmgt
[kd271f5.ibm]$ lsnwmgt
Interface Service IP Node1 Service IP Node2 Management IP Network Gateway VLAN ID
ethX0 9.115.160.221 9.115.160.222 9.115.160.220 255.255.248.0 9.115.167.254
EFSSG1000I The command completed successfully
```

You may receive the following error:

```
$ lsnwmgt
EFSSG0026I Cannot execute commands because Management Service is stopped.
Use startmgtsrv to restart the service.
```

This is an indication that the node you are currently connected to is not the active management node. Plug the keyboard and monitor into the other node, login again and retry the **lsnwmgt** command.

To change the file module cluster IP to its new value, use the **chnwmgt** command:

Here is an example:

```
>$ chnwmgt -mgtip 9.115.160.210 -- netmask 255.255.255.0 -gateway 9.115.160.254
```

Checking the physical status of the Ethernet ports:

The following procedures require physical access to the system. If your link is not connected, perform the following actions to check the port status each time until it is corrected or connected.

- Examine the Ethernet ports LEDs. The activity LED flashes when there is activity on the connection. The link state LED must be permanently on. If it is off, the link is not connected.
- Verify that each end of the cables is securely connected.
- Verify that the port on the Ethernet switch or hub is configured correctly.
- Connect the cable to a different port on your Ethernet network.
- If the status is obtained using the USB flash drive, review all the node errors that are reported.
- Replace the Ethernet cable.
- Follow the hardware replacement procedures for a node canister.
- Follow the hardware replacement procedures for a file module.

If you are unable to change the service IP address, for example, because you cannot use a USB flash drive in the environment, see “Procedure: Accessing a Storwize V7000 Gen1 canister using a directly attached Ethernet cable” on page 276.

Fibre Channel connectivity between file modules and control enclosure

This procedure is used to troubleshoot Fibre Channel connectivity issues between the file modules and the Storwize V7000 control enclosure. The Fibre Channel paths are the paths used for transferring data between the file module and the Storwize V7000 control enclosure.

Before you begin

Before beginning this troubleshooting procedure, review the events listed under the **Block** tab. Perform any recovery actions for events that are listed there.

About this task

Each file module has a dual port Fibre Channel adapter card located in PCI slot 2. Both ports are used to connect to the Storwize V7000 control enclosure with a connection going to each control canister as shown in Figure 41 on page 73 or Figure 42 on page 74.

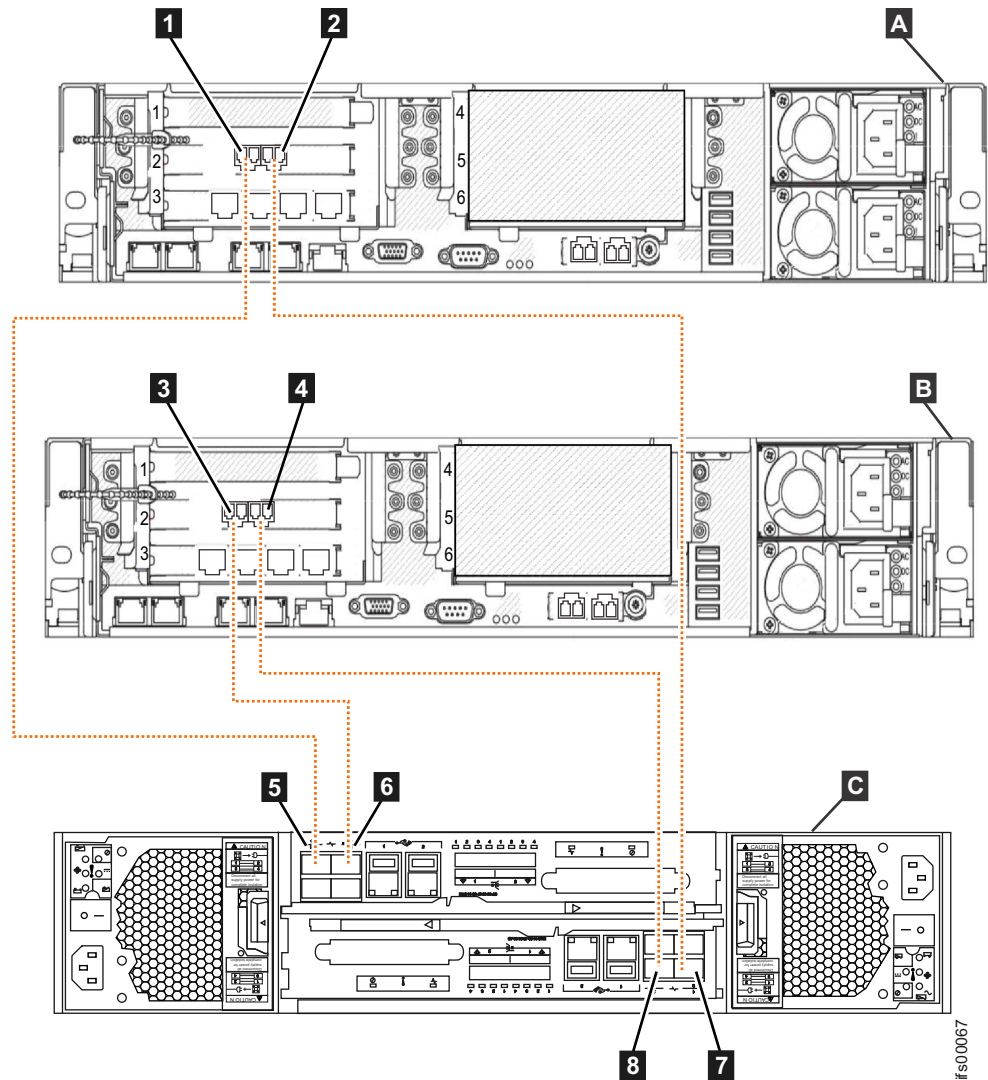


Figure 41. Connecting the file modules to the Storwize V7000 Gen1 control enclosure using Fibre Channel cables

- **A** File module 1
- **B** File module 2
- **C** Storwize V7000 control enclosure
- **1** File module1 - Fibre Channel port 1
- **2** File module 1 - Fibre Channel port 2
- **3** File module 2 - Fibre Channel port 1
- **4** File module 2 - Fibre Channel port 2
- **5** Upper node canister - Fibre Channel port 1
- **6** Upper node canister - Fibre Channel port 2
- **7** Lower node canister - Fibre Channel port 1
- **8** Lower node canister - Fibre Channel port 2

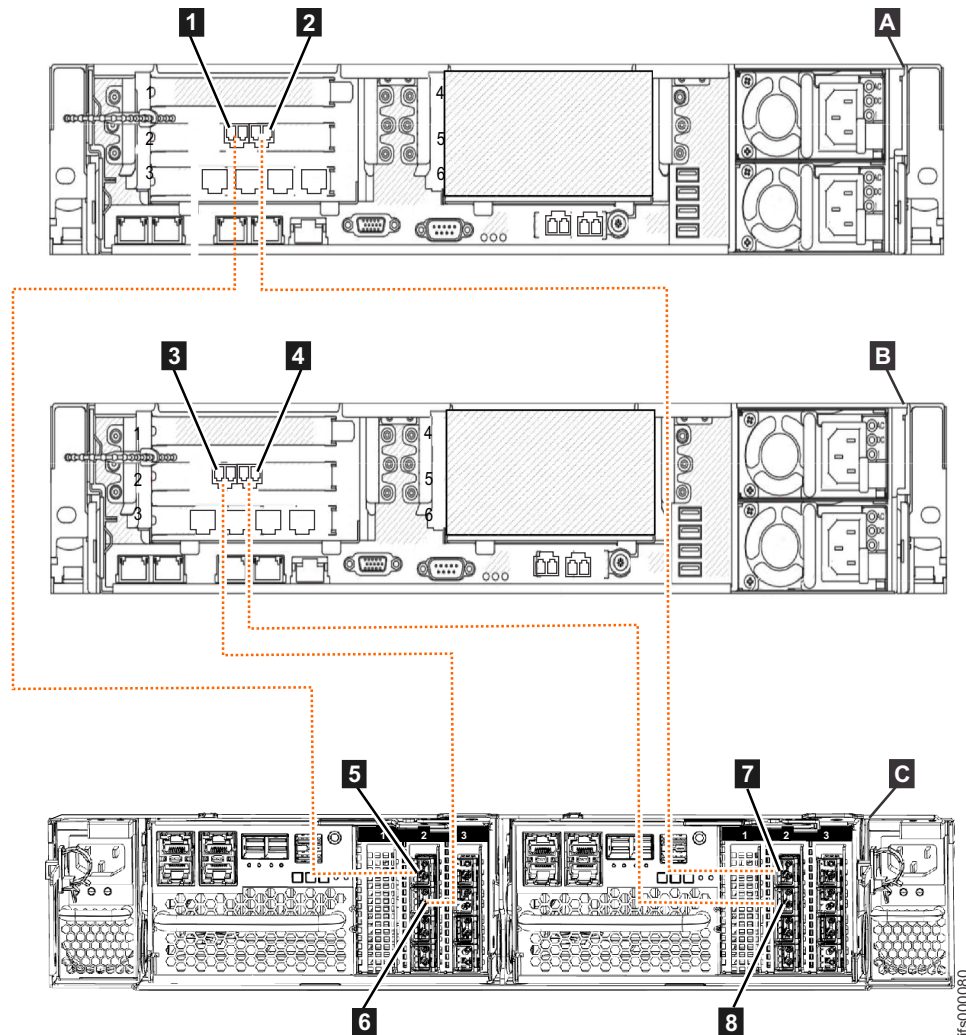


Figure 42. Connecting the file modules to a Storwize V7000 Gen2 control enclosure that has a Fibre Channel interface adapter in PCI slot 2 of each node canister

- **A** File module 1
- **B** File module 2
- **C** Storwize V7000 Gen2 control enclosure (2076-524)
- **1** File module 1 - Fibre Channel port 1
- **2** File module 1 - Fibre Channel port 2
- **3** File module 2 - Fibre Channel port 1
- **4** File module 2 - Fibre Channel port 2
- **5** Node canister 1 (left) - Fibre Channel port 1
- **6** Node canister 1 (left) - Fibre Channel port 2
- **7** Node canister 2 (right) - Fibre Channel port 1
- **8** Node canister 2 (right) - Fibre Channel port 2

Table 38 on page 75 describes the diagrams shown in Figure 41 on page 73 and Figure 42.

Table 38. How to connect Fibre Channel cables from file modules to the control enclosure.

File module	Control enclosure
A File module 1	C Control enclosure
1 Fibre Channel slot 2, port 1	5 Node canister 1 Fibre Channel port 1
2 Fibre Channel slot 2, port 2	7 Node canister 2 Fibre Channel port 1
B File module 2	C Control enclosure
3 Fibre Channel slot 2, port 1	6 Node canister 1 Fibre Channel port 2
4 Fibre Channel slot 2, port 2	8 Node canister 2 Fibre Channel port 2

The Storwize V7000 control enclosure contains an upper and lower (inverted) canister.

In isolating problems, be sure to review the labels on the rear of the systems for exact port plugging.

Software detected problems via event codes:

If a software event code directed you to this procedure, use the **Monitoring > System** page in the management GUI to identify the effected file module or refer to the following procedure to determine the logical to physical mapping of the event, then proceed to the physical hardware isolation procedures.

The isolation of Fibre Channel connections based on a single error event is not simple. As Figure 41 on page 73 shows, there are two file modules attached to the control enclosure; however, the logical host name of these systems does not map directly to the connections. The logical host name of the file module depends on which file module is used for initial USB flash drive installation. For example, in Figure 41 on page 73, file module **B** can have a host name of **mgmt001st001** if the installation was initiated on that node or it might have a host name of **mgmt002st001** if the installation was initiated on the second file module. Each error event is reported against the logical host name where the problem occurred.

For isolation of Fibre Channel connections, it is important with a single file module that both Fibre Channel connections go to the same port number on both Storwize V7000 node canisters. Port 1 always goes to node canister 1, and port 2 goes to node canister 2.

Use the following table for correlating the error code with the physical connections, then follow the procedures after the table for enabling the LED indicator on the front of the file module.

Table 39. Error code port location mapping

Error code	Description	File Module Fibre Channel Location	Storage Node Canister Fibre Channel Port
4B0800C	Link failure. Fibre Channel adapter 1, port 1 not up.	PCI slot #2 – port 1 (right port when facing rear of system)	Node canister 1, port 1. OR node canister 1, port 2.
4B0801C	Link failure. Fibre Channel adapter 1, port 2 not up.	PCI slot #2 – port 2 (left port when facing rear of system)	Node canister 2, port 1. OR node canister 2, port 2.

Table 39. Error code port location mapping (continued)

Error code	Description	File Module Fibre Channel Location	Storage Node Canister Fibre Channel Port
4B0803C	Slow connection on Fibre Channel adapter 1, port 1.	PCI slot #2 – port 1 (right port when facing rear of system)	Node canister 1, port 1. OR node canister 1, port 2.
4B0804C	Slow connection on Fibre Channel adapter 1, port 2.	PCI slot #2 – port 2 (left port when facing rear of system)	Node canister 2, port 1. OR node canister 2, port 2.

To enable the LED indicator for the node reporting the problem, use the **Monitoring > System** page in the management GUI or follow this procedure:

1. Log onto the active file module via the CLI interface.
2. Run the command: **locatenode #HOSTNAME on #SECONDS**. **HOSTNAME** is the hostname associated with the error... either **mgmt001st001** or **mgmt002st001**. **#SECONDS** is the number of seconds for the LED indicator to be turned on.

Physical connection and repair:

Each file module has a dual port Fibre Channel adapter card located in PCI slot 2. Both ports are used to connect to the Storwize V7000 system with a connection going to each Storwize V7000 node canister.

Table 40. Fibre Channel cabling from the file module to the control enclosure.

File Module Node # 1		File Module Storage Node # 2	
PCI slot #2, port 1	PCI slot #2, port 2	PCI slot #2, port 1	PCI slot #2, port 2
Connects to Storwize V7000	Connects to Storwize V7000	Connects to Storwize V7000	Connects to Storwize V7000
Node canister 2 – Fibre Channel port 1	Node canister 1 – Fibre Channel port 1	Node canister 2 – Fibre Channel port 2	Node canister 1 – Fibre Channel port 2

If a problem is detected with a Fibre Channel path between the storage node and the control enclosure, check the LED indicators next to the Fibre Channel connection ports on both the file module and the Storwize V7000 node canister.

Table 41. LED states and associated actions. For the Fibre Channel adapters on the file module check the amber LED lights next to the port.

LED State	Definition and Action
Solid amber LED	This state indicates a good connection status.
Slow flashing amber LED	This state indicates a good connection at the Fibre Channel port but a broken connection at the Storwize V7000 node canister. This broken connection is most likely either a Fibre Channel cable or the Fibre Channel port is bad on the Storwize V7000 node canister.

Table 41. LED states and associated actions. For the Fibre Channel adapters on the file module check the amber LED lights next to the port. (continued)

LED State	Definition and Action
Rapid flashing amber LED	This state indicates the Fibre Channel adapter is attempting to resync the Fibre Channel connection. This situation is normally seen after a Fibre Channel connection is unplugged and then plugged back in.
No LED	There is no connection detected at all at the file module Fibre Channel port. This broken connection is most likely caused by a Fibre Channel cable or the Fibre Channel adapter has failed.

Table 42. Fibre Channel connection on the node canister LED state and associated actions

LED State	Definition and Action
Solid green LED	This state indicates a good connection status.
No LED	There is no connection detected at all at the node canister Fibre Channel port.

The recommended repair actions for Fibre Channel connections are as follows:

1. Reseat the Fibre Channel cable at both the Fibre Channel connection and the node canister.
2. Replace the Fibre Channel cable.
3. Replace the Fibre Channel adapter in the file module. Refer to “Removing a PCI adapter from a PCI riser-card assembly” on page 143 and “Installing a PCI adapter in a PCI riser-card assembly” on page 145
4. Replace the Storwize V7000 node canister. Refer to “Replacing a Storwize V7000 Gen1 node canister” on page 297.

Understanding LED hardware indicators

This topic provides information for understanding the LED status of all system components. If you do not have an LED issue or direct access to the system, proceed to the next troubleshooting topic.

File node hardware indicators for 2073-720

Use this information to evaluate the system LEDs, which can often identify the source of an error.

Light path diagnostics is a system of LEDs on various external and internal components of the server. When an error occurs, LEDs are lit throughout the server. By viewing the LEDs in a particular order, you can often identify the source of the error.

When LEDs are lit to indicate an error, they remain lit when the server is turned off, provided that the server is still connected to power and the power supply is operating correctly.

Before you work inside the server to view light path diagnostics LEDs, read the safety information.

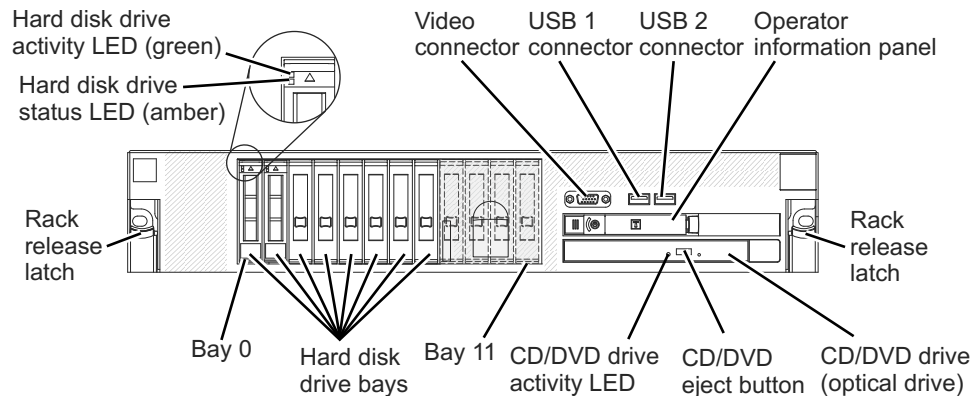
If an error occurs, view the light path diagnostics LEDs in the following order:

1. Look at the operator information panel on the front of the server.

If the information LED is lit, it indicates that information about a suboptimal condition in the server is available in the IMM event log or in the system event log.

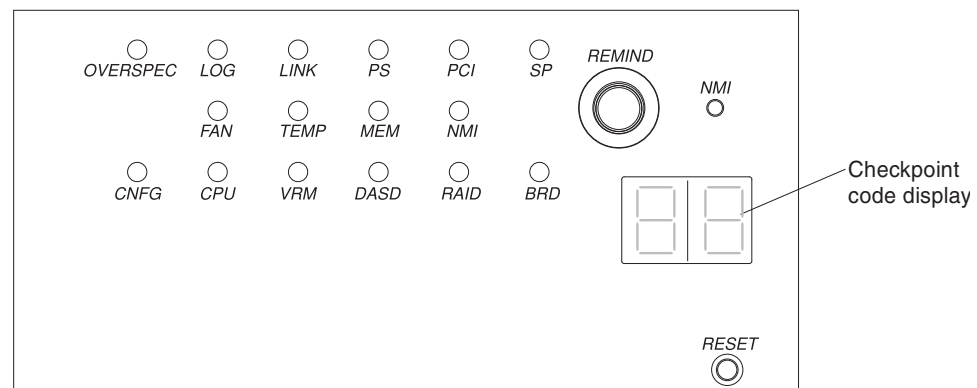
If the system-error LED is lit, it indicates that an error has occurred; go to step 2.

The following illustration shows the operator information panel on the front of the file node.



2. To view the light path diagnostics panel, slide the latch to the left on the front of the operator information panel and pull the panel forward. This reveals the light path diagnostics panel. Lit LEDs on this panel indicate the type of error that has occurred.

The following illustration shows the light path diagnostics panel.



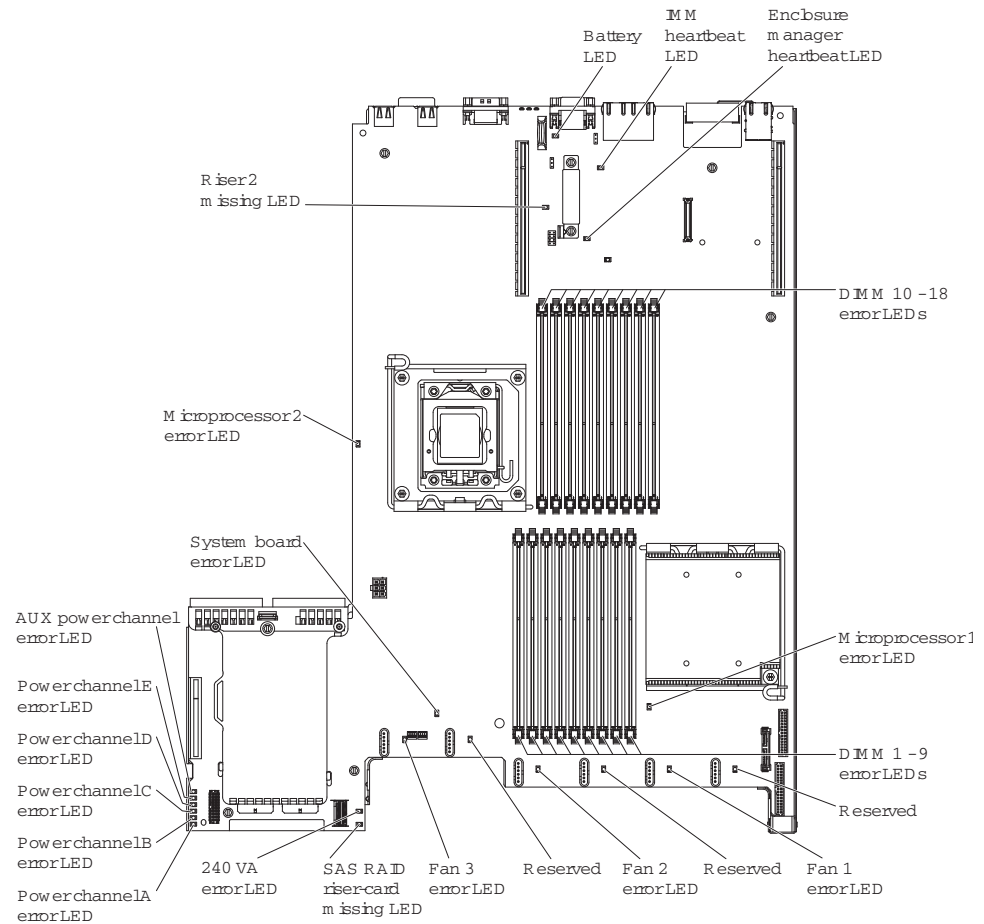
Note any LEDs that are lit, and then push the light path diagnostics panel back into the server.

Note:

- Do not run the server for an extended period of time while the light path diagnostics panel is pulled out of the server.
- Light path diagnostics LEDs remain lit only while the server is connected to power.

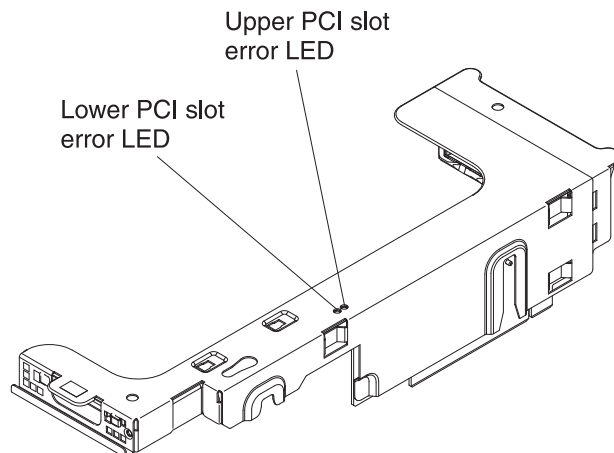
Look at the system service label on the top of the server, which gives an overview of internal components that correspond to the LEDs on the light path diagnostics panel. This information and the information in Light path diagnostics LEDs can often provide enough information to diagnose the error.

- Remove the server cover and look inside the server for lit LEDs. A lit LED on or beside a component identifies the component that is causing the error.
The following illustration shows the LEDs on the system board.



12v channel error LEDs indicate an overcurrent condition. Refer to the procedure "Solving power problems" in the "Troubleshooting the System x3650" in the *IBM Storwize V7000 Unified Information Center* to identify the components that are associated with each power channel, and the order in which to troubleshoot the components.

The following illustration shows the LEDs on the riser card.



4. Check the Light path diagnostics LEDs for the correct combination of power LEDs that should be displayed during typical operation.

Light path diagnostics LEDs

LEDs on the light path diagnostics panel of the file module indicate the cause of a problem. The topic describes the suggested actions to correct the detected problems.

Table 43. LED indicators, corresponding problem causes, and corrective actions

<ul style="list-style-type: none"> Follow the suggested actions in the order in which they are listed in the Action column until the problem is solved. If an action step is preceded by "(Trained technician only)," that step must be completed only by a trained technician. 		
LED	Problem	Action
Check log LED	An error occurred and cannot be isolated without completing certain procedures.	<ol style="list-style-type: none"> Check the IMM2 system event log and the system-error log for information about the error. Save the log if necessary and clear the log afterward.
System-error LED	An error occurred.	<ol style="list-style-type: none"> Check the light path diagnostics LEDs and follow the instructions. Check the IMM2 system event log and the system-error log for information about the error. Save the log if necessary and clear the log afterward.
PS	When only the PS LED is lit, a power supply has failed.	<p>The system might detect a power supply error. Complete the following steps to correct the problem:</p> <ol style="list-style-type: none"> Check the power-supply with a lit yellow LED. Make sure that the power supplies are seated correctly and plugged in a good AC outlet. Remove one of the power supplies to isolate the failed power supply. Make sure that both power supplies installed in the file module are of the same AC input voltage. Replace the failed power supply.
	PS + CONFIG When both the PS and CONFIG LEDs are lit, the power supply configuration is invalid.	<p>If the PS LED and the CONFIG LED are lit, the system issues an invalid power configuration error. Make sure that both power supplies installed in the file module are of the same rating or wattage.</p>
OVER SPEC	The system consumption reaches the power supply over current protection point or the power supplies are damaged.	<ol style="list-style-type: none"> If the Power Rail (A, B, C, D, E, F, G, and H) error was not detected, complete the following steps: <ol style="list-style-type: none"> Use the IBM Power Configurator utility to determine current system power consumption. For more information and to download the utility, go to http://www-03.ibm.com/systems/bladecenter/resources/powerconfig.html. Replace the failed power supply. If the Power Rail (A, B, C, D, E, F, G, and H) error was also detected, follow actions in the "Power problems" under the Troubleshooting tables and "Solving power problems" in the <i>Problem Determination and Service Guide</i>.

Table 43. LED indicators, corresponding problem causes, and corrective actions (continued)

<ul style="list-style-type: none"> Follow the suggested actions in the order in which they are listed in the Action column until the problem is solved. If an action step is preceded by "(Trained technician only)," that step must be completed only by a trained technician. 		
LED	Problem	Action
PCI	An error occurred on a PCI card, a PCI bus, or on the system board. An extra LED is lit next to a failing PCI slot.	<ol style="list-style-type: none"> If the CONFIG LED is not lit, complete the following steps to correct the problem: <ol style="list-style-type: none"> Check the riser-card LEDs, the ServeRAID error LED, and the optional network adapter error LED to identify the component that caused the error. Check the system-error log for information about the error. If you cannot isolate the failing component by using the LEDs and the information in the system-error log, remove one component at a time; and restart the file module after each component is removed. Replace the following components, in the order that is shown, restarting the file module each time: <ul style="list-style-type: none"> PCI riser cards ServeRAID adapter Optional network adapter (Trained technician only) System board If the failure remains, go to http://www.ibm.com/systems/support/supportsite.wss/docdisplay?brandind=5000008&Indocid=SERV-CALL. If the PCI LED and the CONFIG LED are lit, complete the following steps to correct the problem: <ol style="list-style-type: none"> Check the microprocessor that is installed is Intel E5-2690. Remove the high-power (>25 Watt) adapter. Check the system-error logs for information about the error. Replace any component that is identified in the error log.
NMI	An interrupt that cannot be masked occurred, or the NMI button was pressed.	<ol style="list-style-type: none"> Check the system-error log for information about the error. Restart the file module.

Table 43. LED indicators, corresponding problem causes, and corrective actions (continued)

<ul style="list-style-type: none"> Follow the suggested actions in the order in which they are listed in the Action column until the problem is solved. If an action step is preceded by "(Trained technician only)," that step must be completed only by a trained technician. 		
LED	Problem	Action
CONFIG	A hardware configuration error occurred.	<ol style="list-style-type: none"> If the CONFIG LED and the PS LED are lit, the system issues an invalid power configuration error. Make sure that both power supplies installed in the file module are of the same rating or wattage. If the CONFIG LED and the PCI LED are lit, complete the following steps to correct the problem: <ol style="list-style-type: none"> Check the microprocessor that is installed is Intel E5-2690. Remove the high-power (>25 Watt) adapter. Check the system-error logs for information about the error. Replace any component that is identified in the error log. If the CONFIG LED and the CPU LED are lit, complete the following steps to correct the problem: <ol style="list-style-type: none"> Check the microprocessors that were installed to make sure that they are compatible with each other. (Trained technician only) Replace the incompatible microprocessor. Check the system-error logs for information about the error. Replace any component that is identified in the error log. If the CONFIG LED and the MEM LED are lit, check the system-event log in the Setup utility or IMM2 error messages. For more information, see the <i>Problem Determination and Service Guide</i>. If the CONFIG LED and the HDD LED are lit, complete the following steps to correct the problem: <ol style="list-style-type: none"> Check the microprocessor that is installed is Intel E5-2690. If it is, check that the 2.5-inch hard disk drives installed are lesser than eight. Check the system-error logs for information about the error. Replace any component that is identified in the error log.
LINK	Reserved.	

Table 43. LED indicators, corresponding problem causes, and corrective actions (continued)

<ul style="list-style-type: none"> Follow the suggested actions in the order in which they are listed in the Action column until the problem is solved. If an action step is preceded by "(Trained technician only)," that step must be completed only by a trained technician. 		
LED	Problem	Action
CPU	When only the CPU LED is lit, a microprocessor has failed. When both the CPU and CONFIG LEDs are lit, the microprocessor configuration is invalid.	<ol style="list-style-type: none"> If the CONFIG LED is not lit, a microprocessor failure occurs, complete the following steps: <ol style="list-style-type: none"> (Trained technician only) Make sure that the failing microprocessor and its heat sink, which are indicated by a lit LED on the system board, are installed correctly. (Trained technician only) Replace the failing microprocessor. For more information, go to http://www.ibm.com/systems/support/supportsite.wss/docdisplay?brandind=5000008&Indocid=SERV-CALL. If the CONFIG LED and the CPU LED are lit, the system issues an invalid microprocessor configuration error. Complete the following steps to correct the problem: <ol style="list-style-type: none"> Check the microprocessors that were installed to make sure that they are compatible with each other. (Trained technician only) Replace the incompatible microprocessor. Check the system-error logs for information about the error. Replace any component that is identified in the error log.
MEM	When only the MEM LED is lit, a memory error has occurred. When both the MEM and CONFIG LEDs are lit, the memory configuration is invalid.	<p>Note: Each time that you install or remove a DIMM, you must disconnect the file module from the power source; then, wait 10 seconds before you restart the file module.</p> <ol style="list-style-type: none"> If the CONFIG LED is not lit, the system might detect a memory error. Complete the following steps to correct the problem: <ol style="list-style-type: none"> Update the file module firmware to the latest level. For more information, see the <i>Problem Determination and Service Guide</i>. Reseat or swap the DIMMs. Check the system-event log in the Setup utility or IMM error messages. For more information, see the <i>Problem Determination and Service Guide</i>. Replace the failing DIMM. If the MEM LED and the CONFIG LED are lit, check the system-event log in the Setup utility or IMM error messages. For more information, see the <i>Problem Determination and Service Guide</i>.

Table 43. LED indicators, corresponding problem causes, and corrective actions (continued)

<ul style="list-style-type: none"> • Follow the suggested actions in the order in which they are listed in the Action column until the problem is solved. • If an action step is preceded by "(Trained technician only)," that step must be completed only by a trained technician. 		
LED	Problem	Action
TEMP	The system or the system component temperature has exceeded a threshold level. A failing fan can cause the TEMP LED to be lit.	<ol style="list-style-type: none"> 1. Make sure that the heat sink is seated correctly. 2. Determine whether a fan has failed. If it has, replace it. 3. Make sure that the room temperature is not too high. 4. Make sure that the air vents are not blocked. 5. Make sure that the heat sink, the fan on the adapter, or the optional network adapter is seated correctly. If the fan has failed, replace it. 6. If the failure remains, go to http://www.ibm.com/systems/support/supportsite.wss/docdisplay?brandind=5000008&Indocid=SERV-CALL.
FAN	A fan that failed, is operating too slowly, or is removed. The TEMP LED might also be lit.	<ol style="list-style-type: none"> 1. Reseat the failing fan, which is indicated by a lit LED near the fan connector on the system board. 2. Replace the failing fan.
BOARD	An error occurred on the system board.	<ol style="list-style-type: none"> 1. Check the LEDs on the system board to identify the component that caused the error. The BOARD LED can be lit due to any of the following reasons: <ul style="list-style-type: none"> • Battery • (Trained technician only) System board 2. Check the system-error log for information about the error. 3. Replace the failing component: <ul style="list-style-type: none"> • Battery • (Trained technician only) System board

Table 43. LED indicators, corresponding problem causes, and corrective actions (continued)

<ul style="list-style-type: none"> Follow the suggested actions in the order in which they are listed in the Action column until the problem is solved. If an action step is preceded by "(Trained technician only)," that step must be completed only by a trained technician. 		
LED	Problem	Action
HDD	A hard disk drive has failed or is missing.	<ol style="list-style-type: none"> If the CONFIG LED is not lit, complete the following steps to correct the problem: <ol style="list-style-type: none"> Check the LEDs on the hard disk drives for the drive with a lit status LED and reseal the hard disk drive. Reseat the hard disk drive backplane. For more information, see the "hard disk drive problems" under the Troubleshooting tables in the <i>Problem Determination and Service Guide</i>. If the error remains, replace the following components one at a time, in the order that is listed, restarting the file module after each: <ol style="list-style-type: none"> Replace the hard disk drive. Replace the hard disk drive backplane. If the problem remains, go to http://www.ibm.com/systems/support/supportsite.wss/docdisplay?brandind=5000008&Indocid=SERV-CALL. If the HDD LED and the CONFIG LED are lit, complete the following steps to correct the problem: <ol style="list-style-type: none"> Check that the microprocessor installed is Intel E5-2690. If it is, check that the 2.5-inch hard disk drives installed are lesser than eight. Check the system-error logs for information about the error. Replace any component that is identified in the error log.

Power-supply LEDs

LEDs on the operator information panel of the file module indicate the cause of a problem. The topic describes the suggested actions to correct the detected problems.

The following minimum configuration is required for the DC LED on the power supply to be lit:

- Power supply
- Power cord

Note: You must turn on the file module for the DC LED on the power supply to be lit.

The following minimum configuration is required for the file module to start:

- One microprocessor in microprocessor socket 1
- One 2 GB DIMM on the system board
- One power supply
- Power cord
- Four cooling fans (fan 1, 2, 3, and 5)
- One PCI riser-card assembly in PCI connector 1

The following illustration shows the locations of the power-supply LEDs on the AC power supply.

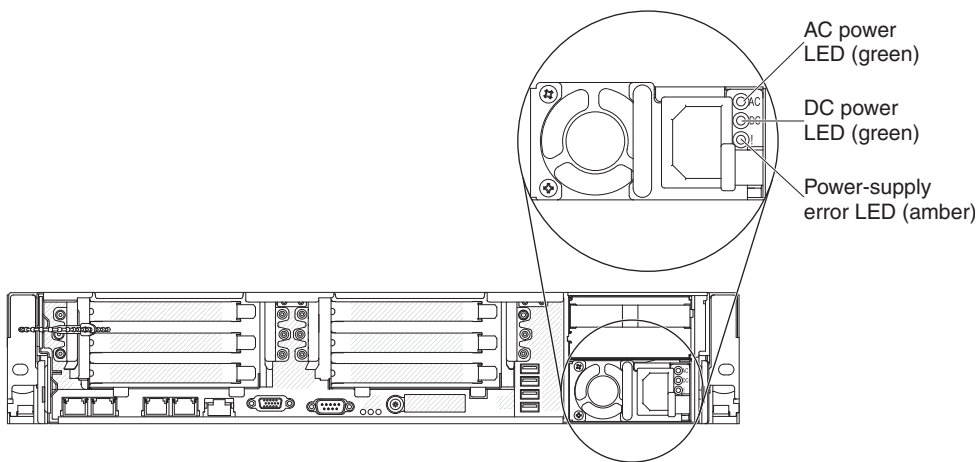


Figure 43. Locations of the power-supply LEDs

The following table describes the problems that are indicated by various combinations of the power-supply LEDs and the power-on LED on the operator information panel and suggested actions to correct the detected problems.

AC power-supply LEDs			Description	Action	Notes
AC	DC	Error (!)			
On	On	Off	Normal operation.		
Off	Off	Off	No AC power to the file module or a problem with the AC power source.	<ol style="list-style-type: none"> 1. Check the AC power to the file module. 2. Make sure that the power cord is connected to a functioning power source. 3. Restart the file module. If the error remains, check the power-supply LEDs. 4. If the problem remains, replace the power-supply. 	This is a normal condition when no AC power is present.
Off	Off	On	The power supply has failed.	Replace the power supply.	
Off	On	Off	The power supply has failed.	Replace the power supply.	
Off	On	On	The power supply has failed.	Replace the power supply.	

AC power-supply LEDs			Description	Action	Notes
AC	DC	Error (!)			
On	Off	Off	Power supply not fully seated, faulty system board, or the power supply has failed.	<ol style="list-style-type: none"> 1. Reseat the power supply. 2. If the OVER SPEC LED on the light path diagnostics is lit, follow the actions in Light path diagnostics LEDs. 3. If the OVER SPEC LED on the light path diagnostics is not lit, check the error LEDs on the system board and the IMM2 error messages. 	Typically indicates a power-supply is not fully seated.
On	Off	On	The power supply has failed.	Replace the power supply.	
On	On	On	The power supply has failed.	Replace the power supply.	

Enclosure hardware indicators

The LEDs provide a general idea of the volume system status.

For specifics about the status of control enclosures, expansion enclosures, node canisters, and expansion canisters, see Chapter 1, “Storwize V7000 Unified hardware components,” on page 1. Also refer to “Procedure: Understanding the system status using the LEDs” on page 255.

Management GUI interface

The management GUI is a browser-based GUI for configuring and managing all aspects of your system. It provides extensive facilities to help troubleshoot and correct problems.

About this task

You use the management GUI to manage and service your system. The **Monitoring > Events** panel provides access to problems that must be fixed and maintenance procedures that step you through the process of correcting the problem.

Two tabs are available for monitoring events:

- A **Block** tab for monitoring the SAN volume events and the file system volume events from the control enclosure.
- A **File** tab for monitoring the NAS events from the Storwize V7000 Unified file modules.

When you click the **Block** tab, a **Next recommended action** is shown. Perform the next recommended action before attempting any other recommended actions.

The information on the Events panel can be filtered three ways:

Recommended action (default)

Shows only the alerts that require attention and have an associated fix procedure. Alerts are listed in priority order and should be fixed sequentially by using the available fix procedures. For each problem that is selected, you can:

- Run a fix procedure.
- View the properties.

Unfixed messages and alerts

Displays only the alerts and messages that are not fixed. For each entry that is selected, you can:

- Run a fix procedure on any alert with an error code.
- Mark an event as fixed.
- Filter the entries to show them by specific minutes, hours, or dates.
- Reset the date filter.
- View the properties.

Show all

Displays all event types whether they are fixed or unfixed. For each entry that is selected, you can:

- Run a fix procedure on any alert with an error code.
- Mark an event as fixed.
- Filter the entries to show them by specific minutes, hours, or dates.
- Reset the date filter.
- View the properties.

Some events require a certain number of occurrences in 25 hours before they are displayed as unfixed. If they do not reach this threshold in 25 hours, they are flagged as expired. Monitoring events are below the coalesce threshold and are usually transient.

You can also sort events by time or error code. When you sort by error code, the most serious events, those with the lowest numbers, are displayed first. You can select any event that is listed and select **Actions > Properties** to view details about the event.

- Recommended Actions. For each problem that is selected, you can:
 - Run a fix procedure.
 - View the properties.
- Event log. For each entry that is selected, you can:
 - Run a fix procedure.
 - Mark an event as fixed.
 - Filter the entries to show them by specific minutes, hours, or dates.
 - Reset the date filter.
 - View the properties.

When to use the management GUI

The management GUI is the primary tool that is used to service your system.

Regularly monitor the status of the system using the management GUI. If you suspect a problem, use the management GUI first to diagnose and resolve the problem.

Use the views that are available in the management GUI to verify the status of the system, the hardware devices, the physical storage, and the available volumes. The

Monitoring > Events panel provides access to all problems that exist on the system. Use the **Recommended Actions** filter to display the most important events that need to be resolved.

If there is a service error code for the alert, you can run a fix procedure that assists you in resolving the problem. These fix procedures analyze the system and provide more information about the problem. They suggest actions to take and step you through the actions that automatically manage the system where necessary. Finally, they check that the problem is resolved.

If there is an error that is reported, always use the fix procedures within the management GUI to resolve the problem. Always use the fix procedures for both system configuration problems and hardware failures. The fix procedures analyze the system to ensure that the required changes do not cause volumes to be inaccessible to the hosts. The fix procedures automatically perform configuration changes that are required to return the system to its optimum state.

Accessing the Storwize V7000 Unified management GUI

This procedure describes how to access the Storwize V7000 Unified management GUI.

About this task

You must use a supported web browser. Verify that you are using a supported web browser. Checking your web browser settings for the management GUI from the Storwize V7000 Information Center.

www.ibm.com/storage/support/storwize/v7000/unified

You can use the management GUI to manage your system as soon as you have completed the USB flash drive initialization.

Procedure

1. Start a supported web browser and point the browser to the management IP address of the file module.
The management IP address is set during the USB flash drive initialization.
2. When the connection is successful, you see a login panel.
3. Log on by using your user name and password. The default user name is admin.
4. When you have logged on, select **Monitoring > Events**.
5. Ensure that the events log is filtered by using **Recommended actions**.
6. Select the recommended action and run the fix procedure.
7. Continue to work through the alerts in the order suggested, if possible.

Results

After all the alerts are fixed, check the status of your system to ensure that it is operating as intended.

If you encounter problems logging on the management GUI or connecting to the management GUI, see “Problem: Unable to log on to the management GUI” on page 240 or “Problem: Unable to connect to the Storwize V7000 Gen1 management GUI” on page 239.

Diagnosing and resolving problems with fix procedures

You can use fix procedures to diagnose and resolve problems with the Storwize V7000 Unified.

About this task

For example, to repair a Storwize V7000 Unified system, you might complete the following tasks:

- Analyze the event log (if it is available, or view node errors)
- Replace failed components
- Verify the status of a repaired device
- Restore a device to an operational state in the system
- Mark the error as fixed in the event log

Fix procedures help simplify these tasks by automating as many of the tasks as possible.

Many of the file module fix procedures are not automated. In these cases, you are directed to a documented procedure in the Storwize V7000 Unified Information Center.

The example uses the management GUI to repair a Storwize V7000 Unified system.

Procedure

Complete the following steps to start the fix procedure.

1. Click **Monitoring > Events** and ensure that you are filtering the event log to display **Recommended actions**.
The list might contain any number of errors that must be repaired. If there is more than one error on the list, the error at the top of the list has the highest priority and must always be fixed first. If you do not fix the higher priority errors first, you might not be able to fix the lower priority errors.
2. Select the error at the top of the list or select the **Next recommended action**.
3. Click **Run Fix Procedure**.
The pane displays the error code and provides a description of the condition.
4. Click **Next** to go forward or **Cancel** to return to the previous pane. One or more panes might be displayed with instructions for you to replace parts or complete other repair activity.
5. If you are not able to complete the actions now, click **Cancel** until you return to the previous pane. Click **Cancel** until you are returned to the Next Recommended Actions pane. When you return to the fix procedures, the repair can be restarted from step 1. After you complete all the instructions, click **OK**. When the last repair action is completed, the procedures might attempt to restore failed devices to the system.
6. After you complete the fix, you see the statement Click OK to mark the error as fixed. Click **OK**. This action marks the error as fixed in the event log and prevents this instance of the error from being listed again.
7. When you see the statement The repair has been completed., click **Exit**. If other errors must be fixed, those errors are displayed and the fix procedures continue.
8. If no errors remain, you are shown the following statement: There are no unfixed errors in the event log.

Chapter 4. File module

This topic provides information about troubleshooting the file module, which includes error codes, problem scenarios, software troubleshooting, and removal and replacement instructions.

General file module procedures

This section covers file module general maintenance and repair issues.

Rebooting a file module

Use this procedure to initiate a file module reboot.

Before you begin

Events can occur on a file module that require the hardware to be rebooted.

Procedure

1. To shut down and restart a node by using the management GUI, follow these steps:
 - a. Click **Monitoring > System Details**.
 - b. Click the **Interface Nodes** tab.
 - c. In the left pane, select the node to reboot. In the right pane, click **Actions > Restart**.

Note: If the file module to be rebooted is the active management node, the management GUI will also shut down and cause the management GUI to become unresponsive. After the management services have failed over to the other file module, a refresh of the management GUI in the browser reestablishes the connection.

2. To shut down and reboot a node by using the command-line interface (CLI) command, enter:

```
stopcluster -node mgmt00Xst001 -restart
```

where *X* is the logical ID of the node to reboot.

3. The node reboot restarts all services that were previously running *g* but does not resume the node on to the cluster.

```
resumenode mgmt00Xst001
```

Removing a file module to perform a maintenance action

You can remove an IBM Storwize V7000 Unified file module to perform maintenance. The procedure that you follow differs slightly, depending on whether you must unplug the power cables.

Before you begin

If you receive an alert event that requires you to service a file module, use the following procedure to remove the file module from the system and perform the required service.

About this task

Some field replaceable units (FRUs) are redundant and hot swappable, such as power supplies. When replacing a hot-swap FRU, you have the option of leaving the file module turned on and the power cables connected. Always follow the remove and replace procedure for the FRU. The procedure for the FRU indicates whether the FRU is hot swappable.

If the removal and replacement procedure does not indicate whether the FRU is hot swappable, assume that it is not. In that case use the file module-removal procedure that requires you to disconnect the power cords.

Note: Before removing an file module, you must suspend the file module.

Procedure

- Remove a file module from the system to replace a hot swappable FRU, as described in “Removing a file module without disconnecting power” on page 93.
- Remove a file module from the system, turn off the node, and disconnect the power cords, as described in “Removing a file module and disconnecting power.”

Removing a file module and disconnecting power

You must remove an IBM Storwize V7000 file module from the file cluster and disconnect it from its power line cords before performing a maintenance action that requires the file module to have no power.

About this task

To identify and perform a service action on any file module that requires you to turn off the power before performing the service action, perform the following procedure.

Procedure

1. Access and log in to the Storwize V7000 Unified system from the command-line interface.
2. Suspend the file module. Use the `suspendnode` command on one of the file modules that you need to maintain, as shown in the following examples, to stop a file module from providing the services.

- `suspendnode mgmt001st001`
- `suspendnode mgmt002st001`

A suspended file module does not participate in the cluster and does not host any records for the clustered trivial database (CTDB). The IP addresses of a file module are taken over by the other file module and no services are lost. You can review the status of the file module by using the `lsnode` command with the `-r` option. Review the row for the file module that was suspended and the column for the **Connection Status**.

3. Use the `stopcluster` command to remove the file module from the system and shut down the file module.

If you are shutting down the `mgmt001st001` file module, for example, issue the following command:

```
stopcluster -n mgmt001st001
```

4. After the file module shuts down and the power indicator light on the front of the file module is flashing slowly, pull the file module out on its rails.

Note: Label and disconnect both power cords and all external cables from the file module.

5. Remove the file module from the rack if necessary, or locate and use the service ladder, if necessary, to perform the maintenance action on the file module when it is fully extended from the rack.
6. Locate and perform the correct removal and replacement procedure.
Attention: You can replace only one of the disk drives in the file module. If you must replace both disk drives, contact IBM Remote Technical Support.
7. After replacing the failing part and replacing the file module cover, replace the file module in the rack, if necessary, and reconnect the power cords.
After reconnecting the power cords, the power indicator LED on the front of the file module begins to flash quickly.
8. Push the file module back into the rack.
9. After the power indicator LED on the front of the file module begins to flash slowly, press the power switch that surrounds the indicator light to turn on the file module.
As the file module reboots, the Storwize V7000 Unified system reintegrates it back into the cluster.
10. After the file module is fully booted back into the system, resume the file module using the `resumenode` command that was previously suspended and shutdown.

Note: Set the machine serial number and product name as described in “Setting the machine serial number” on page 181 prior to resuming the node.

Removing a file module without disconnecting power

You can work on an IBM Storwize V7000 Unified file module to perform a maintenance action that does not require you to remove its power cords.

About this task

Perform the following procedure to remove and replace a hot swappable field replaceable unit (FRU) in a file module when you do not have to remove the file module from the rack to work on it.

Procedure

1. Access and log in to the Storwize V7000 Unified system from the command-line interface.
2. Issue the **suspendnode** command to remove the file module from the system so that you can work on it.

To remove the `mgmt001st001` file module from the system, for example, issue the following command:

```
# suspendnode mgmt001st001
```

3. Wait for the Storwize V7000 Unified system to stop the file module at the clustered trivial database (CTDB) level. The command does not unmount any mounted file systems.

A stopped file module does not participate in the cluster and does not host any records for the clustered trivial database. The IP address of an file module is taken over by another file module and no services are hosted.

You can issue the **lsnode -r** command to view the state of the file module.

The results from running the **lsnode -r** command are similar to the following example:

```
# lsnode -r
```

Hostname	IP	Description	Role
mgmt001st001	10.254.8.2	active management node	management,interface,storage
mgmt002st001	10.254.8.3	passive management node	management,interface,storage

Product Version	Connection	status	GPFS status	CTDB status	Last updated
1.3.0.2-02	OK	active	active	active	1/17/12 4:39 PM
1.3.0.2-02	SUSPEND	active	active	SUSPEND_MAINTENANCE	1/17/12 4:39 PM

4. Pull the file module out from the rack on its rails.
5. Locate and use the service ladder, if necessary, to perform the maintenance action on the file module when it is fully extended from the rack.
6. Locate and perform the correct removal and replacement procedure.
Attention: You can replace only one of the disk drives in the file module. If you must replace both disk drives, contact IBM Remote Technical Support.
7. After replacing the failing part and replacing the file module cover, push the file module back in the rack.
8. Use the **resumenode** command to add the file module back into the system so that it can begin to host services.

To add the mgmt001st001 file module back into the system, for example, issue the following command:

```
# resumenode mgmt001st001
```

9. After the Storwize V7000 Unified system reintegrates the file module back into the cluster, the ctdb status command shows that the service is active on the file module.

Removing and replacing file module components

The IBM Storwize V7000 Unified system contains parts that are both customer replaceable units (CRUs) and field replaceable units (FRUs). CRUs can be installed by the customer, but all FRUs must be installed by trained service technicians.

About this task

Installation guidelines

To help you work safely with IBM Storwize V7000 Unified file modules, read the safety information in , Safety information statements, and these guidelines.

Before you remove or replace a component, read the following information:

- When you install a file module, take the opportunity to download and apply the most recent firmware updates. This step helps to ensure that any known issues are addressed and that your file module is ready to function at maximum levels of performance.
- Before you install any hardware, make sure that the file module is working correctly. Start the file module, and make sure that the Linux operating system starts. If the file module is not working correctly, see Chapter 3, “Getting started troubleshooting,” on page 47 for diagnostic information.
- Observe good housekeeping in the area where you are working. Place removed covers and other parts in a safe place.
- If you must start the file module while the cover is removed, make sure that no one is near the file module and that no tools or other objects have been left inside the file module.
- Do not attempt to lift an object that you think is too heavy for you. If you have to lift a heavy object, observe the following precautions:
 - Make sure that you can stand safely without slipping.
 - Distribute the weight of the object equally between your feet.

- Use a slow lifting force. Never move suddenly or twist when you lift a heavy object.
- To avoid straining the muscles in your back, lift by standing or by pushing up with your leg muscles.
- Make sure that you have an adequate number of properly grounded electrical outlets for the PDUs.
- Back up all important data before you make changes to disk drives.
- Have a small flat-blade screwdriver available.
- To view the error LEDs on the system board and internal components, leave the file module connected to power.
- You do not have to turn off the file module to install or replace hot-swap fans, redundant hot-swap ac power supplies, or hot-plug Universal Serial Bus (USB) devices. However, you must turn off the file module before performing any steps that involve removing or installing adapter cables or non-hot-swap optional devices or components.
- Blue on a component indicates touch points, where you can grip the component to remove it from or install it in the file module, open or close a latch, and so on.
- Orange on a component or an orange label on or near a component indicates that the component can be hot-swapped, which means that if the file module and operating system support hot-swap capability, you can remove or install the component while the file module is running. (Orange can also indicate touch points on hot-swap components.) See the instructions for removing or installing a specific hot-swap component for any additional procedures that you might have to perform before you remove or install the component.
- When you are finished working on the file module, reinstall all safety shields, guards, labels, and ground wires.

Node reliability guidelines

To help ensure proper cooling and system reliability, make sure that:

- Each of the drive bays has a drive or a filler panel and electromagnetic compatibility (EMC) shield installed in it.
- If the server has redundant power, each of the power-supply bays has a power supply installed in it.
- There is adequate space around the server to allow the server cooling system to work properly. Leave approximately 50 mm (2.0 in.) of open space around the front and rear of the server. Do not place objects in front of the fans. For proper cooling and airflow, replace the server cover before turning on the server. Operating the server for extended periods of time (more than 30 minutes) with the server cover removed might damage server components.
- You have followed the cabling instructions that come with optional adapters.
- You have replaced a failed fan within 48 hours.
- You have replaced a hot-swap drive within 2 minutes of removal.
- You operate the server with the air baffles installed. Operating the server without the air baffles might cause the microprocessor to overheat.

Working inside the file module with the power on

Attention: Static electricity that is released to internal file module components when the file module is powered-on might cause the file module to halt, which could result in the loss of data. To avoid this potential problem, always use an electrostatic-discharge wrist strap or other grounding system when working inside the file module with the power on.

The file module supports hot-plug, hot-add, and hot-swap devices and is designed to operate safely while it is turned on and the cover is removed. Follow these guidelines when you work inside a file module that is turned on:

- Avoid wearing loose-fitting clothing on your forearms. Button long-sleeved shirts before working inside the file module; do not wear cuff links while you are working inside the file module.
- Do not allow your necktie or scarf to hang inside the file module.
- Remove jewelry, such as bracelets, necklaces, rings, and loose-fitting wrist watches.
- Remove items from your shirt pocket, such as pens and pencils, that could fall into the file module as you lean over it.
- Avoid dropping any metallic objects, such as paper clips, hairpins, and screws, into the file module.

Handling static-sensitive devices

Attention: Static electricity can damage the server and other electronic devices. To avoid damage, keep static-sensitive devices in their static-protective packages until you are ready to install them.

To reduce the possibility of damage from electrostatic discharge, observe the following precautions:

- Limit your movement. Movement can cause static electricity to build up around you.
- The use of a grounding system is recommended. For example, wear an electrostatic-discharge wrist strap, if one is available. Always use an electrostatic-discharge wrist strap or other grounding system when working inside the server with the power on.
- Handle the device carefully, holding it by its edges or its frame.
- Do not touch solder joints, pins, or exposed circuitry.
- Do not leave the device where others can handle and damage it.
- While the device is still in its static-protective package, touch it to an unpainted metal surface on the outside of the server for at least 2 seconds. This drains static electricity from the package and from your body.
- Remove the device from its package and install it directly into the server without setting down the device. If it is necessary to set down the device, put it back into its static-protective package. Do not place the device on the server cover or on a metal surface.
- Take additional care when handling devices during cold weather. Heating reduces indoor humidity and increases static electricity.

Returning a device or component

When returning a device or component, follow all packaging instructions and use any supplied packaging materials for shipping.

Resolving hard disk drive problems

Use this information to address various hard disk drive issues.

About this task

<ul style="list-style-type: none"> • Before running a procedure, refer to “Removing a file module to perform a maintenance action” on page 91. • Follow the suggested actions for a Symptom in the order in which they are listed in the Action column until the problem is solved. • If an action step is preceded by “(Trained service technician only)”, that step must be performed only by a trained service technician. 	
Symptom	Action
A hard disk drive has failed and the associated amber hard disk drive status LED is lit.	Replace the failed hard disk drive.
An installed hard disk drive is not recognized.	<ol style="list-style-type: none"> 1. Observe the associated amber hard disk drive status LED. If the LED is lit, it indicates a drive fault. 2. If the LED is lit, remove the drive from the bay, wait 45 seconds, then reinsert the drive, ensuring that the drive assembly connects to the hard disk drive backplane. 3. Observe the associated green hard disk drive activity LED and the amber status LED: <ul style="list-style-type: none"> • If the green activity LED is flashing and the amber status LED is not lit, the drive is recognized by the controller and is working correctly. Run the DSA hard disk drive test to determine whether the drive is detected. • If the green activity LED is flashing and the amber status LED is flashing slowly, the drive is recognized by the controller and is rebuilding. • If neither LED is lit or flashing, check the hard disk drive backplane (go to step 4). • If the green activity LED is flashing and the amber status LED is lit, replace the drive. If the activity of the LEDs remains the same, go to step 4. If the activity of the LEDs changes, return to step 1. 4. Ensure that the hard disk drive backplane is correctly seated. When it is correctly seated, the drive assemblies correctly connect to the backplane without bowing or causing movement of the backplane. 5. Move the hard disk drives to different bays to determine if the drive or the backplane is not functioning. 6. Re-seat the backplane power cable and repeat steps 1 through 3. 7. Re-seat the backplane signal cable and repeat steps 1 through 3. 8. Suspect the backplane signal cable or the backplane: <ul style="list-style-type: none"> • If the server has eight hot-swap bays: <ol style="list-style-type: none"> a. Replace the affected backplane signal cable. b. Replace the affected backplane. • If the server has 12 hot-swap bays: <ol style="list-style-type: none"> a. Replace the backplane signal cable. b. Replace the backplane. c. Replace the SAS expander card.
Multiple hard disk drives fail.	<p>Ensure that the hard disk drive, SAS RAID controller, and server device drivers and firmware are of the latest version.</p> <p>Important: Some cluster solutions require specific code levels or coordinated code updates. If the device is part of a cluster solution, check whether the latest code version is supported before you update the code.</p>
Multiple hard disk drives are offline.	<ol style="list-style-type: none"> 1. Review the storage subsystem logs for indications of problems within the storage subsystem, such as backplane or cable problems.

<ul style="list-style-type: none"> • Before running a procedure, refer to “Removing a file module to perform a maintenance action” on page 91. • Follow the suggested actions for a Symptom in the order in which they are listed in the Action column until the problem is solved. • If an action step is preceded by “(Trained service technician only)”, that step must be performed only by a trained service technician. 	
Symptom	Action
A replacement hard disk drive does not rebuild.	<ol style="list-style-type: none"> 1. Ensure that the hard disk drive is recognized by the controller (the green hard disk drive activity LED is flashing). 2. Review the SAS RAID controller documentation to determine the correct configuration parameters and settings.
A green hard disk drive activity LED does not accurately represent the actual state of the associated drive.	<ol style="list-style-type: none"> 1. If the green hard disk drive activity LED does not flash when the drive is in use, run the DSA Preboot diagnostic programs to collect error logs. Refer to the “Diagnostics” or “Running the diagnostic programs” section in “Troubleshooting the System x3650” in the <i>IBM Storwize V7000 Unified Information Center</i>. 2. Use one of the following procedures: <ul style="list-style-type: none"> • If the drive passes the test, replace the backplane. • If the drive fails the test, replace the drive.
An amber hard disk drive status LED does not accurately represent the actual state of the associated drive.	<ol style="list-style-type: none"> 1. If the amber hard disk drive LED and the RAID controller software do not indicate the same status for the drive, complete the following steps: <ol style="list-style-type: none"> a. Turn off the server. b. Re-seat the SAS controller. c. Re-seat the backplane signal cable, backplane power cable, and SAS expander card (if the server has 12 drive bays). d. Re-seat the hard disk drive. e. Turn on the server and observe the activity of the hard disk drive LEDs.

Displaying node mirror and hard drive status

The Storwize V7000 Unified system provides a method to check the node mirror status and hard drive status for each file module.

About this task

As a privileged user, you can run a perl script to verify whether or not mirroring is configured. By displaying the mirror status, you can view information that shows the location of each hard drive, the status values of each hard drive, and any errors, if applicable. If the mirror status is re-synchronizing, information that shows the percentage complete for the resynchronization is displayed.

Procedure

1. Ensure that you are logged into the file module as root.
2. To display mirror status and hard drive status, run the following perl script:

```
# sc /opt/IBM/sonas/bin/cnrspromptnode.pl -a -c "/opt/IBM/sonas/bin/cnrsQueryNodeDrives.pl"
```

File modules in this Storwize V7000 Unified Cluster

Node	Node Name	Node Details
1.	mgmt001st001	x3650m3 KQ186WX
2.	mgmt002st001	x3650m3 KQ186WV

B. Back to Menus
Choice:

Figure 44. Selecting a file module to display node status

3. Select the number for a file module to display its status. For example, type **1** to select **mgmt001st001**. Press **Enter** to display the information in Figure 45 on page 100, which shows an example of a healthy status for the mirroring and drive status. The output shows a file module with two hard disk drives.

```

Mirror Information:
  Volume ID                : 3
  Status of volume         : Okay (OKY)
  RAID level               : 1
  Size (in MB)             : 285148
  Physical hard disks (Target ID) : 6 5
  Current operation        : None
  Physical disk I/Os       : Not quiesced

Drive Information:
Total number of drives found: 2

Target on ID #5
  Device is a Hard disk
  Enclosure #              : 1
  Slot #                   : 1
  Connector ID             : 1
  Target ID                : 5
  State                    : Online (ONL)
  Size (in MB)/(in sectors) : 286102/585937500
  Manufacturer             : IBM-ESXS
  Model Number             : XXXXXXXXXXXXX
  Firmware Revision        : XXXX
  Serial No                : XXXXXXXXXXXXXXXXXXXXX
  Drive Type               : SAS
  Protocol                 : SAS
  Error Information
    SMART Error Count      : none
    SMART ASC              : none
    SMART ASCQ             : none

Target on ID #6
  Device is a Hard disk
  Enclosure #              : 1
  Slot #                   : 0
  Connector ID             : 0
  Target ID                : 6
  State                    : Online (ONL)
  Size (in MB)/(in sectors) : 286102/585937500
  Manufacturer             : IBM-ESXS
  Model Number             : XXXXXXXXXXXXX
  Firmware Revision        : XXXX
  Serial No                : XXXXXXXXXXXXXXXXXXXXX
  Drive Type               : SAS
  Protocol                 : SAS
  Error Information
    SMART Error Count      : none
    SMART ASC              : none
    SMART ASCQ             : none

```

Figure 45. Displaying node status

4. Review the section for **Mirror Information** and the value for **Status of volume**, then see Table 44 for the possible values for **Status of volume**.

Table 44. Status of volume

Status of volume	Description
Okay (OKY)	The volume is Active and the drives are functioning correctly. The user data is protected if the volume is integrated mirroring or integrated mirroring enhanced.
Degraded (DGD)	The volume is Active. The user data is not fully protected due to a configuration change or drive failure.
Rebuilding (RBLD) or Resyncing (RSY)	A data resynchronization or rebuild might be in progress.

Table 44. Status of volume (continued)

Status of volume	Description
Inactive, Okay (OKY)	The volume is inactive and the drives are functioning correctly. The user data is protected if the current RAID level is RAID 1 (IM) or RAID 1E (IME).
Inactive, Degraded (DGD)	The volume is inactive and the user data is not fully protected due to a configuration change or drive failure; a data resync or rebuild might be in progress.

5. Review the section for **Drive information** and the value for **State** Figure 45 on page 100, then see Table 45 to see the possible values for **State** of the drives.

Table 45. State of drives

Status of drives	Description
Online (ONL)	The drive is operational and is part of a logical drive.
Hot Spare (HSP)	The drive is a hot spare that is available for replacing a failed drive in an array.
Ready (RDY)	The drive is ready for use as a normal disk drive; or it is available to be assigned to a disk array or hot spare pool.
Available (AVL)	The drive might or might not be ready, and it is not suitable for inclusion in an array or hot spare pool (for example, it did not spin up, its block size is incorrect, or its media is removable).
Failed (FLD)	The drive was part of a logical drive or was a hot spare drive, and it has failed. It has been taken offline.
Standby (SBY)	This status is used to tag all non-hard disk drive devices.
Missing (MIS)	The hard drive might be removed.
Out of Sync (OSY)	A data resynchronization or rebuild might be in progress.

6. See Figure 46 on page 102 for an example that shows that mirroring is re-synchronizing. If a hard disk drive is removed and reinserted, the array starts to resynchronize automatically.

Notelist: You can tell that the mirroring is re-synchronizing when the following conditions are true:

- **State of volume** is **Resyncing (RSY)**
- **Current operation** is **Synchronize**
- **Percentage complete** is displayed

A mirror/volume consists of two hard drives. In Figure 46 on page 102, the section for **Mirror Information** has a status line called **Physical hard disk (Target ID)**. The line shows which drives are part of the mirror/volume.

The **Status of volume** shows **Resyncing (RSY)**.

The mirror consists of **Physical hard disk (Target ID)** of 6 and 9. Drive 9 is in a **State of Out of Sync (OSY)**. The **Mirror Information** will also show you the percentage complete for the resynchronization. For example, the percentage complete in Figure 46 on page 102 is 5.23%.

```

Mirror Information:
Volume ID                : 3
Status of volume      : Resyncing (RSY) <---
RAID level               : 1
Size (in MB)             : 285148
Physical hard disks (Target ID)
Current operation     : Synchronize <---
Physical disk I/Os       : Not quiesced
Volume size (in sectors) : 583983104
Number of remaining sectors : 553462899
Percentage complete   : 5.23% <---

```

Drive Information:

Total number of drives found: 2

Target on ID #5

```

Device is a Hard disk
Enclosure #              : 1
Slot #                   : 1
Connector ID             : 1
Target ID                 : 5
State                    : Ready (RDY)
Size (in MB)/(in sectors) : 286102/585937500
Manufacturer              : IBM-ESXS
Model Number             : XXXXXXXXXXXXX
Firmware Revision        : XXXX
Serial No                 : XXXXXXXXXXXXXXXXXXXXX
Drive Type                : SAS
Protocol                 : SAS
Error Information
  SMART Error Count      : none
  SMART ASC              : none
  SMART ASCQ             : none

```

Target on ID #6

```

Device is a Hard disk
Enclosure #              : 1
Slot #                   : 0
Connector ID             : 0
Target ID                 : 6
State                    : Online (ONL)
Size (in MB)/(in sectors) : 286102/585937500
Manufacturer              : IBM-ESXS
Model Number             : XXXXXXXXXXXXX
Firmware Revision        : XXXX
Serial No                 : XXXXXXXXXXXXXXXXXXXXX
Drive Type                : SAS
Protocol                 : SAS
Error Information
  SMART Error Count      : none
  SMART ASC              : none
  SMART ASCQ             : none

```

Figure 46. Example that shows that mirroring is re-synchronizing

If a drive were not synchronized, the status might appear like the status shown in Figure 47 on page 103:

```

Target on ID #5
Device is a Hard disk
Enclosure #           : 1
Slot #               : 1
Connector ID         : 1
Target ID            : 5
State                : Out of Sync (OSY) <---
Size (in MB)/(in sectors) : 286102/585937500
Manufacturer         : IBM-ESXS
Model Number         : XXXXXXXXXXXX
Firmware Revision    : XXXX
Serial No            : XXXXXXXXXXXXXXXXXXXX
Drive Type           : SAS
Protocol             : SAS
Error Information
  SMART Error Count   : none
  SMART ASC           : none
  SMART ASCQ          : none

```

Figure 47. Example that shows that a drive is not synchronized

7. See Figure 48 on page 104 for an example of status when there is no mirror.
 If mirroring is not enabled, the output under **Mirror Information** displays a message that says: **The mirror is not created/configured.**
 If the mirror is not created, refer to “Troubleshooting the System x3650” in the *IBM Storwize V7000 Unified Information Center* for information on launching the LSI configuration tool.

```

Mirror Information:
  NOTICE: The mirror is not created/configured.      <---

Drive Information:
Total number of drives found: 2
Target on ID #4
  Device is a Hard disk
  Enclosure #           : 1
  Slot #                : 1
  Connector ID          : 1
  Target ID             : 4
  State                 : Ready (RDY)
  Size (in MB)/(in sectors) : 286102/585937500
  Manufacturer          : IBM-ESXS
  Model Number          : XXXXXXXXXXXXX
  Firmware Revision     : XXXX
  Serial No             : XXXXXXXXXXXXXXXXXXXXX
  Drive Type            : SAS
  Protocol              : SAS
  Error Information
    SMART Error Count   : none
    SMART ASC           : none
    SMART ASCQ          : none

Target on ID #6
  Device is a Hard disk
  Enclosure #           : 1
  Slot #                : 0
  Connector ID          : 0
  Target ID             : 6
  State                 : Ready (RDY)
  Size (in MB)/(in sectors) : 286102/585937500
  Manufacturer          : IBM-ESXS
  Model Number          : XXXXXXXXXXXXX
  Firmware Revision     : XXXX
  Serial No             : XXXXXXXXXXXXXXXXXXXXX
  Drive Type            : SAS
  Protocol              : SAS
  Error Information
    SMART Error Count   : none
    SMART ASC           : none
    SMART ASCQ          : none

```

Figure 48. Example that shows that the mirror is not created

8. See Figure 49 on page 105 for an example of a Self-Monitoring, Analysis and Reporting Technology (**SMART**) error found for a hard drive. SMART adds monitoring and troubleshooting functionality by automatically checking a disk drive's health and reporting potential problems. If any **SMART** errors are detected for a hard drive, you can see the status in the section for **Error Information** as shown in Figure 49 on page 105.

Note: In Figure 49 on page 105 the hard disk drive with **Target ID #6** has a **ASC/ ASCQ** error of **05/00**.

For isolation and the repair of hard disk problems, refer to “Troubleshooting the System x3650” in the *IBM Storwize V7000 Unified Information Center*.

For a list of **SMART** (ASC/ASCQ) error codes and their descriptions, go to “SMART ASC/ASCQ error codes and messages” on page 105.

```

Mirror Information:
  Volume ID                : 4
  Status of volume         : Resyncing (RSY)
  RAID level               : 1
  Size (in MB)             : 285148
  Physical hard disks (Target ID) : 6 9
  Current operation        : Synchronize
  Physical disk I/Os       : Not quiesced

Drive Information:
Total number of drives found: 2

Target on ID #6
  Device is a Hard disk
  Enclosure #              : 1
  Slot #                   : 0
  Connector ID             : 0
  Target ID                : 6
  State                    : Online (ONL)
  Size (in MB)/(in sectors) : 286102/585937500
  Manufacturer             : IBM-ESXS
  Model Number             : MBD2300RC
  Firmware Revision        : SB19
  Serial No                : D009P9A01SJC
  Drive Type               : SAS
  Protocol                 : SAS
  Error Information
    SMART Error Count      : 1
    SMART ASC              : 05*      <----
    SMART ASCQ             : 00*      <----
*See Infocenter for SMART ASC/ASCQ error codes and messages

Target on ID #9
  Device is a Hard disk
  Enclosure #              : 1
  Slot #                   : 1
  Connector ID             : 1
  Target ID                : 9
  State                    : Out of Sync (OSY)
  Size (in MB)/(in sectors) : 286102/585937500
  Manufacturer             : IBM-ESXS
  Model Number             : MBD2300RC
  Firmware Revision        : SB19
  Serial No                : D009P990184N
  Drive Type               : SAS
  Protocol                 : SAS
  Error Information
    SMART Error Count      : none
    SMART ASC              : none
    SMART ASCQ             : none

```

Figure 49. Example of a SMART error

SMART ASC/ASCQ error codes and messages

Table 46 on page 106 shows descriptions of common Self-Monitoring, Analysis and Reporting Technology (SMART) ASC/ASCQ error codes that are classified for a direct access device. The ASC (additional sense code) and ASCQ (additional sense code qualifier) are known as SCSI additional sense data codes, as defined by SCSI standards. SMART adds monitoring and troubleshooting functionality by automatically checking a disk drive's health and reporting potential problems.

Note: Values in the following table such as “5D” are the same as the “5DH” displayed in the tool; some values such as “0” might have additional padding, so that “0” will be the same as “00.”

Table 46. SMART ASC/ASCQ error codes and messages

ASC	ASCQ	Description
00	00	NO ADDITIONAL SENSE INFORMATION
00	06	I/O PROCESS TERMINATED
00	16	OPERATION IN PROGRESS
00	17	CLEANING REQUESTED
00	1D	ATA PASS THROUGH INFORMATION AVAILABLE
00	1E	CONFLICTING SA CREATION REQUEST
00	1F	LOGICAL UNIT TRANSITIONING TO ANOTHER POWER CONDITION
01	00	NO INDEX/SECTOR SIGNAL
02	00	NO SEEK COMPLETE
03	00	PERIPHERAL DEVICE WRITE FAULT
04	00	LOGICAL UNIT NOT READY
04	01	LOGICAL UNIT IS IN PROCESS OF BECOMING READY
04	02	LOGICAL UNIT NOT READY, INITIALIZING COMMAND REQUIRED
04	03	LOGICAL UNIT NOT READY, MANUAL INTERVENTION REQUIRED
04	04	LOGICAL UNIT NOT READY, FORMAT IN PROGRESS
04	05	LOGICAL UNIT NOT READY, REBUILD IN PROGRESS
04	06	LOGICAL UNIT NOT READY, RECALCULATION IN PROGRESS
04	07	LOGICAL UNIT NOT READY, OPERATION IN PROGRESS
04	09	LOGICAL UNIT NOT READY, SELF-TEST IN PROGRESS
04	0A	LOGICAL UNIT NOT ACCESSIBLE, ASYMMETRIC ACCESS STATE TRANSITION
04	0B	LOGICAL UNIT NOT ACCESSIBLE, TARGET PORT IN STANDBY STATE
04	0C	LOGICAL UNIT NOT ACCESSIBLE, TARGET PORT IN UNAVAILABLE STATE
04	10	LOGICAL UNIT NOT READY, AUXILIARY MEMORY NOT ACCESSIBLE
04	11	LOGICAL UNIT NOT READY, NOTIFY (ENABLE SPINUP) REQUIRED
04	13	LOGICAL UNIT NOT READY, SA CREATION IN PROGRESS
04	14	LOGICAL UNIT NOT READY, SPACE ALLOCATION IN PROGRESS
04	1A	LOGICAL UNIT NOT READY, START STOP UNIT COMMAND IN PROGRESS
05	00	LOGICAL UNIT DOES NOT RESPOND TO SELECTION
06	00	NO REFERENCE POSITION FOUND
07	00	MULTIPLE PERIPHERAL DEVICES SELECTED
08	00	LOGICAL UNIT COMMUNICATION FAILURE
08	01	LOGICAL UNIT COMMUNICATION TIME-OUT
08	02	LOGICAL UNIT COMMUNICATION PARITY ERROR

Table 46. SMART ASC/ASCQ error codes and messages (continued)

ASC	ASCQ	Description
08	03	LOGICAL UNIT COMMUNICATION CRC ERROR (ULTRA-DMA /32)
08	04	UNREACHABLE COPY TARGET
09	00	TRACK FOLLOWING ERROR
09	04	HEAD SELECT FAULT
0A	00	ERROR LOG OVERFLOW
0B	00	WARNING
0B	01	WARNING - SPECIFIED TEMPERATURE EXCEEDED
0B	02	WARNING - ENCLOSURE DEGRADED
0B	03	WARNING - BACKGROUND SELF-TEST FAILED
0B	04	WARNING - BACKGROUND PRE-SCAN DETECTED MEDIUM ERROR
0B	05	WARNING - BACKGROUND MEDIUM SCAN DETECTED MEDIUM ERROR
0B	06	WARNING - NON-VOLATILE CACHE NOW VOLATILE
0B	07	WARNING - DEGRADED POWER TO NON-VOLATILE CACHE
0B	08	WARNING - POWER LOSS EXPECTED
0C	02	WRITE ERROR - AUTO REALLOCATION FAILED
0C	03	WRITE ERROR - RECOMMEND REASSIGNMENT
0C	04	COMPRESSION CHECK MISCOMPARE ERROR
0C	05	DATA EXPANSION OCCURRED DURING COMPRESSION
0C	06	BLOCK NOT COMPRESSIBLE
0C	0B	AUXILIARY MEMORY WRITE ERROR
0C	0C	WRITE ERROR - UNEXPECTED UNSOLICITED DATA
0C	0D	WRITE ERROR - NOT ENOUGH UNSOLICITED DATA
0D	00	ERROR DETECTED BY THIRD PARTY TEMPORARY INITIATOR
0D	01	THIRD PARTY DEVICE FAILURE
0D	02	COPY TARGET DEVICE NOT REACHABLE
0D	03	INCORRECT COPY TARGET DEVICE TYPE
0D	04	COPY TARGET DEVICE DATA UNDERRUN
0D	05	COPY TARGET DEVICE DATA OVERRUN
0E	00	INVALID INFORMATION UNIT
0E	01	INFORMATION UNIT TOO SHORT
0E	02	INFORMATION UNIT TOO LONG
0E	03	INVALID FIELD IN COMMAND INFORMATION UNIT
10	00	ID CRC OR ECC ERROR
10	01	LOGICAL BLOCK GUARD CHECK FAILED
10	02	LOGICAL BLOCK APPLICATION TAG CHECK FAILED
10	03	LOGICAL BLOCK REFERENCE TAG CHECK FAILED
11	00	UNRECOVERED READ ERROR
11	01	READ RETRIES EXHAUSTED

Table 46. SMART ASC/ASCQ error codes and messages (continued)

ASC	ASCQ	Description
11	02	ERROR TOO LONG TO CORRECT
11	03	MULTIPLE READ ERRORS
11	04	UNRECOVERED READ ERROR - AUTO REALLOCATE FAILED
11	0A	MISCORRECTED ERROR
11	0B	UNRECOVERED READ ERROR - RECOMMEND REASSIGNMENT
11	0C	UNRECOVERED READ ERROR - RECOMMEND REWRITE THE DATA
11	0D	DE-COMPRESS CRC ERROR
11	0E	CANNOT DECOMPRESS USING DECLARED ALGORITHM
11	12	AUXILIARY MEMORY READ ERROR
11	13	READ ERROR - FAILED RETRANSMISSION REQUEST
11	14	READ ERROR - LBA MARKED BAD BY APPLICATION CLIENT
12	00	ADDRESS MARK NOT FOUND FOR ID FIELD
13	00	ADDRESS MARK NOT FOUND FOR DATA FIELD
14	00	RECORDED ENTITY NOT FOUND
14	01	RECORD NOT FOUND
14	05	RECORD NOT FOUND - RECOMMEND REASSIGNMENT
14	06	RECORD NOT FOUND - DATA AUTO-REALLOCATED
15	00	RANDOM POSITIONING ERROR
15	01	MECHANICAL POSITIONING ERROR
15	02	POSITIONING ERROR DETECTED BY READ OF MEDIUM
16	00	DATA SYNCHRONIZATION MARK ERROR
16	01	DATA SYNC ERROR - DATA REWRITTEN
16	02	DATA SYNC ERROR - RECOMMEND REWRITE
16	03	DATA SYNC ERROR - DATA AUTO-REALLOCATED
16	04	DATA SYNC ERROR - RECOMMEND REASSIGNMENT
17	00	RECOVERED DATA WITH NO ERROR CORRECTION APPLIED
17	01	RECOVERED DATA WITH RETRIES
17	02	RECOVERED DATA WITH POSITIVE HEAD OFFSET
17	03	RECOVERED DATA WITH NEGATIVE HEAD OFFSET
17	05	RECOVERED DATA USING PREVIOUS SECTOR ID
17	06	RECOVERED DATA WITHOUT ECC - DATA AUTO-REALLOCATED
17	07	RECOVERED DATA WITHOUT ECC - RECOMMEND REASSIGNMENT
17	08	RECOVERED DATA WITHOUT ECC - RECOMMEND REWRITE
17	09	RECOVERED DATA WITHOUT ECC - DATA REWRITTEN
18	00	RECOVERED DATA WITH ERROR CORRECTION APPLIED
18	01	RECOVERED DATA WITH ERROR CORR. & RETRIES APPLIED
18	02	RECOVERED DATA - DATA AUTO-REALLOCATED
18	05	RECOVERED DATA - RECOMMEND REASSIGNMENT

Table 46. SMART ASC/ASCQ error codes and messages (continued)

ASC	ASCQ	Description
18	06	RECOVERED DATA - RECOMMEND REWRITE
18	07	RECOVERED DATA WITH ECC - DATA REWRITTEN
19	00	DEFECT LIST ERROR
19	01	DEFECT LIST NOT AVAILABLE
19	02	DEFECT LIST ERROR IN PRIMARY LIST
19	03	DEFECT LIST ERROR IN GROWN LIST
1A	00	PARAMETER LIST LENGTH ERROR
1B	00	SYNCHRONOUS DATA TRANSFER ERROR
1C	00	DEFECT LIST NOT FOUND
1C	01	PRIMARY DEFECT LIST NOT FOUND
1C	02	GROWN DEFECT LIST NOT FOUND
1D	00	MISCOMPARE DURING VERIFY OPERATION
1D	01	MISCOMPARE VERIFY OF UNMAPPED LBA
1E	00	RECOVERED ID WITH ECC CORRECTION
1F	00	PARTIAL DEFECT LIST TRANSFER
20	00	INVALID COMMAND OPERATION CODE
20	01	ACCESS DENIED - INITIATOR PENDING-ENROLLED
20	02	ACCESS DENIED - NO ACCESS RIGHTS
20	03	ACCESS DENIED - INVALID MGMT ID KEY
20	08	ACCESS DENIED - ENROLLMENT CONFLICT
20	09	ACCESS DENIED - INVALID LU IDENTIFIER
20	0A	ACCESS DENIED - INVALID PROXY TOKEN
20	0B	ACCESS DENIED - ACL LUN CONFLICT
21	00	LOGICAL BLOCK ADDRESS OUT OF RANGE
21	01	INVALID ELEMENT ADDRESS
22	00	ILLEGAL FUNCTION (USE 20 00, 24 00, OR 26 00)
24	00	INVALID FIELD IN CDB
24	01	CDB DECRYPTION ERROR
24	08	INVALID XCDB
25	00	LOGICAL UNIT NOT SUPPORTED
26	00	INVALID FIELD IN PARAMETER LIST
26	01	PARAMETER NOT SUPPORTED
26	02	PARAMETER VALUE INVALID
26	03	THRESHOLD PARAMETERS NOT SUPPORTED
26	04	INVALID RELEASE OF PERSISTENT RESERVATION
26	05	DATA DECRYPTION ERROR
26	06	TOO MANY TARGET DESCRIPTORS
26	07	UNSUPPORTED TARGET DESCRIPTOR TYPE CODE
26	08	TOO MANY SEGMENT DESCRIPTORS
26	09	UNSUPPORTED SEGMENT DESCRIPTOR TYPE CODE

Table 46. SMART ASC/ASCQ error codes and messages (continued)

ASC	ASCQ	Description
26	0A	UNEXPECTED INEXACT SEGMENT
26	0B	INLINE DATA LENGTH EXCEEDED
26	0C	INVALID OPERATION FOR COPY SOURCE OR DESTINATION
26	0D	COPY SEGMENT GRANULARITY VIOLATION
26	0E	INVALID PARAMETER WHILE PORT IS ENABLED
27	00	WRITE PROTECTED
27	01	HARDWARE WRITE PROTECTED
27	02	LOGICAL UNIT SOFTWARE WRITE PROTECTED
27	07	SPACE ALLOCATION FAILED WRITE PROTECT
28	00	NOT READY TO READY CHANGE, MEDIUM MAY HAVE CHANGED
28	01	IMPORT OR EXPORT ELEMENT ACCESSED
29	00	POWER ON, RESET, OR BUS DEVICE RESET OCCURRED
29	01	POWER ON OCCURRED
29	02	SCSI BUS RESET OCCURRED
29	03	BUS DEVICE RESET FUNCTION OCCURRED
29	04	DEVICE INTERNAL RESET
29	05	TRANSCIEVER MODE CHANGED TO SINGLE-ENDED
29	06	TRANSCIEVER MODE CHANGED TO LVD
29	07	I_T NEXUS LOSS OCCURRED
2A	00	PARAMETERS CHANGED
2A	01	MODE PARAMETERS CHANGED
2A	02	LOG PARAMETERS CHANGED
2A	03	RESERVATIONS PREEMPTED
2A	04	RESERVATIONS RELEASED
2A	05	REGISTRATIONS PREEMPTED
2A	06	ASYMMETRIC ACCESS STATE CHANGED
2A	07	IMPLICIT ASYMMETRIC ACCESS STATE TRANSITION FAILED
2A	08	PRIORITY CHANGED
2A	09	CAPACITY DATA HAS CHANGED
2A	0A	ERROR HISTORY I_T NEXUS CLEARED
2A	0B	ERROR HISTORY SNAPSHOT RELEASED
2A	10	TIMESTAMP CHANGED
2A	14	SA CREATION CAPABILITIES DATA HAS CHANGED
2B	00	COPY CANNOT EXECUTE SINCE HOST CANNOT DISCONNECT
2C	00	COMMAND SEQUENCE ERROR
2C	05	ILLEGAL POWER CONDITION REQUEST
2C	07	PREVIOUS BUSY STATUS
2C	08	PREVIOUS TASK SET FULL STATUS
2C	09	PREVIOUS RESERVATION CONFLICT STATUS

Table 46. SMART ASC/ASCQ error codes and messages (continued)

ASC	ASCQ	Description
2C	0C	ORWRITE GENERATION DOES NOT MATCH
2F	00	COMMANDS CLEARED BY ANOTHER INITIATOR
2F	01	COMMANDS CLEARED BY POWER LOSS NOTIFICATION
2F	02	COMMANDS CLEARED BY DEVICE SERVER
30	00	INCOMPATIBLE MEDIUM INSTALLED
30	01	CANNOT READ MEDIUM - UNKNOWN FORMAT
30	02	CANNOT READ MEDIUM - INCOMPATIBLE FORMAT
30	03	CLEANING CARTRIDGE INSTALLED
30	04	CANNOT WRITE MEDIUM - UNKNOWN FORMAT
30	05	CANNOT WRITE MEDIUM - INCOMPATIBLE FORMAT
30	06	CANNOT FORMAT MEDIUM - INCOMPATIBLE MEDIUM
30	07	CLEANING FAILURE
30	0A	CLEANING REQUEST REJECTED
31	00	MEDIUM FORMAT CORRUPTED
31	01	FORMAT COMMAND FAILED
32	00	NO DEFECT SPARE LOCATION AVAILABLE
32	01	DEFECT LIST UPDATE FAILURE
34	00	ENCLOSURE FAILURE
35	00	ENCLOSURE SERVICES FAILURE
35	01	UNSUPPORTED ENCLOSURE FUNCTION
35	02	ENCLOSURE SERVICES UNAVAILABLE
35	03	ENCLOSURE SERVICES TRANSFER FAILURE
35	04	ENCLOSURE SERVICES TRANSFER REFUSED
35	05	ENCLOSURE SERVICES CHECKSUM ERROR
37	00	ROUNDED PARAMETER
38	07	THIN PROVISIONING SOFT THRESHOLD REACHED
39	00	SAVING PARAMETERS NOT SUPPORTED
3A	00	MEDIUM NOT PRESENT
3A	01	MEDIUM NOT PRESENT - TRAY CLOSED
3A	02	MEDIUM NOT PRESENT - TRAY OPEN
3A	03	MEDIUM NOT PRESENT - LOADABLE
3A	04	MEDIUM NOT PRESENT - MEDIUM AUXILIARY MEMORY ACCESSIBLE
3B	0D	MEDIUM DESTINATION ELEMENT FULL
3B	0E	MEDIUM SOURCE ELEMENT EMPTY
3B	11	MEDIUM MAGAZINE NOT ACCESSIBLE
3B	12	MEDIUM MAGAZINE REMOVED
3B	13	MEDIUM MAGAZINE INSERTED
3B	14	MEDIUM MAGAZINE LOCKED
3B	15	MEDIUM MAGAZINE UNLOCKED

Table 46. SMART ASC/ASCQ error codes and messages (continued)

ASC	ASCQ	Description
3D	00	INVALID BITS IN IDENTIFY MESSAGE
3E	00	LOGICAL UNIT HAS NOT SELF-CONFIGURED YET
3E	01	LOGICAL UNIT FAILURE
3E	02	TIMEOUT ON LOGICAL UNIT
3E	03	LOGICAL UNIT FAILED SELF-TEST
3E	04	LOGICAL UNIT UNABLE TO UPDATE SELF-TEST LOG
3F	00	TARGET OPERATING CONDITIONS HAVE CHANGED
3F	01	MICROCODE HAS BEEN CHANGED
3F	02	CHANGED OPERATING DEFINITION
3F	03	INQUIRY DATA HAS CHANGED
3F	04	COMPONENT DEVICE ATTACHED
3F	05	DEVICE IDENTIFIER CHANGED
3F	06	REDUNDANCY GROUP CREATED OR MODIFIED
3F	07	REDUNDANCY GROUP DELETED
3F	08	SPARE CREATED OR MODIFIED
3F	09	SPARE DELETED
3F	0A	VOLUME SET CREATED OR MODIFIED
3F	0B	VOLUME SET DELETED
3F	0C	VOLUME SET DEASSIGNED
3F	0D	VOLUME SET REASSIGNED
3F	0E	REPORTED LUNS DATA HAS CHANGED
3F	0F	ECHO BUFFER OVERWRITTEN
3F	10	MEDIUM LOADABLE
3F	11	MEDIUM AUXILIARY MEMORY ACCESSIBLE
3F	12	ISCSI IP ADDRESS ADDED
3F	13	ISCSI IP ADDRESS REMOVED
3F	14	ISCSI IP ADDRESS CHANGED
40	00	RAM FAILURE
40	NN	DIAGNOSTIC FAILURE ON COMPONENT NN
41	00	DATA PATH FAILURE
42	00	POWER-ON OR SELF-TEST FAILURE
43	00	MESSAGE ERROR
44	00	INTERNAL TARGET FAILURE
44	71	ATA DEVICE FAILED SET FEATURES
45	00	SELECT OR RESELECT FAILURE
46	00	UNSUCCESSFUL SOFT RESET
47	00	SCSI PARITY ERROR
47	01	DATA PHASE CRC ERROR DETECTED
47	02	SCSI PARITY ERROR DETECTED DURING ST DATA PHASE
47	03	INFORMATION UNIT IUCRC ERROR DETECTED

Table 46. SMART ASC/ASCQ error codes and messages (continued)

ASC	ASCQ	Description
47	04	ASYNCHRONOUS INFORMATION PROTECTION ERROR DETECTED
47	05	PROTOCOL SERVICE CRC ERROR
47	06	PHY TEST FUNCTION IN PROGRESS
47	7F	SOME COMMANDS CLEARED BY ISCSI PROTOCOL EVENT
48	00	INITIATOR DETECTED ERROR MESSAGE RECEIVED
49	00	INVALID MESSAGE ERROR
4A	00	COMMAND PHASE ERROR
4B	00	DATA PHASE ERROR
4B	01	INVALID TARGET PORT TRANSFER TAG RECEIVED
4B	02	TOO MUCH WRITE DATA
4B	03	ACK/NAK TIMEOUT
4B	04	NAK RECEIVED
4B	05	DATA OFFSET ERROR
4B	06	INITIATOR RESPONSE TIMEOUT
4B	07	CONNECTION LOST
4C	00	LOGICAL UNIT FAILED SELF-CONFIGURATION
4D	NN	TAGGED OVERLAPPED COMMANDS (NN = TASK TAG)
4E	00	OVERLAPPED COMMANDS ATTEMPTED
53	00	MEDIA LOAD OR EJECT FAILED
53	02	MEDIUM REMOVAL PREVENTED
55	01	SYSTEM BUFFER FULL
55	02	INSUFFICIENT RESERVATION RESOURCES
55	03	INSUFFICIENT RESOURCES
55	04	INSUFFICIENT REGISTRATION RESOURCES
55	05	INSUFFICIENT ACCESS CONTROL RESOURCES
55	06	AUXILIARY MEMORY OUT OF SPACE
55	0B	INSUFFICIENT POWER FOR OPERATION
5A	00	OPERATOR REQUEST OR STATE CHANGE INPUT
5A	01	OPERATOR MEDIUM REMOVAL REQUEST
5A	02	OPERATOR SELECTED WRITE PROTECT
5A	03	OPERATOR SELECTED WRITE PERMIT
5B	00	LOG EXCEPTION
5B	01	THRESHOLD CONDITION MET
5B	02	LOG COUNTER AT MAXIMUM
5B	03	LOG LIST CODES EXHAUSTED
5C	00	RPL STATUS CHANGE
5C	01	SPINDLES SYNCHRONIZED
5C	02	SPINDLES NOT SYNCHRONIZED
5D	00	FAILURE PREDICTION THRESHOLD EXCEEDED
5D	05	HARDWARE IMPENDING FAILURE HARD DISK DRIVE ERROR

Table 46. SMART ASC/ASCQ error codes and messages (continued)

ASC	ASCQ	Description
5D	10	HARDWARE IMPENDING FAILURE GENERAL HARD DRIVE FAILURE
5D	11	HARDWARE IMPENDING FAILURE DRIVE ERROR RATE TOO HIGH
5D	12	HARDWARE IMPENDING FAILURE DATA ERROR RATE TOO HIGH
5D	13	HARDWARE IMPENDING FAILURE SEEK ERROR RATE TOO HIGH
5D	14	HARDWARE IMPENDING FAILURE TOO MANY BLOCK REASSIGNS
5D	15	HARDWARE IMPENDING FAILURE ACCESS TIMES TOO HIGH
5D	16	HARDWARE IMPENDING FAILURE START UNIT TIMES TOO HIGH
5D	17	HARDWARE IMPENDING FAILURE CHANNEL PARAMETRICS
5D	18	HARDWARE IMPENDING FAILURE CONTROLLER DETECTED
5D	19	HARDWARE IMPENDING FAILURE THROUGHPUT PERFORMANCE
5D	1A	HARDWARE IMPENDING FAILURE SEEK TIME PERFORMANCE
5D	1B	HARDWARE IMPENDING FAILURE SPIN-UP RETRY COUNT
5D	1C	HARDWARE IMPENDING FAILURE DRIVE CALIBRATION RETRY COUNT
5D	20	CONTROLLER IMPENDING FAILURE GENERAL HARD DRIVE FAILURE
5D	21	CONTROLLER IMPENDING FAILURE DRIVE ERROR RATE TOO HIGH
5D	22	CONTROLLER IMPENDING FAILURE DATA ERROR RATE TOO HIGH
5D	23	CONTROLLER IMPENDING FAILURE SEEK ERROR RATE TOO HIGH
5D	24	CONTROLLER IMPENDING FAILURE TOO MANY BLOCK REASSIGNS
5D	25	CONTROLLER IMPENDING FAILURE ACCESS TIMES TOO HIGH
5D	26	CONTROLLER IMPENDING FAILURE START UNIT TIMES TOO HIGH
5D	27	CONTROLLER IMPENDING FAILURE CHANNEL PARAMETRICS
5D	28	CONTROLLER IMPENDING FAILURE CONTROLLER DETECTED
5D	29	CONTROLLER IMPENDING FAILURE THROUGHPUT PERFORMANCE
5D	2A	CONTROLLER IMPENDING FAILURE SEEK TIME PERFORMANCE
5D	2B	CONTROLLER IMPENDING FAILURE SPIN-UP RETRY COUNT
5D	2C	CONTROLLER IMPENDING FAILURE DRIVE CALIBRATION RETRY COUNT
5D	30	DATA CHANNEL IMPENDING FAILURE GENERAL HARD DRIVE FAILURE
5D	31	DATA CHANNEL IMPENDING FAILURE DRIVE ERROR RATE TOO HIGH
5D	32	DATA CHANNEL IMPENDING FAILURE DATA ERROR RATE TOO HIGH
5D	33	DATA CHANNEL IMPENDING FAILURE SEEK ERROR RATE TOO HIGH

Table 46. SMART ASC/ASCQ error codes and messages (continued)

ASC	ASCQ	Description
5D	34	DATA CHANNEL IMPENDING FAILURE TOO MANY BLOCK REASSIGNS
5D	35	DATA CHANNEL IMPENDING FAILURE ACCESS TIMES TOO HIGH
5D	36	DATA CHANNEL IMPENDING FAILURE START UNIT TIMES TOO HIGH
5D	37	DATA CHANNEL IMPENDING FAILURE CHANNEL PARAMETRICS
5D	38	DATA CHANNEL IMPENDING FAILURE CONTROLLER DETECTED
5D	39	DATA CHANNEL IMPENDING FAILURE THROUGHPUT PERFORMANCE
5D	3A	DATA CHANNEL IMPENDING FAILURE SEEK TIME PERFORMANCE
5D	3B	DATA CHANNEL IMPENDING FAILURE SPIN-UP RETRY COUNT
5D	3C	DATA CHANNEL IMPENDING FAILURE DRIVE CALIBRATION RETRY COUNT
5D	40	SERVO IMPENDING FAILURE GENERAL HARD DRIVE FAILURE
5D	41	SERVO IMPENDING FAILURE DRIVE ERROR RATE TOO HIGH
5D	42	SERVO IMPENDING FAILURE DATA ERROR RATE TOO HIGH
5D	43	SERVO IMPENDING FAILURE SEEK ERROR RATE TOO HIGH
5D	44	SERVO IMPENDING FAILURE TOO MANY BLOCK REASSIGNS
5D	45	SERVO IMPENDING FAILURE ACCESS TIMES TOO HIGH
5D	46	SERVO IMPENDING FAILURE START UNIT TIMES TOO HIGH
5D	47	SERVO IMPENDING FAILURE CHANNEL PARAMETRICS
5D	48	SERVO IMPENDING FAILURE CONTROLLER DETECTED
5D	49	SERVO IMPENDING FAILURE THROUGHPUT PERFORMANCE
5D	4A	SERVO IMPENDING FAILURE SEEK TIME PERFORMANCE
5D	4B	SERVO IMPENDING FAILURE SPIN-UP RETRY COUNT
5D	4C	SERVO IMPENDING FAILURE DRIVE CALIBRATION RETRY COUNT
5D	50	SPINDLE IMPENDING FAILURE GENERAL HARD DRIVE FAILURE
5D	51	SPINDLE IMPENDING FAILURE DRIVE ERROR RATE TOO HIGH
5D	52	SPINDLE IMPENDING FAILURE DATA ERROR RATE TOO HIGH
5D	53	SPINDLE IMPENDING FAILURE SEEK ERROR RATE TOO HIGH
5D	54	SPINDLE IMPENDING FAILURE TOO MANY BLOCK REASSIGNS
5D	55	SPINDLE IMPENDING FAILURE ACCESS TIMES TOO HIGH
5D	56	SPINDLE IMPENDING FAILURE START UNIT TIMES TOO HIGH
5D	57	SPINDLE IMPENDING FAILURE CHANNEL PARAMETRICS
5D	58	SPINDLE IMPENDING FAILURE CONTROLLER DETECTED
5D	59	SPINDLE IMPENDING FAILURE THROUGHPUT PERFORMANCE
5D	5A	SPINDLE IMPENDING FAILURE SEEK TIME PERFORMANCE
5D	5B	SPINDLE IMPENDING FAILURE SPIN-UP RETRY COUNT
5D	5C	SPINDLE IMPENDING FAILURE DRIVE CALIBRATION RETRY COUNT

Table 46. SMART ASC/ASCQ error codes and messages (continued)

ASC	ASCQ	Description
5D	60	FIRMWARE IMPENDING FAILURE GENERAL HARD DRIVE FAILURE
5D	61	FIRMWARE IMPENDING FAILURE DRIVE ERROR RATE TOO HIGH
5D	62	FIRMWARE IMPENDING FAILURE DATA ERROR RATE TOO HIGH
5D	63	FIRMWARE IMPENDING FAILURE SEEK ERROR RATE TOO HIGH
5D	64	FIRMWARE IMPENDING FAILURE TOO MANY BLOCK REASSIGNS
5D	65	FIRMWARE IMPENDING FAILURE ACCESS TIMES TOO HIGH
5D	66	FIRMWARE IMPENDING FAILURE START UNIT TIMES TOO HIGH
5D	67	FIRMWARE IMPENDING FAILURE CHANNEL PARAMETRICS
5D	68	FIRMWARE IMPENDING FAILURE CONTROLLER DETECTED
5D	69	FIRMWARE IMPENDING FAILURE THROUGHPUT PERFORMANCE
5D	6A	FIRMWARE IMPENDING FAILURE SEEK TIME PERFORMANCE
5D	6B	FIRMWARE IMPENDING FAILURE SPIN-UP RETRY COUNT
5D	6C	FIRMWARE IMPENDING FAILURE DRIVE CALIBRATION RETRY COUNT
5D	FF	FAILURE PREDICTION THRESHOLD EXCEEDED (FALSE)
5E	00	LOW POWER CONDITION ON
5E	01	IDLE CONDITION ACTIVATED BY TIMER
5E	02	STANDBY CONDITION ACTIVATED BY TIMER
5E	03	IDLE CONDITION ACTIVATED BY COMMAND
5E	04	STANDBY CONDITION ACTIVATED BY COMMAND
5E	05	IDLE_B CONDITION ACTIVATED BY TIMER
5E	06	IDLE_B CONDITION ACTIVATED BY COMMAND
5E	07	IDLE_C CONDITION ACTIVATED BY TIMER
5E	08	IDLE_C CONDITION ACTIVATED BY COMMAND
5E	09	STANDBY_Y CONDITION ACTIVATED BY TIMER
5E	0A	STANDBY_Y CONDITION ACTIVATED BY COMMAND
65	00	VOLTAGE FAULT
67	0A	SET TARGET PORT GROUPS COMMAND FAILED
67	0B	ATA DEVICE FEATURE NOT ENABLED
74	08	DIGITAL SIGNATURE VALIDATION FAILURE
74	0C	UNABLE TO DECRYPT PARAMETER LIST
74	10	SA CREATION PARAMETER VALUE INVALID
74	11	SA CREATION PARAMETER VALUE REJECTED
74	12	INVALID SA USAGE
74	30	SA CREATION PARAMETER NOT SUPPORTED
74	40	AUTHENTICATION FAILED
74	71	LOGICAL UNIT ACCESS NOT AUTHORIZED
74	79	SECURITY CONFLICT IN TRANSLATED DEVICE

Monitoring memory usage on a file module

Use this procedure to monitor memory usage on a file module.

Procedure

1. Log in to the file module and issue the command `lsperfdata -g memory_free_usage -t hour -n <node> | tail`.
2. If the file module shows diminishing memory and is reaching full capacity, initiate a file module reboot. See “Shut down or reboot a file module or clustered system” in the *IBM Storwize V7000 Unified Information Center*.

Errors and messages

System errors and messages can be triggered by conditions that range from simple typing errors to problems with system devices or programs.

About this task

Refer to the following topics for information about errors and messages.

Example

Note: For reference to or repair-information about non-Storwize V7000 Unified components, refer to the user documentation provided with those components.

Understanding error codes

The Storwize V7000 Unified error codes convey specific information in an alphanumeric sequence.

Tip: Search for error codes or event IDs by using EFS on the front. For 66012FC, for example, search on EFS66012FC. For a broader range of results, use a wildcard at the end and shorten the search appropriately. For example, search on EFS66012* or EFS660*, and so on.

Error code information

The following tables show the error code elements: ACDDDDx and provide information on what the various elements represent.

Table 47. Error code information.

Listing the code element information in the sequence of ACDDDDx.

Code element	Information
A	Originating role information
C	Originating hardware or software code
DDDD	Specific error code
x	Severity of the error code

Originating device information

The alphanumeric symbol or code in the A position indicates the originating device.

Table 48. Originating role information.

Listing devices for A in sequence ACDDDD.

A = Originating role information in sequence ACDDDD	
Code	Device
0/1	Management node error codes
2/3	File Module role error codes
4/5	Storage node role error codes
6	Storage node role error codes
8	Ethernet switch error codes.

Originating specific hardware and software codes

The alphanumeric symbol in the C position represents the originating specific hardware and software code.

- For the originating file module and file module specific hardware code (code 0, 2, 4), go to Table 49.
- For the originating file module specific software code (code 1, 3, 5), go to Table 50 on page 119.
- For the storage enclosure hardware code (code 6), go to Table 51 on page 119.
- For the Ethernet switches (code 8): The Ethernet switches are a single field replaceable unit (FRU) and have no unique failing hardware code. The Ethernet switches use 0 for the originating specific hardware or software code.

Table 49. Originating file module and file module specific hardware code – Code 0, 2, 4.

Listing devices for variable C in the specific hardware code sequence of ABBCDDDD.

C = Originating specific hardware code in sequence ABBCDDDD	
Code	Device
0	System x hardware (CPU, memory, powers supplies, etc.)
1	Built-in Ethernet port 0
2	Built-in Ethernet port 1
3	Built-in Ethernet port 2
4	Built-in Ethernet port 3
5	Optional Ethernet port 4 (Dual Port 10G card)
6	Optional Ethernet port 5 (Dual Port 10G card)
7	Optional Ethernet port 6 (Dual Port 10G card)
8	Optional Ethernet port 7 (Dual Port 10G card)
B	Fibre channel adapter 1 (both ports) – Storage node only
C	Fibre channel adapter 2 (both ports) – Storage node only
D	Bonded device (data0 mgmt0)
E	System x internal hard disk drives

Table 50. Originating file module specific software code – Code 1, 3, 5.

Listing devices for variable C in the specific software code sequence of ABBCDDDD.

C = Originating specific software code in sequence ABBCDDDD	
Code	Device
0	Red Hat Linux
1	GPFS
2	CIFS server
3	CTDB
4	SoFS
5	winbind
6	multipathd
7	nscd
8	sshd
9	httpd
A	vsftpd
B	nmbd
C	nfsd
D	cpu
E	multipath/disk

Table 51. Storage enclosure hardware code – Code 6.

Listing devices for variable C in the specific hardware code sequence of ABBCDDDD.

C = Originating specific software code in sequence ABBCDDDD	
Code	Device
0	Generic value for storage enclosure hardware
1	Disk drive in controller drawer
2	RAID controller card 0
3	RAID controller card 1
4	Power supply in controller drawer
5	RAID array/LUN issue in controller drawer
6	Disk drive in expansion drawer
7	Expansion fibre channel card 0
8	Expansion fibre channel card 1
9	Power supply in expansion drawer
A	RAID array/LUN issue in expansion drawer

Severity of the error

The element *x* indicates the severity of the error. The value *x* can be:

- **A for Action:** GUI error messages. The user must perform a specific action.
- **C for Critical:** A critical error occurred which must be corrected by the user or system administrator.

- **D for Debug:** Used only for debug purposes.
- **I for Informational:** No operation action required.
- **W for Warning:** An error occurred that should be investigated and fixed.

Error code example

The following error code example illustrates how to interpret the alphanumeric elements based on the information provided above.

Error code and message:

4E0013C – Controller cache discarded due to firmware version incompatibility.

The following table shows the break down of the error code's alphanumeric elements:

Table 52. Error code break down.

This identifies the variables of 4 E 0 nnn x in the sequence of ACDXXXx.

ACDXXXx	
4E0013C	
4	File Module
E	System x internal hard disk drives
0	Originated with system checkout
nnn	Unique error code
x	Severity of the error

Understanding event IDs

The Storwize V7000 Unified messages follow a specific format, which is detailed here.

About this task

Tip: Search for error codes or event IDs by using EFS on the front. For 66012FC, for example, search on EFS66012FC. For a broader range of results, use a wildcard at the end and shorten the search appropriately. For example, search on EFS66012* or EFS660*, and so on.

The format of system messages is *cnnnnx*. The elements—*cnnnnx*—represent the following information:

- The element *c* is an alphabetic identifier assigned to a component. The message component identifiers are assigned as follows:
 - A** for Common or Access Layer
 - B** for Space
 - C** for GPFS
 - D** for Wizards
 - F** for Statistics
 - G** for CLI
 - H** for Health Center

I for Asynchronous Replication

J for SCM

L for HSM

AK for NDMP

- The element *nnnn* is a 4 digit message number
- The element *x* indicates the severity of the error. The value *x* can be:
 - A for Action:** GUI error messages. The user must perform a specific action.
 - C for Critical:** A critical error occurred which must be corrected by the user or system administrator.
 - D for Debug:** Used only for debug purposes.
 - I for Informational:** No operation action required.
 - W for Warning:** An error occurred that can cause problems in the future. The problem should be investigated and fixed.

File module hardware problems

This section helps you to identify and resolve file module hardware problems.

Removing and replacing parts for the 2073-720

About this task

Illustrations in this section might differ slightly from the actual hardware.

Table 53. Components identified as customer replaceable units (CRUs) and field replaceable units (FRUs)

Types of replaceable parts	Explanation of each type of replaceable part	Procedures categorized under each type of replaceable part
Tier 1 Customer replaceable units (CRUs)	<p>Tier 1 CRUs are your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation.</p> <p>Service agreements can be purchased so that you can ask IBM to replace these.</p>	<p>"Removing the cover" on page 123</p> <p>"Installing the cover" on page 124</p> <p>"Removing the bezel" on page 125</p> <p>"Installing the bezel" on page 126</p> <p>"Disk drive cable connections" on page 129</p> <p>"Removing the battery" on page 131</p> <p>"Installing the battery" on page 135</p> <p>"Removing the air baffle" on page 137</p> <p>"Installing the air baffle" on page 138</p> <p>"Removing the fan bracket" on page 139</p> <p>"Installing the fan bracket" on page 141</p> <p>"Removing a PCI riser-card assembly" on page 142</p> <p>"Installing a PCI riser-card assembly" on page 142</p> <p>"Removing a PCI adapter from a PCI riser-card assembly" on page 143</p> <p>"Installing a PCI adapter in a PCI riser-card assembly" on page 145</p> <p>"Removing a Fibre Channel PCI adapter" on page 146</p> <p>"Installing a Fibre Channel PCI adapter" on page 146</p> <p>"Removing a 10-Gbps Ethernet adapter" on page 146</p> <p>"Installing a 10-Gbps Ethernet adapter" on page 147</p> <p>"Removing a hot-swap hard disk drive" on page 148</p> <p>"Installing a hot-swap hard disk drive" on page 149</p> <p>"Removing the DVD drive" on page 150</p> <p>"Installing the DVD drive" on page 152</p> <p>"Removing a memory module" on page 152</p> <p>"Installing a memory module" on page 153</p> <p>"Removing a hot-swap fan" on page 156</p> <p>"Installing a hot-swap fan" on page 157</p> <p>"Removing a hot-swap ac power supply" on page 158</p> <p>"Installing a hot-swap ac power supply" on page 159</p> <p>"Removing the operator information panel assembly" on page 162</p> <p>"Installing the operator information panel assembly" on page 163</p> <p>"Removing the hot-swap drive backplane" on page 164</p> <p>"Installing the hot-swap drive backplane" on page 165</p> <p>"Removing the 240 VA safety cover" on page 127</p>

Table 53. Components identified as customer replaceable units (CRUs) and field replaceable units (FRUs) (continued)

Types of replaceable parts	Explanation of each type of replaceable part	Procedures categorized under each type of replaceable part
Field replaceable units (FRUs)	FRUs must be installed only by trained service technicians.	<p>"Installing the 240 VA safety cover" on page 128</p> <p>"Removing a microprocessor and heat sink" on page 166</p> <p>"Installing a microprocessor and heat sink" on page 170</p> <p>"Removing and replacing the thermal grease" on page 174</p> <p>"Removing a heat-sink retention module" on page 176</p> <p>"Installing a heat-sink retention module" on page 176</p> <p>"Removing the system board" on page 177</p> <p>"Installing the system board" on page 179</p> <p>"Setting the machine serial number" on page 181</p>

Removing the cover

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

To remove the cover, complete the following steps.

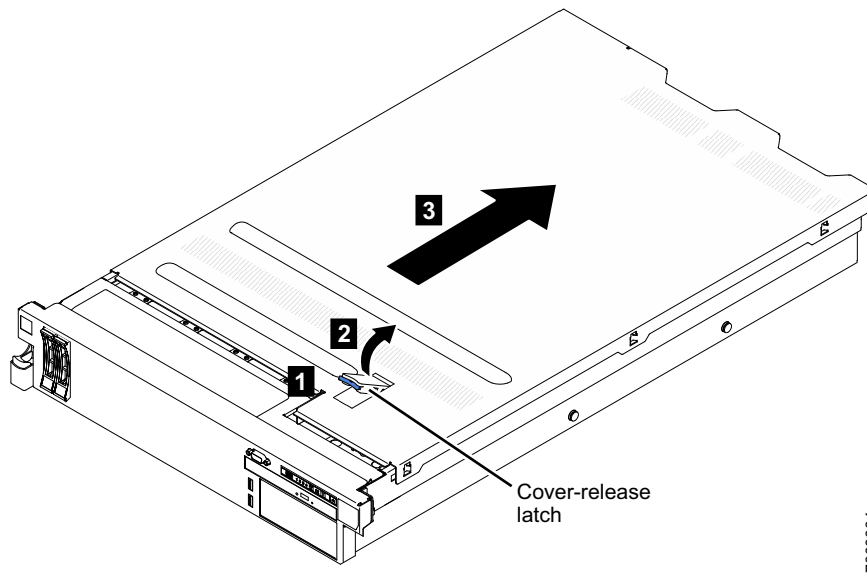


Figure 50. Removing the cover

Procedure

1. Read the safety information that begins on page Safety and "Installation guidelines" on page 94.

2. If you are planning to view the error LEDs that are on the system board and components, leave the file module connected to power and go directly to step 4.
3. If you are planning to install or remove a microprocessor, memory module, PCI adapter, battery, or other non-hot-swap optional device, follow the procedure in “Removing a file module and disconnecting power” on page 92 to suspend the file module from the cluster and shut it down and disconnect all power cords and external cables.
4. Press down on the left and right side latches and pull the file module out of the rack enclosure until both slide rails lock.

Note: You can reach the cables on the back of the file module when the file module is in the locked position.

5. Push the cover-release latch back **1**, then lift it up **2**.
6. Slide the cover back **3**, then lift the cover off the file module and set it aside.

Attention: For proper cooling and airflow, replace the cover before you turn on the file module. Operating the file module for extended periods of time (over 30 minutes) with the cover removed might damage the file module components.

7. If you are instructed to return the cover, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Installing the cover

The following procedure is for a Tier 1 customer replaceable unit (CRU).

Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

To install the cover, complete the following steps.

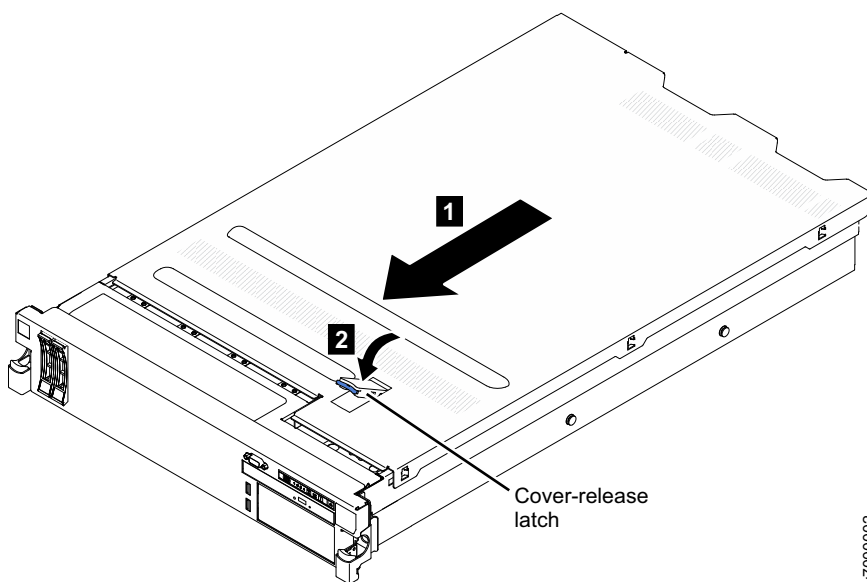


Figure 51. Installing the cover

Procedure

1. Make sure that all cables, adapters, and other components are installed and seated correctly and that there are no loose tools or parts inside the file module. Also, ensure that all internal cables are correctly routed.

Important: Before you slide the cover forward, make sure that all the tabs on the front, rear, and side of the cover engage the chassis correctly. If all the tabs do not engage the chassis correctly, it can cause difficulty in removing the cover later on.

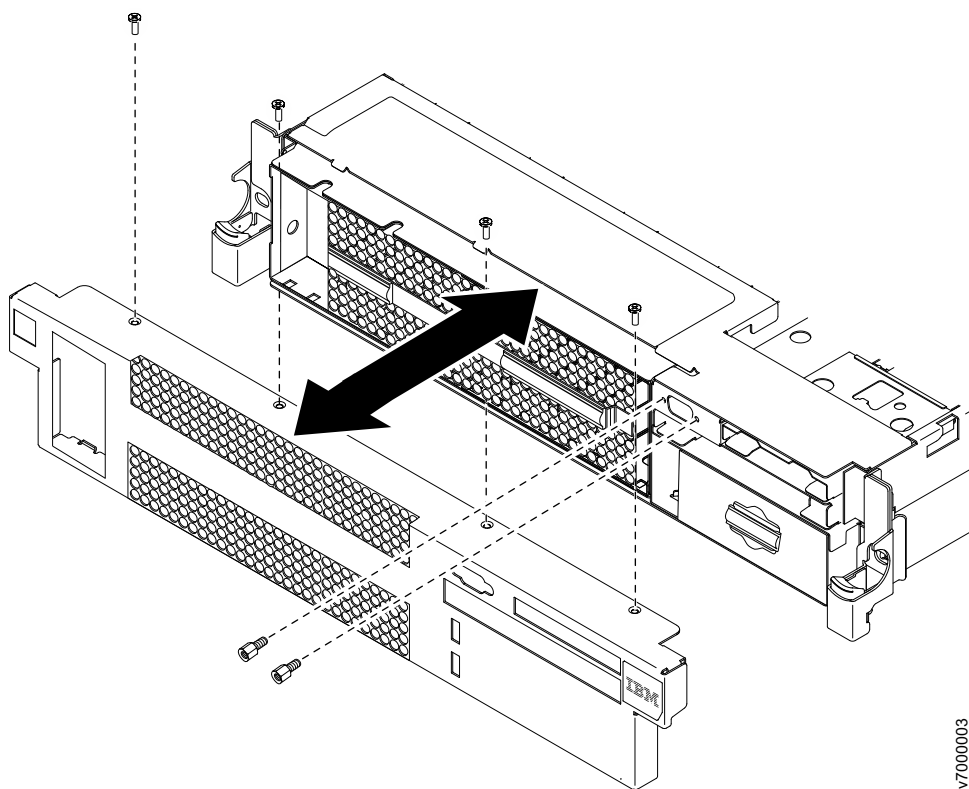
2. Place the cover-release latch in the open (up) position.
3. Insert the bottom tabs of the top cover into the matching slots in the file module chassis.
4. Press down on the cover-release latch to lock the cover in place.
5. Slide the file module into the rack until it latches.
6. Follow the steps at the end of the procedure in “Removing a file module and disconnecting power” on page 92 to suspend the file module to reconnect the file module and resume its use in the cluster.

Removing the bezel

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

To remove the bezel, complete the following steps.



v7000003

Figure 52. Removing the bezel

Procedure

1. Read the safety information that begins on page Safety and "Installation guidelines" on page 94.
2. Remove all the cables that are connected to the front of the file module.
3. Remove the cable retention bolts from the VGA port.
4. Remove the screws from the bezel.
5. Rotate the top of the bezel away from the file module.

Installing the bezel

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

To install the bezel, complete the following steps.

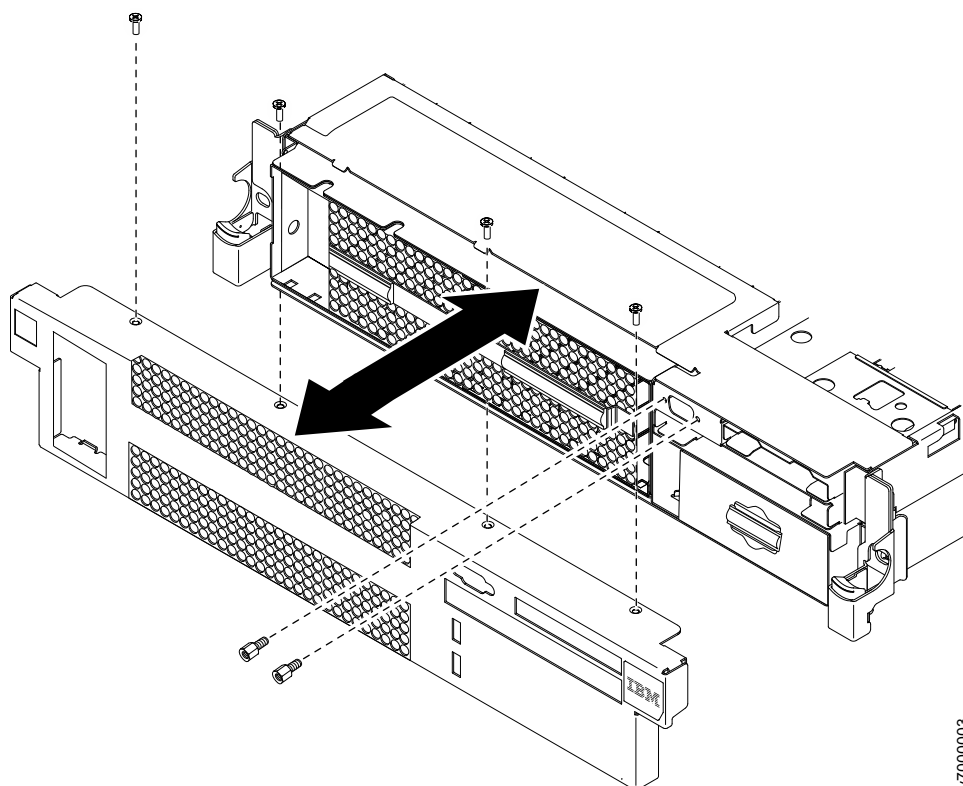


Figure 53. Installing the bezel

Procedure

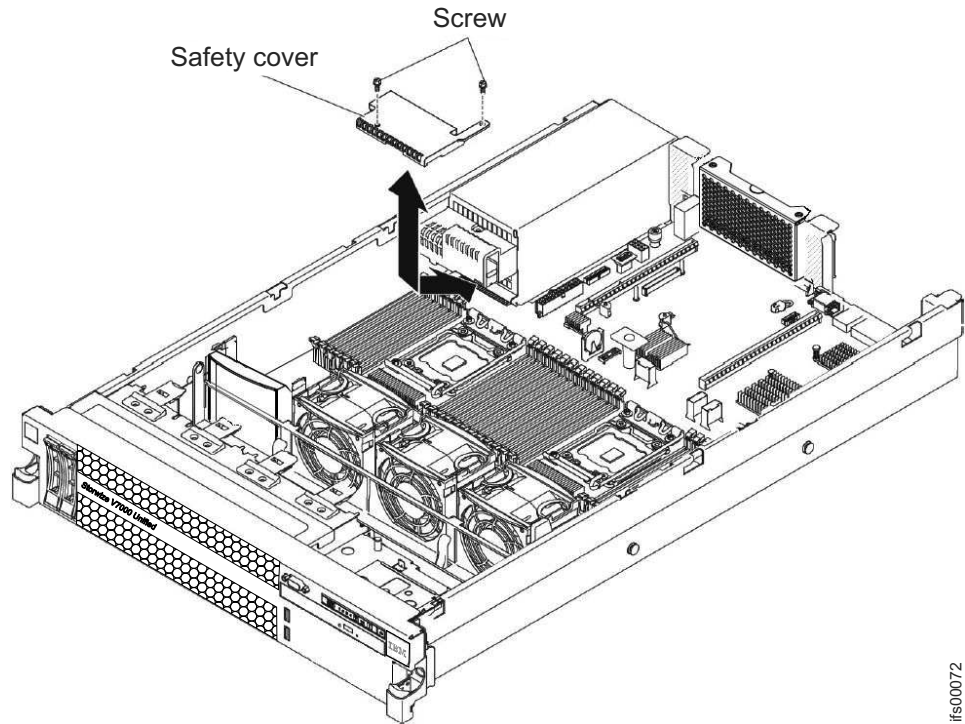
1. Insert the tabs on the bottom of the bezel into the slots on the underside of the chassis and attach it with the screws.
2. Connect any cables that you previously removed from the front of the file module.

Removing the 240 VA safety cover

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

To remove the 240 VA safety cover, complete the following steps:

1. Read the Safety information and "Installation guidelines" on page 94.
2. Follow the procedure in "Removing a file module and disconnecting power" on page 92 to suspend the file module from the cluster and shut it down, and then disconnect all power cords and external cables.
3. Pull the file module out of the rack.
4. Remove the file module cover (see "Removing the cover" on page 123).



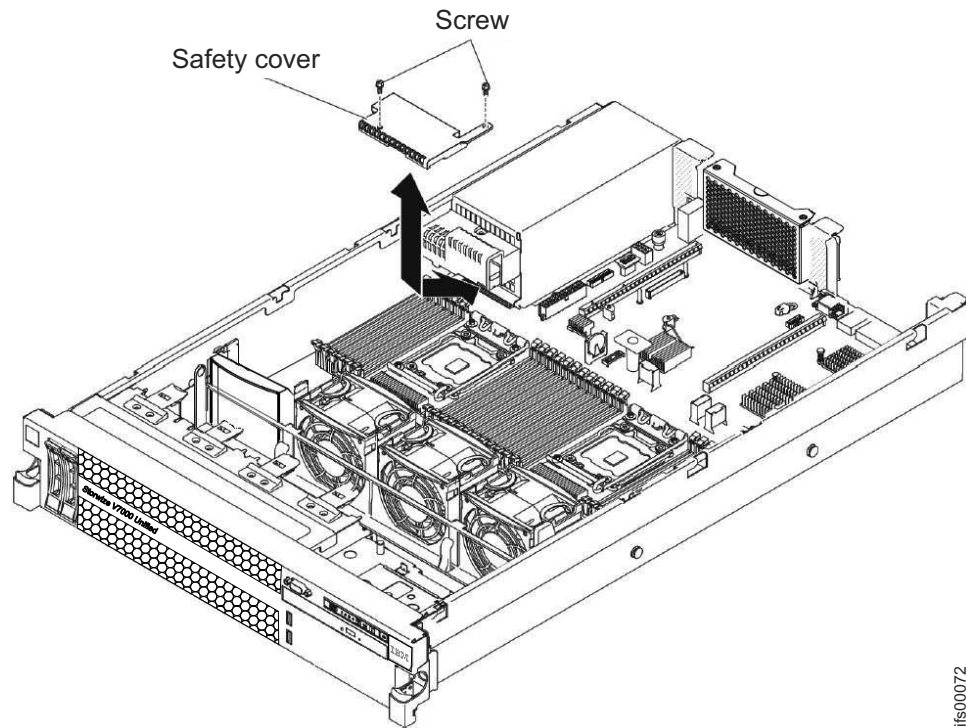
ifs00072

5. Remove the screw from the safety cover.
6. Disconnect the hard disk drive backplane power cables from the connector in front of the safety cover.
7. Slide the cover forward to disengage it from the system board, and then lift it out of the file module.
8. If you are instructed to return the 240 VA safety cover, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Installing the 240 VA safety cover

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

To install the 240 VA safety cover, complete the following steps.



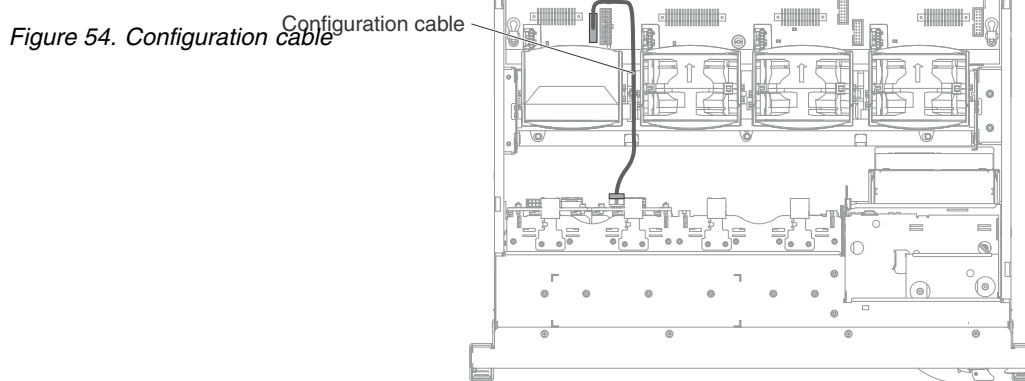
ifs00072

1. Line up and insert the tabs on the bottom of the safety cover into the slots on the system board.
2. Slide the safety cover toward the back of the file module until it is secure.
3. Connect the hard disk drive backplane power cables to the connector in front of the safety cover.
4. Install the screw into the safety cover.
5. Install the file module cover (see "Installing the cover" on page 124).
6. Slide the file module into the rack.
7. Follow the steps at the end of the procedure "Removing a file module and disconnecting power" on page 92 to reconnect the file module and resume its use in the cluster.

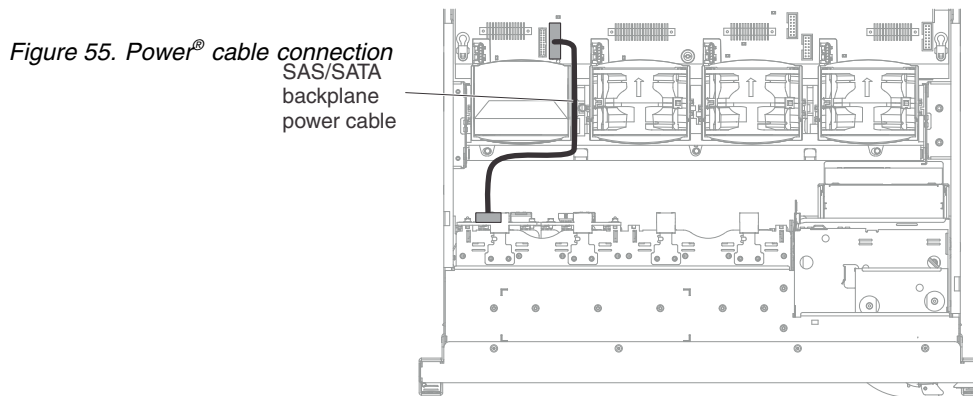
Disk drive cable connections

Use the information described here to know about the cabling structure for the 8 x 2.5-inch hot-swap drive bays.

The following illustration shows the cabling information for the configuration cable in the file module:



The following illustration shows the internal routing for the hard disk drive power cable.

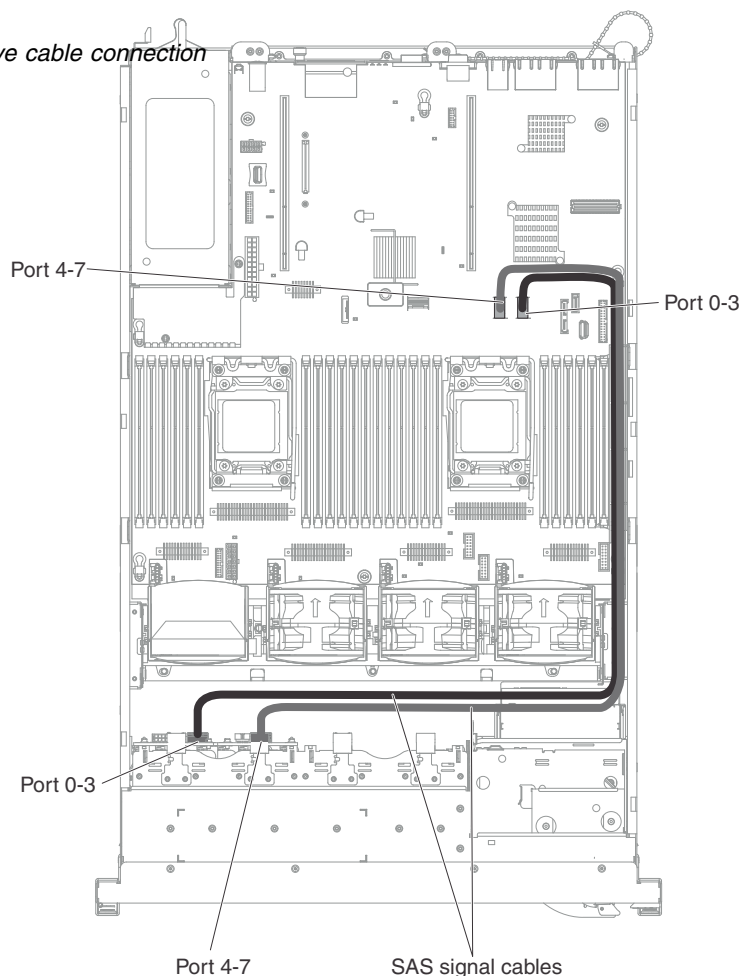


The following illustration shows the internal routing and connectors for the two SAS signal cables.

Note:

1. To connect the SAS signal cables, make sure that you first connect the signal cable, and then the power cable and configuration cable.
2. To disconnect the SAS signal cables, make sure that you first disconnect the power cable, and then the signal cable and configuration cable.

Figure 56. Hard disk drive cable connection



Removing the battery

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

Note: Before running a procedure, refer to “Removing a file module to perform a maintenance action” on page 91.

The following notes describe information that you must consider when replacing the battery:

- IBM has designed this product with your safety in mind. The lithium battery must be handled correctly to avoid possible danger. If you replace the battery, you must adhere to the following instructions.

Note: In the U. S., call 1-800-IBM-4333 for information about battery disposal.

- If you replace the original lithium battery with a heavy-metal battery or a battery with heavy-metal components, be aware of the following environmental consideration. Batteries and accumulators that contain heavy metals must not be

disposed of with normal domestic waste. They will be taken back free of charge by the manufacturer, distributor, or representative, to be recycled or disposed of in a proper manner.

- To order replacement batteries, call 1-800-IBM-SERV within the United States, and 1-800-465-7999 or 1-800-465-6666 within Canada. Outside the US and Canada, call your support center or IBM Business Partner.

Note: After you replace the battery, you must reconfigure the file module and reset the system date and time.

Statement 2



CAUTION:

When you are replacing the lithium battery, use only IBM Part Number 33F8354 or an equivalent type battery that is recommended by the manufacturer. If your system has a module that contains a lithium battery, replace it only with the same module type made by the same manufacturer. The battery contains lithium and can explode if not properly used, handled, or disposed of.

Do not:

- Throw or immerse into water
- Heat to more than 100°C (212°F)
- Repair or disassemble

Dispose of the battery as required by local ordinances or regulations.

To remove the battery, complete the following steps:

Procedure

1. Read the safety information that begins on page Safety and “Installation guidelines” on page 94.
2. Follow any special handling and installation instructions that come with the battery.
3. Follow the procedure in “Removing a file module and disconnecting power” on page 92 to suspend the file module from the cluster and shut it down and disconnect all power cords and external cables.
4. Slide the file module out of the rack.
5. Remove the cover. For more information, see Removing the cover.
6. Disconnect any internal cables, as necessary.
7. Locate the battery on the system board.
8. Remove the battery:
 - a. If there is a rubber cover on the battery holder, use your fingers to lift the battery cover from the battery connector.
 - b. Use one finger to push the battery horizontally away from the PCI riser card in slot 2 and out of its housing.

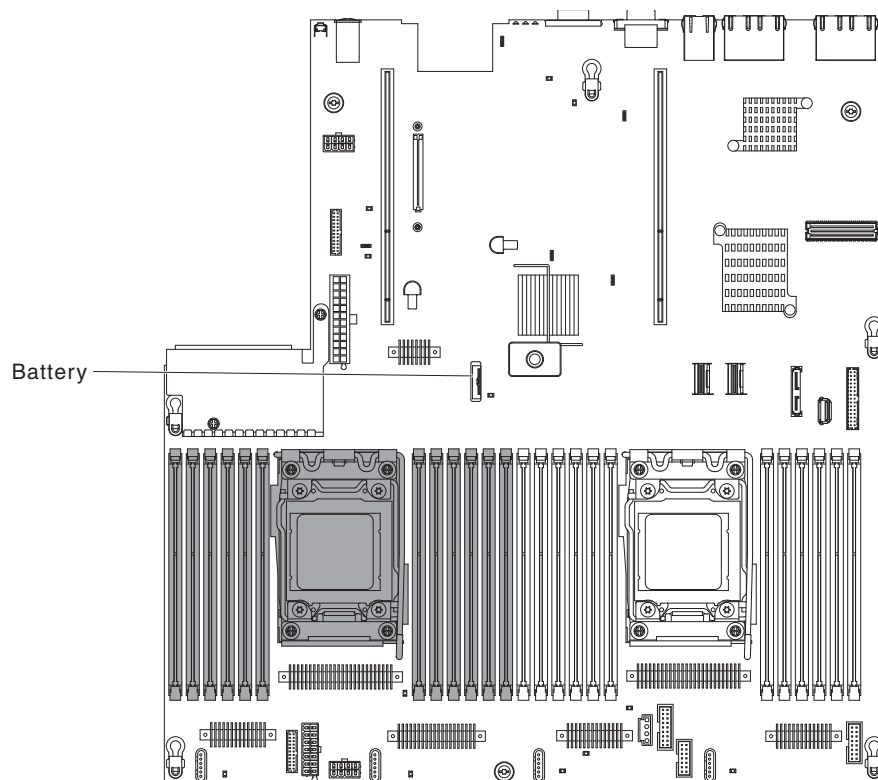


Figure 57. Removing the battery

Attention: You must not tilt or push the battery by using excessive force.
c. Use your thumb and index finger to lift the battery from the socket.

Attention: Do not lift the battery by using excessive force. Failing to remove the battery properly might damage the socket on the system board. Any damage to the socket might require replacing the system board.

9. Dispose of the battery as required by local ordinances or regulations. For more information, see the *IBM Environmental Notices and User's Guide* on the IBM Documentation CD.

Battery return program: This product may contain sealed lead acid, nickel cadmium, nickel metal hydride, lithium, or lithium ion battery. Consult your user manual or service manual for specific battery information. The battery must be recycled or disposed of properly. Recycling facilities may not be available in your area. For information on disposal of batteries outside the United States, go to www.ibm.com/ibm/environment/products/index.shtml or contact your local waste disposal facility.

In the United States, IBM has established a return process for reuse, recycling, or proper disposal of used IBM sealed lead acid, nickel cadmium, nickel metal hydride, and other battery packs from IBM Equipment. For information on proper disposal of these batteries, contact IBM at 1-800-426-4333. Please have the IBM part number listed on the battery available prior to your call.

The following applies for countries within the European Union:



For Taiwan:



Please recycle batteries.

廢電池請回收 SVC00066

Batteries or packaging for batteries are labeled in accordance with European Directive 2006/66/EC concerning batteries and accumulators and waste batteries and accumulators. The Directive determines the framework for the return and recycling of used batteries and accumulators as applicable throughout the European Union. This label is applied to various batteries to indicate that the battery is not to be thrown away, but rather reclaimed upon end of life per this Directive.

Les batteries ou emballages pour batteries sont étiquetés conformément aux directives européennes 2006/66/EC, norme relative aux batteries et accumulateurs en usage et aux batteries et accumulateurs usés. Les directives déterminent la marche à suivre en vigueur dans l'Union Européenne pour le retour et le recyclage des batteries et accumulateurs usés. Cette étiquette est appliquée sur diverses batteries pour indiquer que la batterie ne doit pas être mise au rebut mais plutôt récupérée en fin de cycle de vie selon cette norme.

バッテリーあるいはバッテリー用のパッケージには、EU 諸国に対する廃電気電子機器指令 2006/66/EC のラベルが貼られています。この指令は、バッテリーと蓄電池、および廃棄バッテリーと蓄電池に関するものです。この指令は、使用済みバッテリーと蓄電池の回収とリサイクルの骨子を定めているもので、EU 諸国にわたって適用されます。このラベルは、使用済みになったときに指令に従って適正な処理をする必要があることを知らせるために種々のバッテリーに貼られています。

In accordance with the European Directive 2006/66/EC, batteries and accumulators are labeled to indicate that they are to be collected separately and recycled at end of life. The label on the battery may also include a chemical symbol for the metal concerned in the battery (Pb for lead, Hg for mercury and Cd for cadmium). Users of batteries and accumulators must not dispose of batteries and accumulators as unsorted municipal waste, but use the collection framework available to customers for the return, recycling and treatment of batteries and accumulators. Customer participation is important to minimize any potential effects of batteries and

accumulators on the environment and human health due to the potential presence of hazardous substances. For proper collection and treatment, contact your local IBM representative.

Spain

This notice is provided in accordance with Royal Decree 106/2008 of Spain: The retail price of batteries, accumulators and power cells includes the cost of the environmental management of their waste.

Perchlorate Material - California

Special handling may apply. See <http://www.dtsc.ca.gov/hazardouswaste/perchlorate> for more information.

The foregoing notice is provided in accordance with California Code of Regulations Title 22, Division 4.5 Chapter 33. Best Management Practices for Perchlorate Materials. This product, part or both may include a lithium manganese dioxide battery which contains a perchlorate substance.

Installing the battery

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

The following notes describe information that you must consider when you replace the battery in the file module.

- You must replace the battery with a lithium battery of the same type from the same manufacturer.
- After you replace the battery, you must reconfigure the file module and reset the system date and time.
- To avoid possible danger, read and follow the following safety statement.
- To order replacement batteries, call 1-800-IBM-SERV within the United States, and 1-800-465-7999 or 1-800-465-6666 within Canada. Outside the US and Canada, call your support center or IBM Business Partner.

Statement 2



CAUTION:

When you are replacing the lithium battery, use only IBM Part Number 33F8354 or an equivalent type battery that is recommended by the manufacturer. If your system has a module that contains a lithium battery, replace it only with the same module type made by the same manufacturer. The battery contains lithium and can explode if not properly used, handled, or disposed of.

Do not:

- Throw or immerse into water
- Heat to more than 100°C (212°F)
- Repair or disassemble

Dispose of the battery as required by local ordinances or regulations.

For more information, see the *IBM Environmental Notices and User's Guide* on the IBM Documentation CD.

To install the replacement battery, complete the following steps:

Procedure

1. Follow any special handling and installation instructions that come with the replacement battery.
2. Insert the new battery:
 - a. Tilt the battery so that you can insert it into the socket on the side opposite the battery clip.

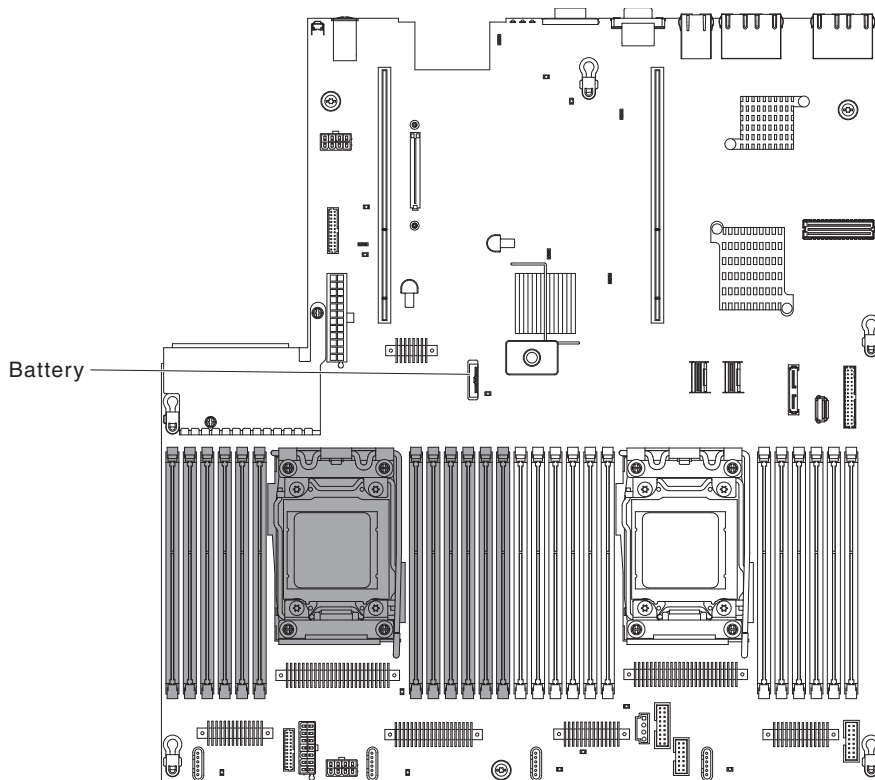


Figure 58. Installing the battery

- b. Press the battery down into the socket until it clicks into place. Make sure that the battery clip holds the battery securely.
 - c. If you removed a rubber cover from the battery holder, use your fingers to install the battery cover on top of the battery connector.
3. Reinstall any adapters that you removed.
4. Reconnect the internal cables that you disconnected.
5. Install the cover, as described in Installing the cover.
6. Slide the file module into the rack.
7. Follow the steps at the end of the procedure in “Removing a file module and disconnecting power” on page 92 to suspend the file module to reconnect the file module and resume its use in the cluster.

Note: You must wait approximately 2.5 minutes after you connect the power cord of the file module to an electrical outlet before the power-control button becomes active.

8. Start the Setup utility and reset the configuration.
 - Set the system date and time.
 - Set the power-on password.
 - Reconfigure the file module.

Removing the air baffle

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

When you work with some replaceable devices, you must first remove the DIMM air baffle to access certain components or connectors on the system board.

To remove the DIMM air baffle, complete the following steps.

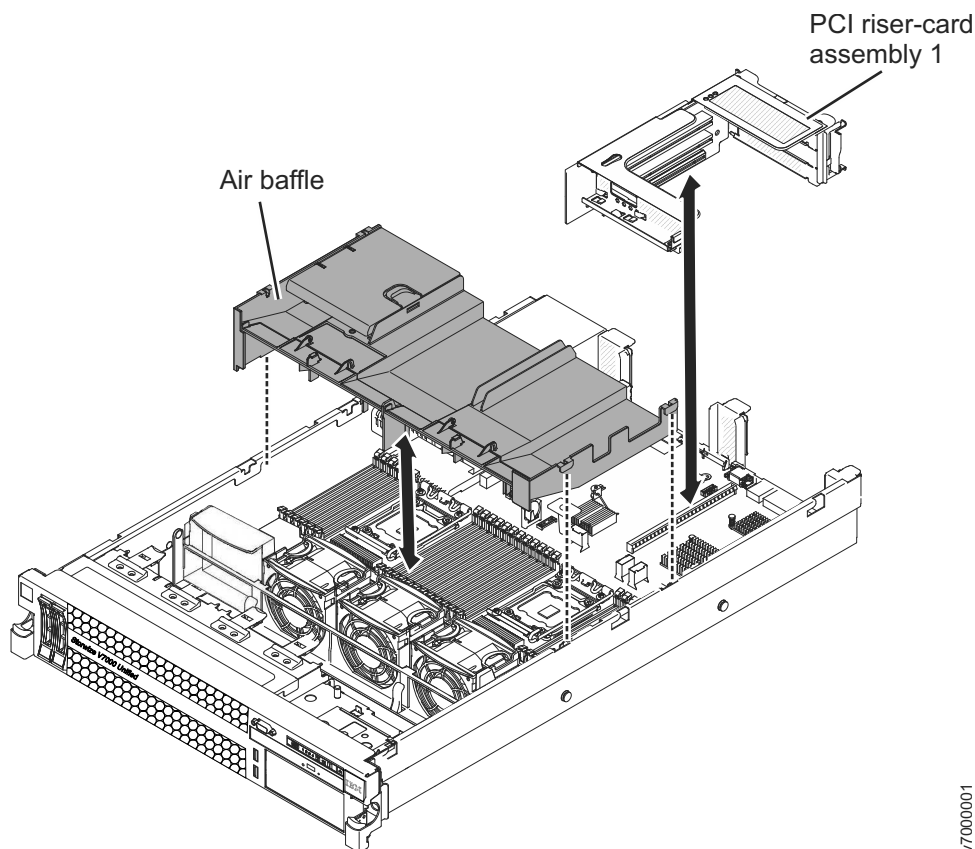


Figure 59. Removing the air baffle

Procedure

1. Read the safety information that begins on page Safety and “Installation guidelines” on page 94.
2. Follow the procedure in “Removing a file module and disconnecting power” on page 92 to suspend the file module from the cluster and shut it down and disconnect all power cords and external cables.
3. Remove the cover. For more information, see Removing the cover.
4. If there is any full-height, full-length card, remove riser-card assembly 1. For more information, see Removing a PCI riser-card assembly.
5. Place your fingers under the front and back of the top of the air baffle, and then lift the air baffle out of the file module.

Attention: For proper cooling and airflow, replace all the air baffles before you turn on the file module. If you operate the file module with any air baffle removed, it might cause damages to the file module components.

Installing the air baffle

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

To install the DIMM air baffle, complete the following steps.

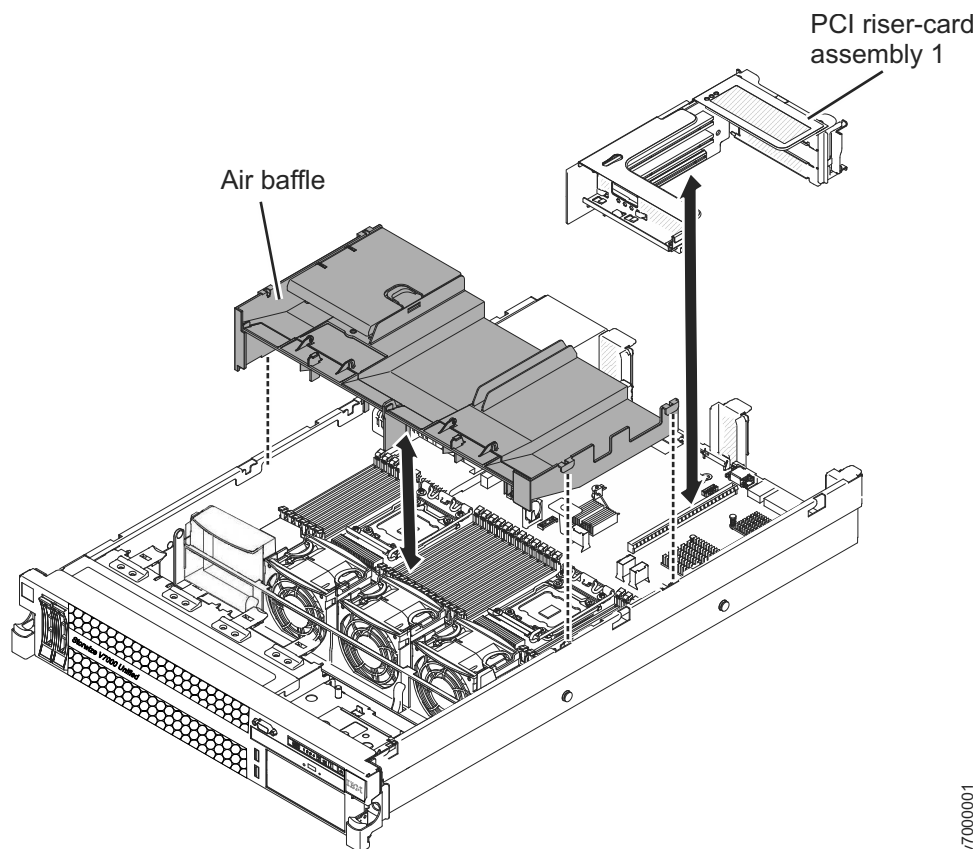


Figure 60. Installing the air baffle

Procedure

1. Read the safety information that begins on page Safety and “Installation guidelines” on page 94.
2. Align the air baffle pins with the two baffle pin slots on both sides of chassis.
3. Lower the air baffle into place, making sure that all cables are out of the way. Press the air baffle down until it is securely seated.

Note: Close the retaining clip on each end of the DIMM connector before you install the air baffle.

4. Install the cover. For more information, see Installing the cover.
5. Slide the file module into the rack.
6. Follow the steps at the end of the procedure in “Removing a file module and disconnecting power” on page 92 to suspend the file module to reconnect the file module and resume its use in the cluster.

Results

Attention: For proper cooling and airflow, replace all air baffles before you turn on the file module. Operating the file module with any air baffle removed might damage file module components.

Removing the fan bracket

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at

your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

To replace some components or to create working room, you might have to remove the fan-bracket assembly.

Note: To remove or install a fan, it is not necessary to remove the fan bracket.

To remove the fan bracket, complete the following steps.

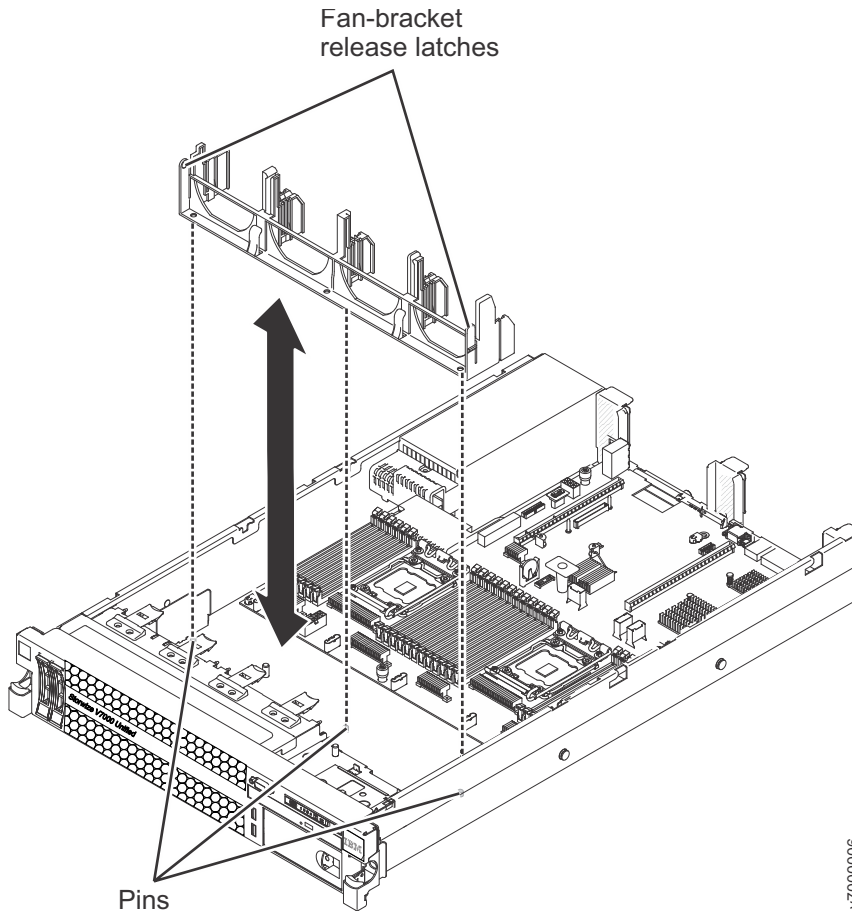


Figure 61. Removing the fan bracket

Procedure

1. Read the safety information that begins on page Safety and "Installation guidelines" on page 94.
2. Follow the procedure in "Removing a file module and disconnecting power" on page 92 to suspend the file module from the cluster and shut it down and disconnect all power cords and external cables.
3. Remove the cover. For more information, see Removing the cover.
4. Remove the fans.
5. Remove the PCI riser-card assemblies. For more information, see Removing a PCI riser-card assembly.

6. Press the fan-bracket release latches toward each other and lift the fan bracket out of the file module.

Installing the fan bracket

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

To install the fan bracket, complete the following steps.

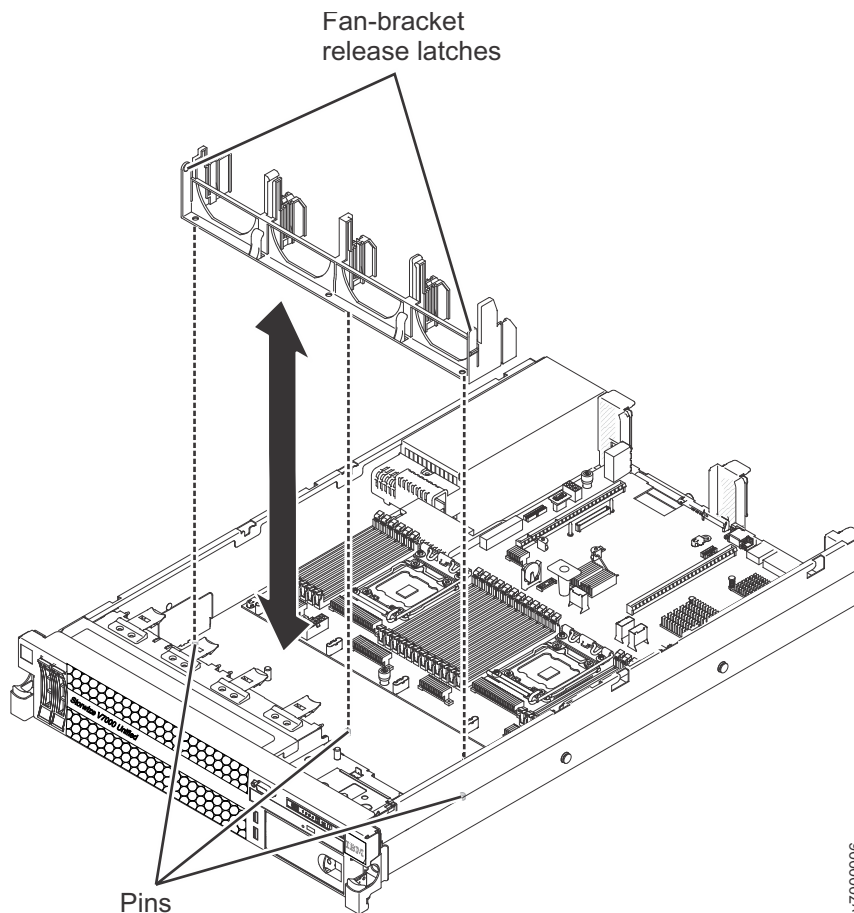


Figure 62. Installing the fan bracket

Procedure

1. Lower the fan bracket into the chassis.
2. Align the holes in the bottom of the bracket with the pins in the bottom of the chassis.
3. Press the bracket into position until the fan-bracket release levers click into place.
4. Replace the fans.
5. Replace the PCI riser-card assemblies. For more information, see [Installing a PCI riser-card assembly](#).
6. Install the cover. For more information, see [Installing the cover](#).

7. Slide the file module into the rack.
8. Follow the steps at the end of the procedure in “Removing a file module and disconnecting power” on page 92 to suspend the file module to reconnect the file module and resume its use in the cluster.

Removing a PCI riser-card assembly

The following procedure is for a Tier 1 customer replaceable unit (CRU).

Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

The file module comes with one riser-card assembly that contains two to three PCI slots. See <http://www.ibm.com/servers/eserver/serverproven/compat/us/> for a list of riser-card assemblies that you can use with the file module.

To remove a riser-card assembly, complete the following steps.

Procedure

1. Read the safety information that begins on page Safety and “Installation guidelines” on page 94.
2. Follow the procedure in “Removing a file module and disconnecting power” on page 92 to suspend the file module from the cluster and shut it down and disconnect all power cords and external cables.
3. Slide the file module out of the rack.
4. Remove the file module cover. For more information, see Removing the cover.
5. Grasp the riser-card assembly at the front tab and rear edge and lift it to remove it from the file module. Place the riser-card assembly on a flat, static-protective surface.

Installing a PCI riser-card assembly

IBM authorized service providers can install a PCI riser-card assembly in the file module. The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

The file module provides two PCI riser-card slots on the system board. The following information indicates the riser-card slots:

- The file module come with one PCI Express® riser-card assembly installed.
- A PCI Express riser-card assembly has a black connector and supports PCI Express adapters.
- PCI riser slot 1 (the farthest slot from the power supplies). You must install a PCI riser-card assembly in slot 1.
- PCI riser slot 2 (the closest slot to the power supplies). You must not install a PCI riser-card assembly in slot 2.

To install a riser-card assembly, complete the following steps.

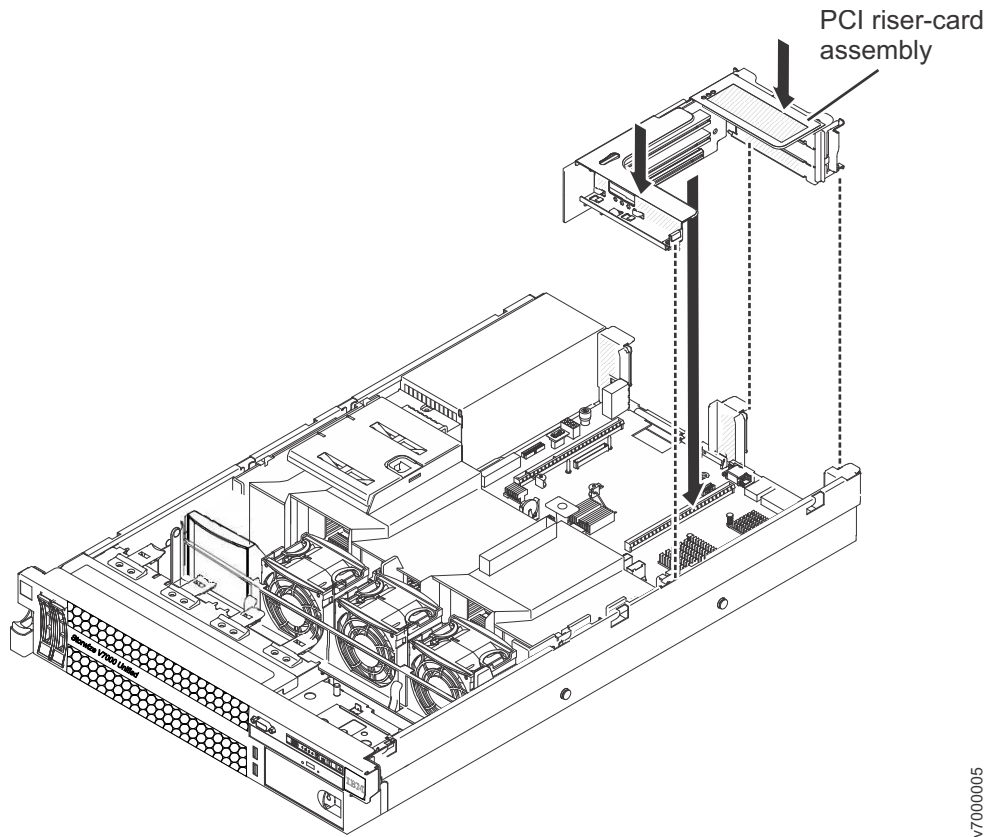


Figure 63. Installing a PCI riser-card assembly

Procedure

1. Reinstall any adapters.
2. Align the PCI riser-card assembly with the selected PCI connector on the system board:

Note: The chassis might sag after you remove the riser assembly. In this case, lift the bottom of the chassis to line up the slots on the side of the assembly to the alignment brackets in the side of the chassis.

- **PCI connector 1:** Carefully fit the two alignment slots on the side of the assembly onto the two alignment brackets in the side of the chassis.
3. Press down on the assembly. Make sure that the riser-card assembly is fully seated in the riser-card connector on the system board.
 4. Install the file module cover. For more information, see *Installing the cover*.
 5. Slide the file module into the rack.
 6. Follow the steps at the end of the procedure in “Removing a file module and disconnecting power” on page 92 to suspend the file module to reconnect the file module and resume its use in the cluster.

Removing a PCI adapter from a PCI riser-card assembly

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

This topic describes removing an adapter from a PCI expansion slot in a PCI riser-card assembly. These instructions apply to PCI adapters such as the Fibre Channel and the Ethernet network adapters.

To remove an adapter from a PCI expansion slot, complete the following steps.

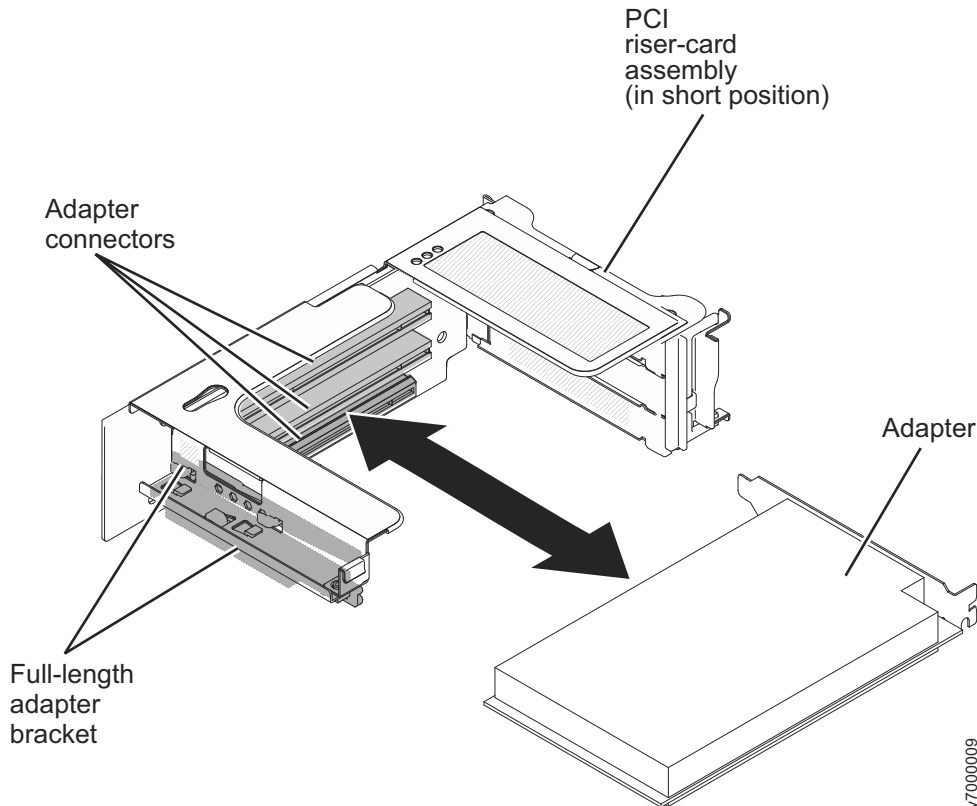


Figure 64. Removing a PCI adapter from a PCI riser-card assembly

Procedure

1. Read the safety information that begins on page Safety and "Installation guidelines" on page 94.
2. Follow the procedure in "Removing a file module and disconnecting power" on page 92 to suspend the file module from the cluster and shut it down and disconnect all power cords and external cables.
3. Press down on the left and right side latches and slide the file module out of the rack enclosure until both slide rails lock, and then remove the cover. For more information, see Removing the cover.
4. Remove the PCI riser-card assembly that contains the adapter, as described in Removing a PCI riser-card assembly.
5. Carefully grasp the adapter by its top edge or upper corners, and pull the adapter from the PCI expansion slot.
6. If you are instructed to return the adapter, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Installing a PCI adapter in a PCI riser-card assembly

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

Important: Some cluster solutions require specific code levels or coordinated code updates. If the device is part of a cluster solution, verify that the latest level of code is supported for the cluster solution before you update the code.

To install an adapter, complete the following steps.

Procedure

1. Install the adapter in the expansion slot.
 - a. Align the adapter with the PCI connector on the riser card and the guide on the external end of the riser-card assembly.
 - b. Press the adapter firmly into the PCI connector on the riser card.

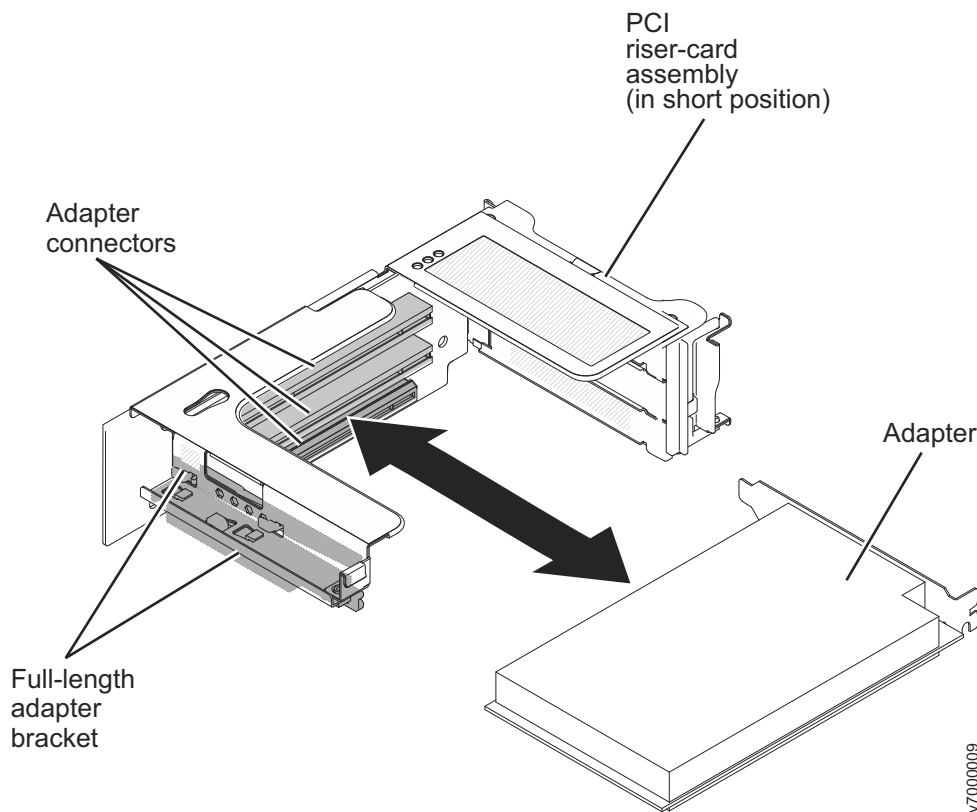


Figure 65. Inserting the adapter into the PCI connector

2. Align the PCI riser-card assembly with the selected PCI connector on the system board:
 - Carefully fit the two alignment slots on the side of the assembly onto the two alignment brackets on the side of the chassis; align the rear of the assembly with the guides on the rear of the file module.

3. Press down on the assembly. Make sure that the riser-card assembly is fully seated in the riser-card connector on the system board.
4. Install the file module cover. For more information, see *Installing the cover*.
5. Slide the file module into the rack.
6. Reconnect the external cables; then, reconnect the power cords and turn on the peripheral devices and the file module.

Removing a Fibre Channel PCI adapter

This removal instruction indicates the slot location for the Fibre Channel PCI adapter.

About this task

The Fibre Channel adapter is in PCI slot 2.

Refer to “Removing a PCI adapter from a PCI riser-card assembly” on page 143 for instructions.

Installing a Fibre Channel PCI adapter

This installation instruction indicates the slot location for the Fibre Channel PCI adapter.

About this task

The Fibre Channel adapter must go in PCI slot 2.

Refer to “Installing a PCI adapter in a PCI riser-card assembly” on page 145 for instructions.

Removing a 10-Gbps Ethernet adapter

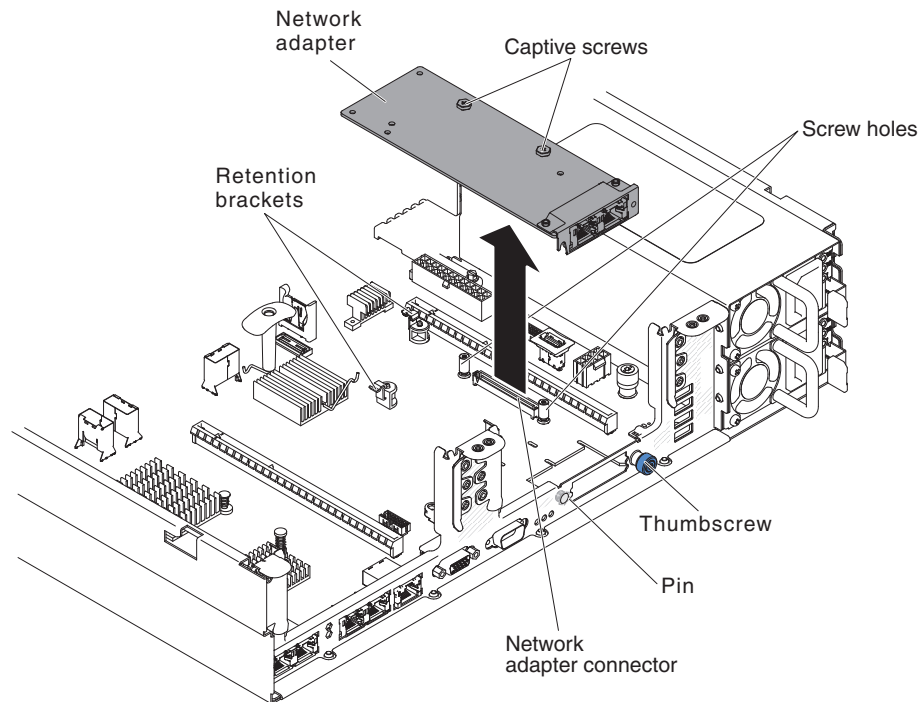
The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

To remove the network adapter, complete the following steps:

Procedure

1. Read the Safety information and “Installation guidelines” on page 94.
2. Follow the procedure in “Removing a file module and disconnecting power” on page 92 to suspend the file module from the cluster and shut it down, and then disconnect all power cords and external cables.
3. Remove the cover (see “Removing the cover” on page 123).
4. Loosen the thumbscrew on the rear of the chassis.



5. Grasp the network adapter and disengage it from the pin, standoffs, retention brackets, and the connector on the system board; then, lift the adapter out of the port openings on the rear of the chassis and remove it from the server.
6. If you are instructed to return the network adapter, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Installing a 10-Gbps Ethernet adapter

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

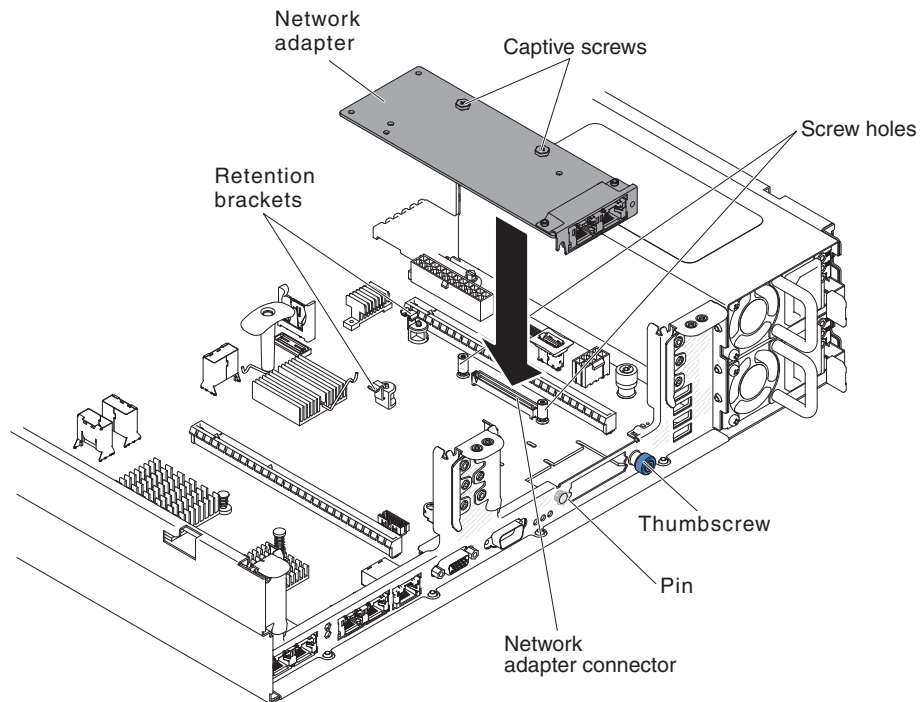
About this task

The file module has a Emulex dual port 10GbE SFP+ Embedded VFA III adapter.

To install the network adapter, complete the following steps:

Procedure

1. Read the Safety information and “Installation guidelines” on page 94.
2. Follow the procedure in “Removing a file module and disconnecting power” on page 92 to suspend the file module from the cluster and shut it down, and then disconnect all power cords and external cables.
3. Remove the cover (see “Removing the cover” on page 123).
4. Touch the static-protective package that contains the new adapter to any unpainted metal surface on the file module, and then remove the adapter from the package.
5. Align the adapter so that the port connectors on the adapter line up with the pin and thumbscrew on the chassis; then, align the connector of the adapter with the adapter connector on the system board.



6. Press the adapter firmly until the pin, standoffs, and retention brackets engage the adapter. Make sure the adapter is securely seated on the connector on the system board.

Attention: Make sure the port connectors on the adapter are aligned properly with the chassis on the rear of the server. An incorrectly seated adapter might cause damage to the system board or the adapter.

7. Fasten the thumbscrew.
8. Install the cover (see "Installing the cover" on page 124).
9. Slide the file module into the rack.
10. Follow the steps at the end of the procedure "Removing a file module and disconnecting power" on page 92 to reconnect the file module and resume its use in the cluster.

Removing a hot-swap hard disk drive

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

To remove a hard disk drive from a hot-swap bay, complete the following steps.

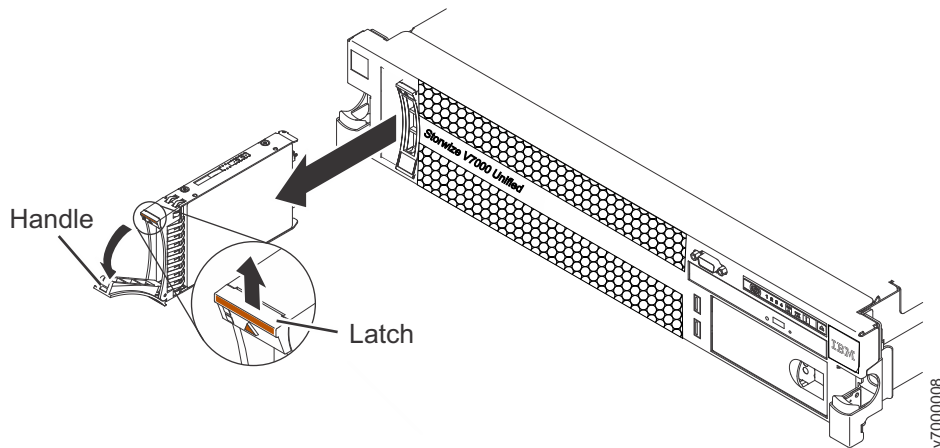


Figure 66. Removing a hot-swap hard disk drive

Attention: To maintain proper system cooling, do not operate the file module for more than 10 minutes without either a drive or a filler panel installed in each bay.

Procedure

1. Read the safety information that begins on page Safety, “Handling static-sensitive devices” on page 96, and “Installation guidelines” on page 94.
2. Press up on the release latch at the top of the drive front.
3. Rotate the handle on the drive downward to the open position.
4. Pull the hot-swap drive assembly out of the bay approximately 25 mm (1 inch). Wait approximately 45 seconds while the drive spins down before you remove the drive assembly completely from the bay.
5. If you are instructed to return the hot-swap drive, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Installing a hot-swap hard disk drive

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

Locate the documentation that comes with the hard disk drive and follow those instructions in addition to the instructions in this section.

To install a drive in a hot-swap bay, complete the following steps.

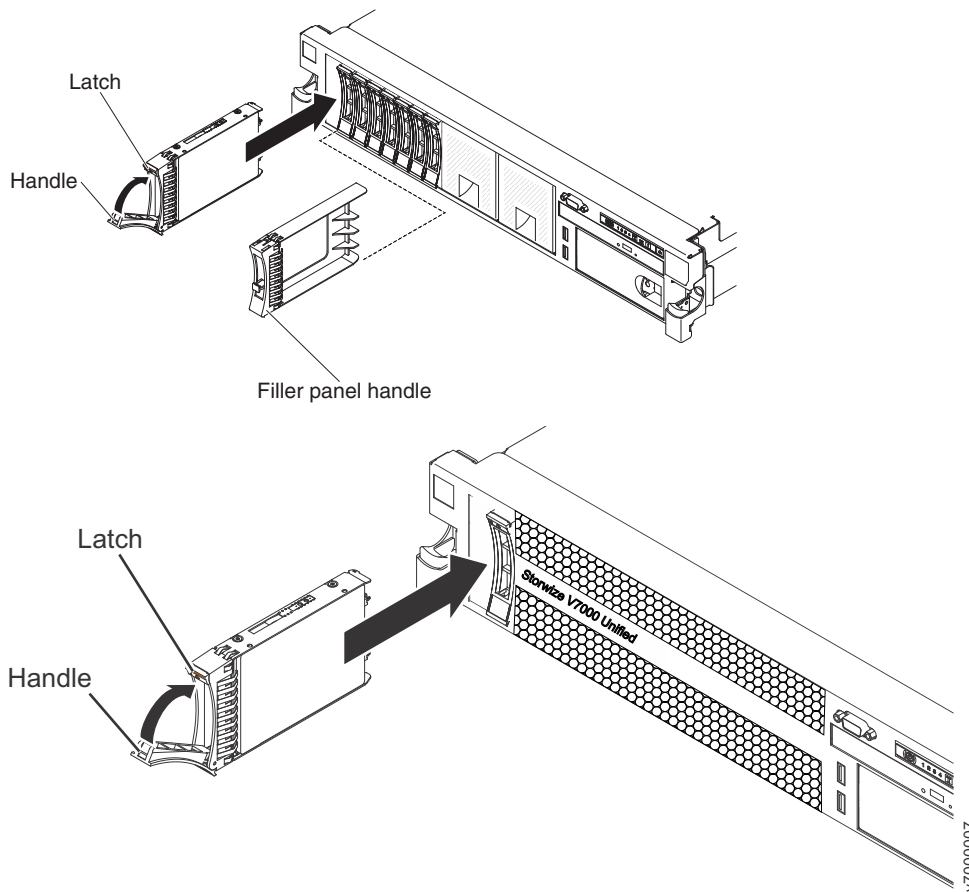


Figure 67. Installing a hot-swap hard disk drive

Attention: To maintain proper system cooling, do not operate the file module for more than 10 minutes without a drive that is installed in each bay.

Procedure

1. Orient the drive as shown in the illustration.
2. Make sure that the tray handle is open.
3. Align the drive assembly with the guide rails in the bay.
4. Gently push the drive assembly into the bay until the drive stops.
5. Push the tray handle to the closed (locked) position.
6. If the system is turned on, check the hard disk drive status LED to verify that the hard disk drive is operating correctly.

Results

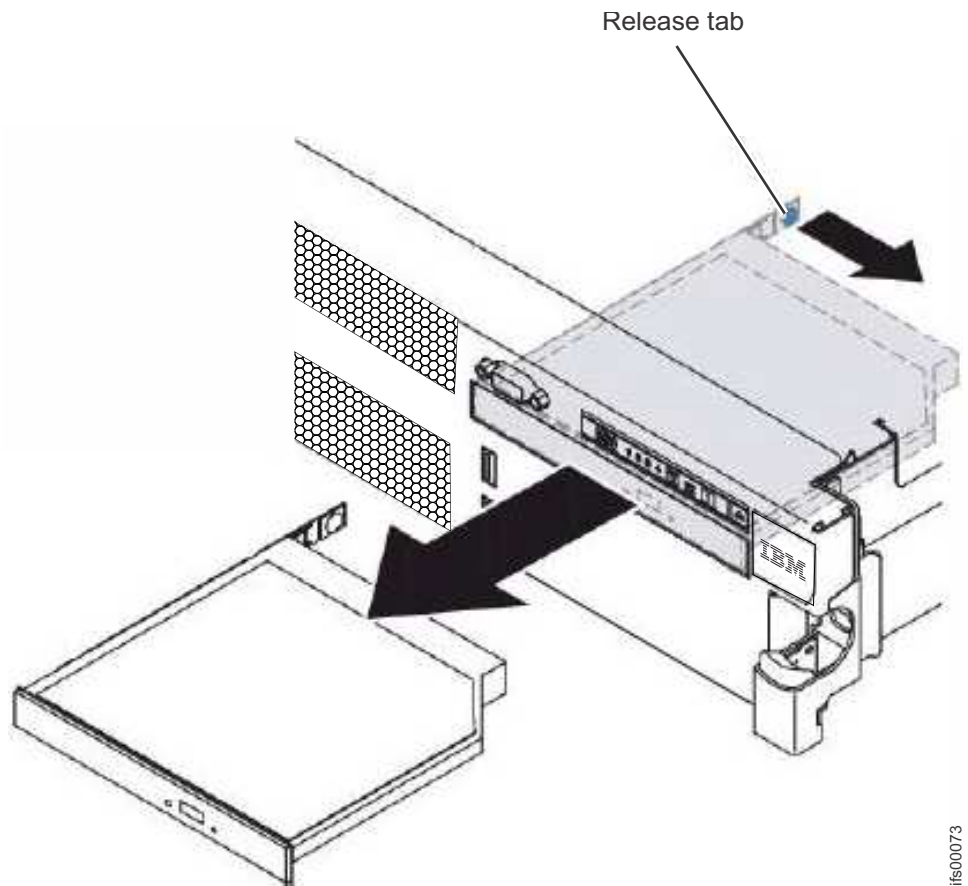
After you replace a failed hard disk drive, the green activity LED flashes as the disk spins up. The yellow LED turns off after approximately 1 minute. If the new drive starts to rebuild, the yellow LED flashes slowly, and the green activity LED remains lit during the rebuild process.

Removing the DVD drive

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

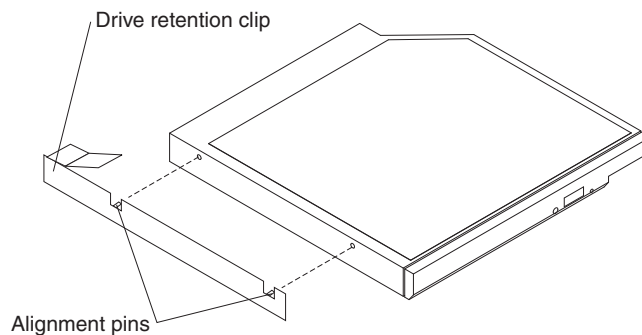
About this task

To remove the DVD drive, complete the following steps.



Procedure

1. Read the Safety information and "Installation guidelines" on page 94.
2. Follow the procedure in "Removing a file module and disconnecting power" on page 92 to suspend the file module from the cluster and shut it down, and then disconnect all power cords and external cables.
3. Slide the file module out of the rack, and then remove the cover (see "Removing the cover" on page 123).
4. Press the release tab down to release the drive; then, while you press the tab, push the drive toward the front of the file module.
5. From the front of the file module, pull the drive out of the bay.



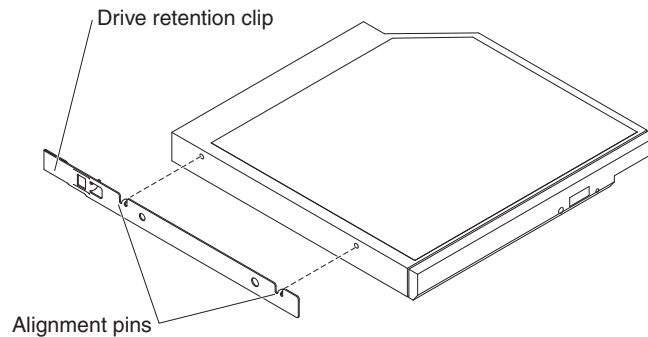
6. Remove the drive retention clip from the drive.
7. If you are instructed to return the DVD drive, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Installing the DVD drive

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

To install the replacement DVD drive, complete the following steps.



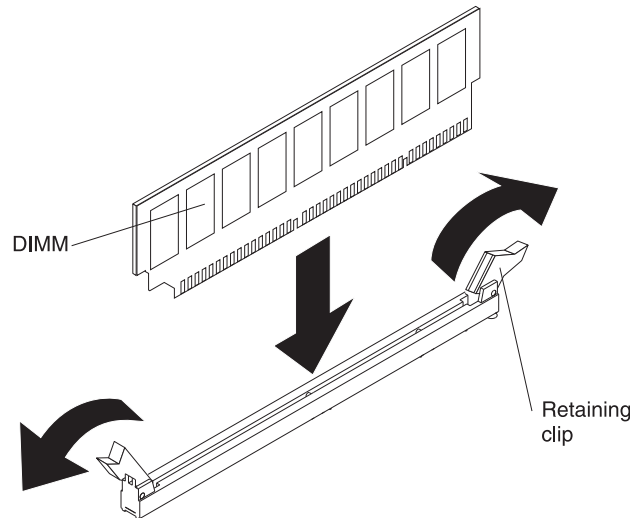
1. Attach the drive-retention clip to the side of the drive.
2. Slide the drive into the CD/DVD drive bay until the drive clicks into place.
3. Install the cover (see "Installing the cover" on page 124).
4. Slide the file module into the rack.
5. Follow the steps at the end of the procedure "Removing a file module and disconnecting power" on page 92 to reconnect the file module and resume its use in the cluster.

Removing a memory module

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

To remove a DIMM, complete the following steps.



Procedure

1. Read the Safety information and “Installation guidelines” on page 94.
2. Follow the procedure in “Removing a file module and disconnecting power” on page 92 to suspend the file module from the cluster and shut it down, and then disconnect all power cords and external cables.
3. Slide the file module out of the rack.
4. Remove the cover (see “Removing the cover” on page 123).
5. Remove the air baffle over the DIMMs (see “Removing the air baffle” on page 137).
6. Open the retaining clip on each end of the DIMM connector and lift the DIMM from the connector.
Attention: To avoid breaking the retaining clips or damaging the DIMM connectors, open and close the clips gently.
7. If you are instructed to return the DIMM, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Installing a memory module

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

See Figure 68 on page 154 for the locations of the DIMM connectors on the system board.

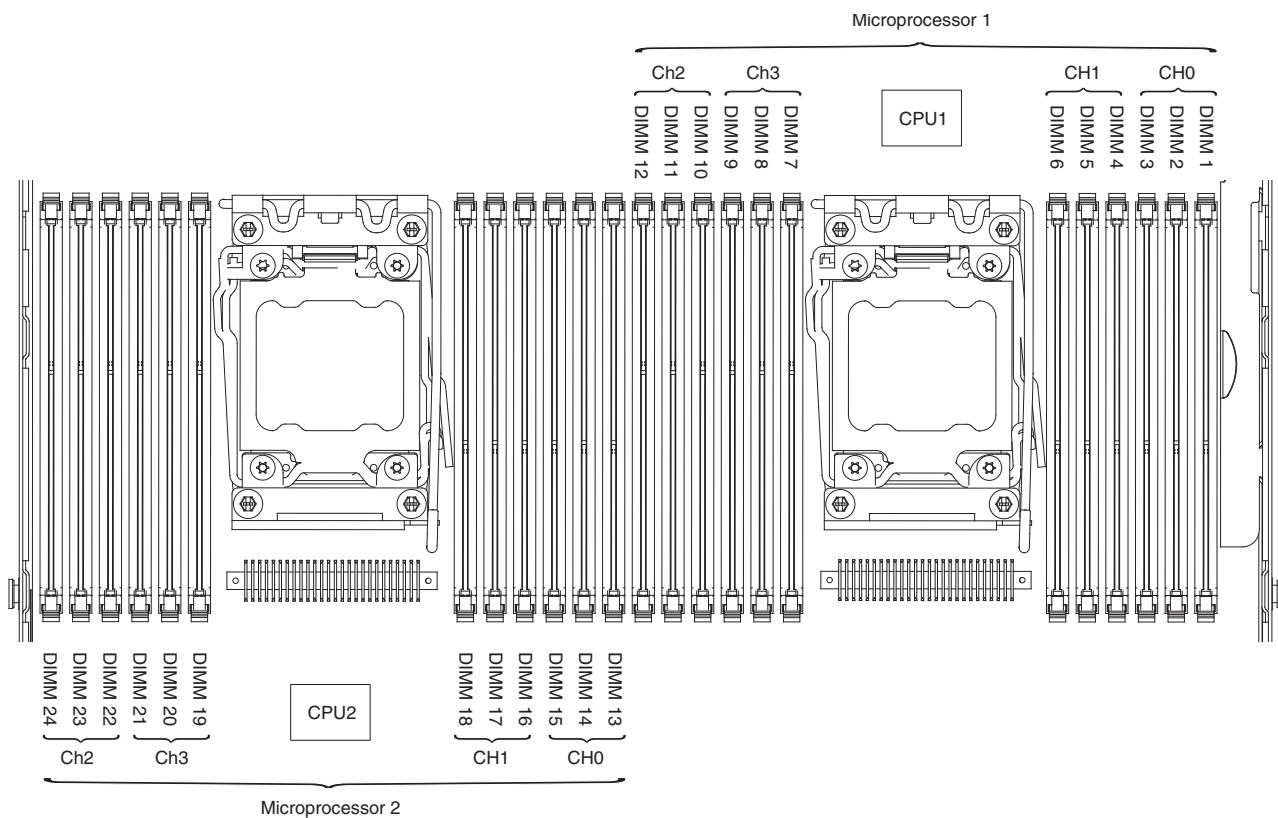
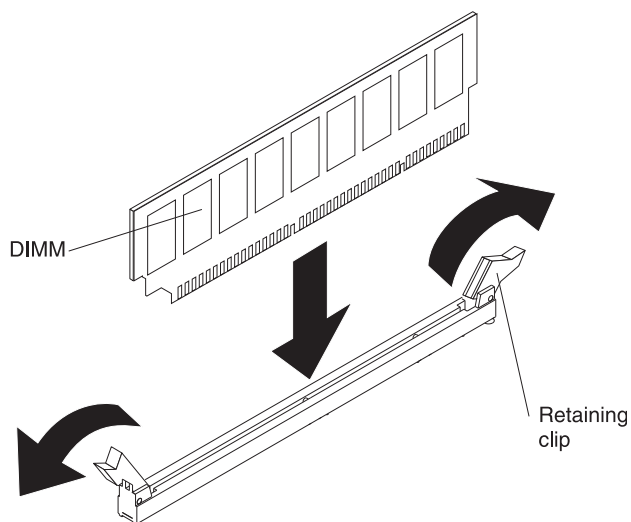


Figure 68. Locations of the DIMM connectors on the system board

To install a DIMM, complete the following procedure.



See Table 54 on page 155 for a listing of the eight DIMM slots populated with the memory RDIMM.

Table 54. DIMM slots populated with the memory RDIMM

Processor	Memory Channel	DIMM Slot Number
1	0	1 - 16GB RDIMM
		2 - 2GB RDIMM
	1	4 - 16GB RDIMM
		5 - 2GB RDIMM
	2	12 - 16GB RDIMM
		11 - 2GB RDIMM
	3	9 - 16GB RDIMM
		8 - 2GB RDIMM
2	0	13 - 16GB RDIMM
		14 - 2GB RDIMM
	1	16 - 16GB RDIMM
		17 - 2GB RDIMM
	2	24 - 16GB RDIMM
		23 - 2GB RDIMM
	3	21 - 16GB RDIMM
		20 - 2GB RDIMM

Note: Do not put any DIMM into DIMM slots 3, 6, 7, 10, or slots 13 to 24.

Procedure

1. Remove the air baffle over the DIMMs (see “Removing the air baffle” on page 137).
2. Open the retaining clip on each end of the DIMM connector.
Attention: To avoid breaking the retaining clips or damaging the DIMM connectors, open and close the clips gently.
3. Touch the static-protective package that contains the DIMM to any unpainted metal surface on the file module, and then remove the DIMM from the package.
4. Turn the DIMM so that the DIMM keys align correctly with the connector.
5. Insert the DIMM into the connector by aligning the edges of the DIMM with the slots at the ends of the DIMM connector. Firmly press the DIMM straight down into the connector by applying pressure on both ends of the DIMM simultaneously. The retaining clips snap into the locked position when the DIMM is firmly seated in the connector.
Attention: If there is a gap between the DIMM and the retaining clips, the DIMM has not been correctly inserted; open the retaining clips, remove the DIMM, and then reinsert it.
6. Repeat steps 1 through 5 until all the new or replacement DIMMs are installed.
7. Replace the air baffle over the DIMMs (see “Installing the air baffle” on page 138), making sure all cables are out of the way.
8. Install the cover (see “Installing the cover” on page 124).
9. Slide the file module into the rack.

10. Follow the steps at the end of the procedure “Removing a file module and disconnecting power” on page 92 to reconnect the file module and resume its use in the cluster.
11. Go to the management GUI and look for any unfixed events related to DIMMs.

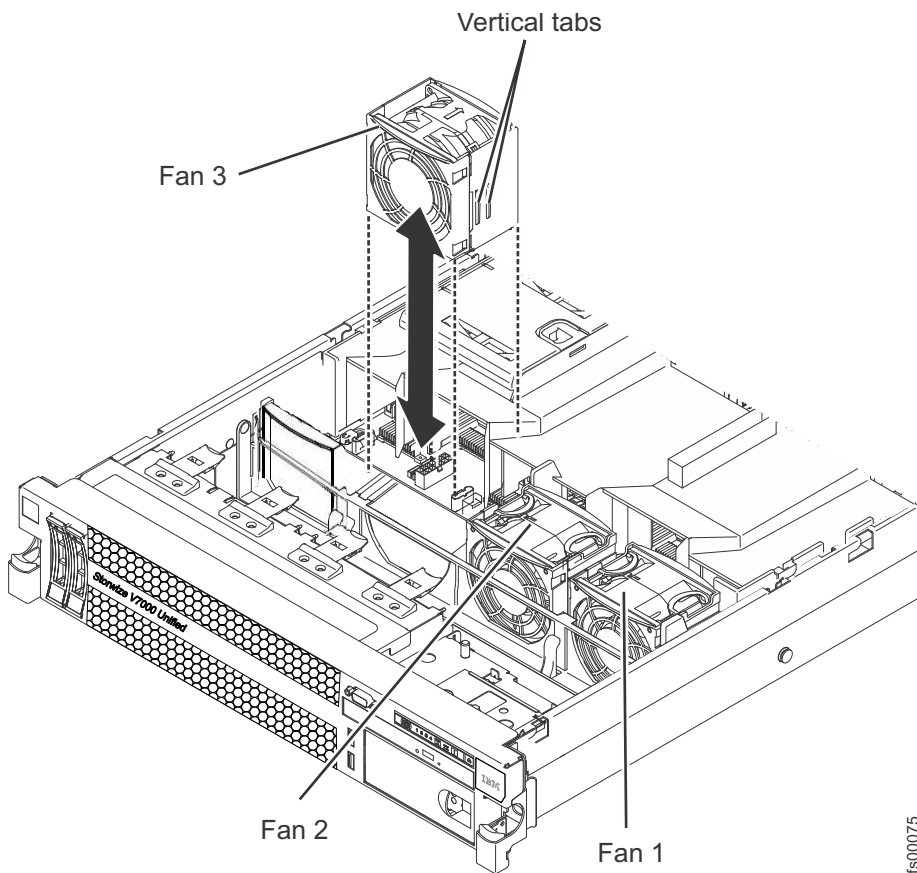
Removing a hot-swap fan

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

Attention: To ensure proper file module operation and cooling, if you remove a dual-motor hot-swap fan with the system running, you must install a replacement dual-motor hot-swap fan within 30 seconds or the system will shut down.

To remove any of the three replaceable dual-motor hot-swap fans, complete the following steps.



Procedure

1. Read the Safety information and “Installation guidelines” on page 94.
2. Leave the file module connected to power.

3. Slide the file module out of the rack and remove the cover (see “Removing the cover” on page 123). The LED on the system board near the connector for the failing dual-motor hot-swap fan will be lit.

Attention: To ensure proper system cooling, do not remove the top cover for more than 30 minutes during this procedure.

4. Grasp the dual-motor hot-swap fan by the finger grips on the sides of the dual-motor hot-swap fan.
5. Rotate the air baffle up.
6. Lift the dual-motor hot-swap fan out of the file module.
7. Replace the dual-motor hot-swap fan within 30 seconds.
8. If you are instructed to return the dual-motor hot-swap fan, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Installing a hot-swap fan

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

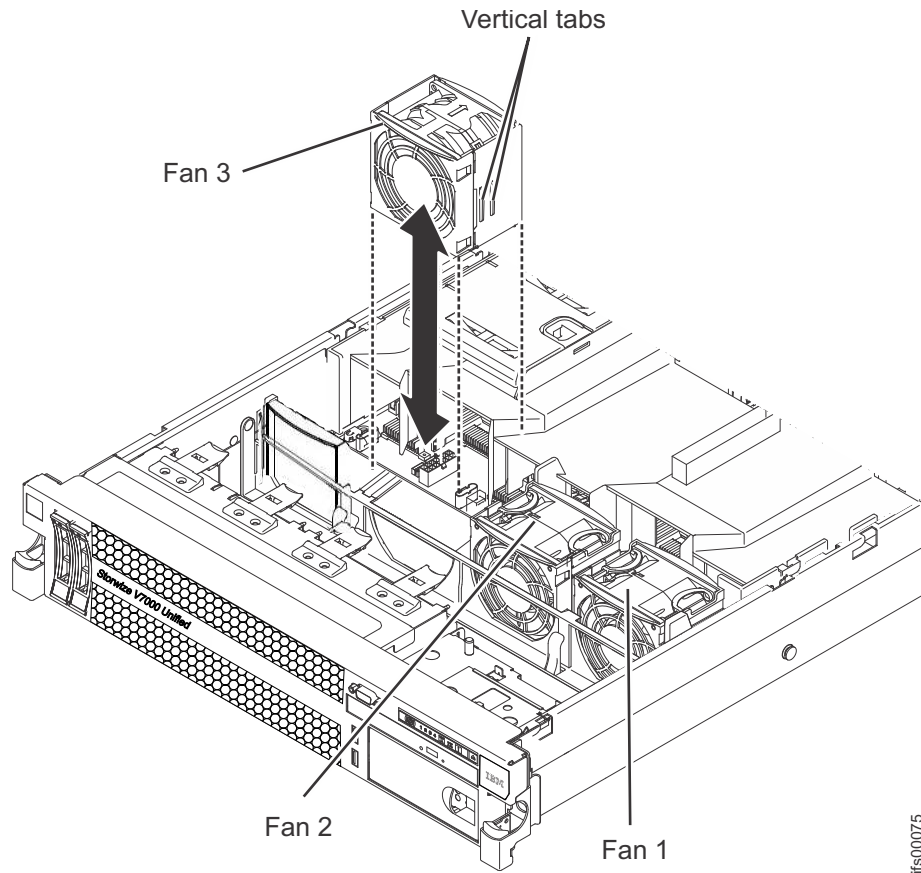
About this task

For proper cooling, the file module requires that all three dual-motor hot-swap fans be installed at all times.

Attention: To ensure proper file module operation, if a dual-motor hot-swap fan fails, replace it immediately. Have a replacement dual-motor hot-swap fan ready to install as soon as you remove the failed dual-motor hot-swap fan.

See System-board internal connectors for the locations of the dual-motor hot-swap fan connectors.

To install any of the three replaceable fans, complete the following steps.



Procedure

1. Rotate the air baffle up.
2. Orient the new dual-motor hot-swap fan over its position in the dual-motor hot-swap fan bracket so that the connector on the bottom aligns with the dual-motor hot-swap fan connector on the system board.
3. Align the vertical tabs on the dual-motor hot-swap fan with the slots on the dual-motor hot-swap fan cage bracket.
4. Push the new dual-motor hot-swap fan into the dual-motor hot-swap fan connector on the system board. Press down on the top surface of the dual-motor hot-swap fan to seat the dual-motor hot-swap fan fully. (Make sure that the LED has turned off.)
5. Repeat steps 1 through 3 until all the new or replacement dual-motor hot-swap fans are installed.
6. Install the cover (see "Installing the cover" on page 124).
7. Slide the file module into the rack.

Removing a hot-swap ac power supply

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

To remove a power supply, complete the following steps:

Procedure

1. Read the Safety information and “Installation guidelines” on page 94.
2. If only one power supply is installed, turn off the server and peripheral devices.
3. Disconnect the power cord from the power supply that you are removing.
4. Grasp the power-supply handle. Refer to Figure 69.

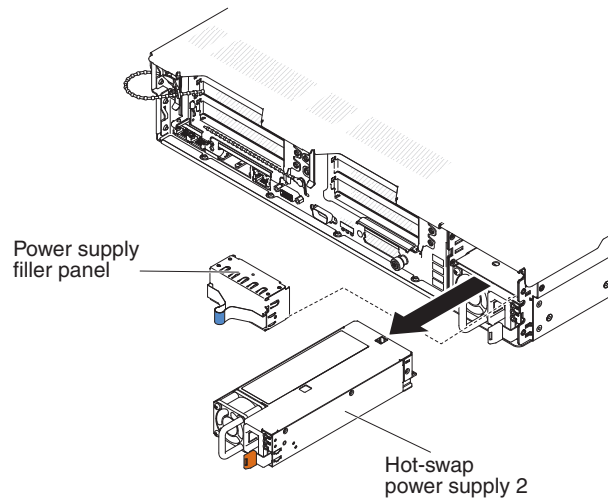


Figure 69. Removing a hot-swap ac power supply

5. Press the orange release latch to the left and hold it in place.
6. Pull the power supply part of the way out of the bay, then release the latch and support the power supply as you pull it the rest of the way out of the bay.
7. If you are instructed to return the power supply, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Installing a hot-swap ac power supply

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

The following notes describe the type of ac power supply that the server supports and other information that you must consider when you install a power supply:

- Make sure that the devices that you are installing are supported. For a list of supported optional devices for the server, see <http://www.ibm.com/systems/info/x86servers/serverproven/compat/us/>.
- Before you install an additional power supply or replace a power supply with one of a different wattage, you may use the IBM Power Configurator utility to determine current system power consumption. For more information and to download the utility, go to <http://www-03.ibm.com/systems/bladecenter/resources/powerconfig.html>.
- The server comes with one hot-swap 12-volt output power supply that connects to power supply bay 1. The input voltage is 100-127 V ac or 200-240 V ac auto-sensing.

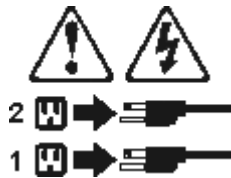
- Power supplies in the server must be with the same power rating or wattage to ensure that the server will operate correctly. For example, you cannot mix 750-watt and 900-watt power supplies in the server.
- Power supply 1 is the default/primary power supply. If power supply 1 fails, you must replace the power supply immediately.
- You can order an optional power supply for redundancy.
- These power supplies are designed for parallel operation. In the event of a power-supply failure, the redundant power supply continues to power the system. The server supports a maximum of two power supplies.

Statement 5



CAUTION:

The power control button on the device and the power switch on the power supply do not turn off the electrical current supplied to the device. The device also might have more than one power cord. To remove all electrical current from the device, ensure that all power cords are disconnected from the power source.



Statement 8

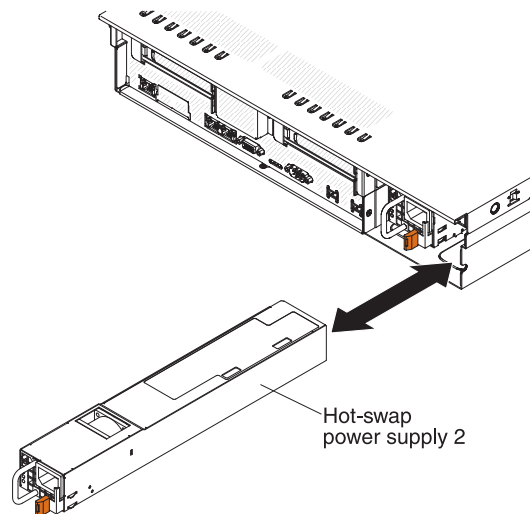


CAUTION:

Never remove the cover on a power supply or any part that has the following label attached.



Hazardous voltage, current, and energy levels are present inside any component that has this label attached. There are no serviceable parts inside these components. If you suspect a problem with one of these parts, contact a service technician.

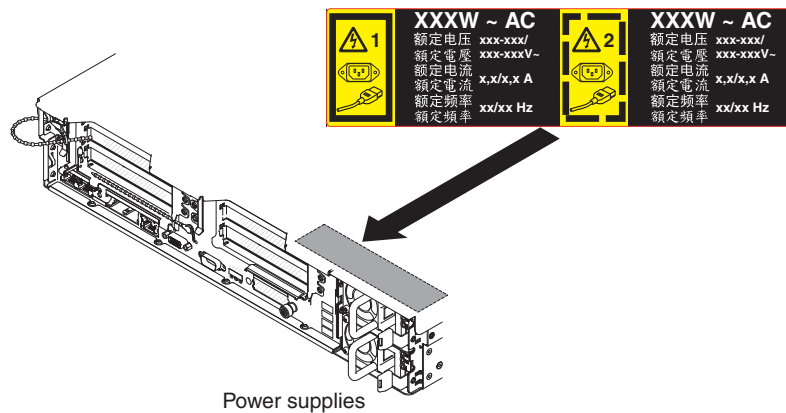


Attention: During normal operation, each power-supply bay must have either a power supply or power-supply filler installed for proper cooling.

To install a hot-swap ac power supply, complete the following steps:

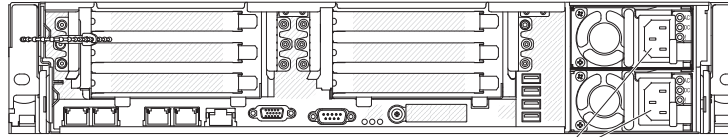
Procedure

1. Read the Safety information and “Installation guidelines” on page 94.
2. Touch the static-protective package that contains the hot-swap power supply to any unpainted metal surface on the server; then, remove the power supply from the package and place it on a static-protective surface.
3. If you are adding a power supply to the server, attach the redundant power information label that comes with this option on the server cover near the power supplies.



4. Slide the power supply into the bay until the retention latch clicks into place.
5. Connect the power cord for the new power supply to the power-cord connector on the power supply.

The following illustration shows the power-cord connectors on the back of the server.



Power cord connectors

6. Route the power cord through the clip next to power-supply and through any cable clamps on the rear of the server, to prevent the power cord from being accidentally pulled out when you slide the server in and out of the rack.
7. Connect the power cord to a properly grounded electrical outlet.
8. Make sure that the error LED on the power supply is not lit, and that the ac power LED on the power supply are lit, indicating that the power supply is operating correctly.
9. If you are replacing a power supply with one of a different wattage, apply the power information label provided with the new power supply over the existing power information label on the server.

额定电压 xxx-xxx/xxx-xxx	额定電壓 x,x/x,x	额定電壓 xxx-xxx/xxx-xxx
额定电流 xx/xx Hz	额定電流 xx/xx Hz	额定電流 xx/xx Hz
额定频率	额定頻率	额定頻率

Marka Registrada
©Registered Trademark
of International Business
Machines Corporation

Product certified in Shenzhen, China
Made in China V 中国制造

Apparaten skall anslutas till jordat uttag
Apparatet må tilkoples jordnet stikkontakt
Laita on liitettävä suojamaadoituskoskettimilla
varustettuun pistorasiaan
This device complies with part 15 of FCC rules.
Operation is subject to the following two
conditions: (1) this device may not cause harmful
interference, and (2) this device must accept any
interference received, including interference that
may cause undesired operation.

製造商 Manufacturer: IBM Corporation
Copyright Code and Parts Contained Herein.
©Copyright IBM Corp. 2010 All Rights Reserved.
Canada ICES/NMB-003 Class/Classe A

この装置は、クラス A 情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。V C C I - A

廢電池請回收

警告使用者:
這是甲類的資訊產品, 在居住的環境中使用時, 可能會造成射頻干擾, 在這種情況下, 使用者會被要求採取某些適當的對策。

伺服器 服务器
型号 MT: XXXX
Model: xxx
SN: SSSSSSS
MFG date: YYYYYMDD
Product ID:
PN:

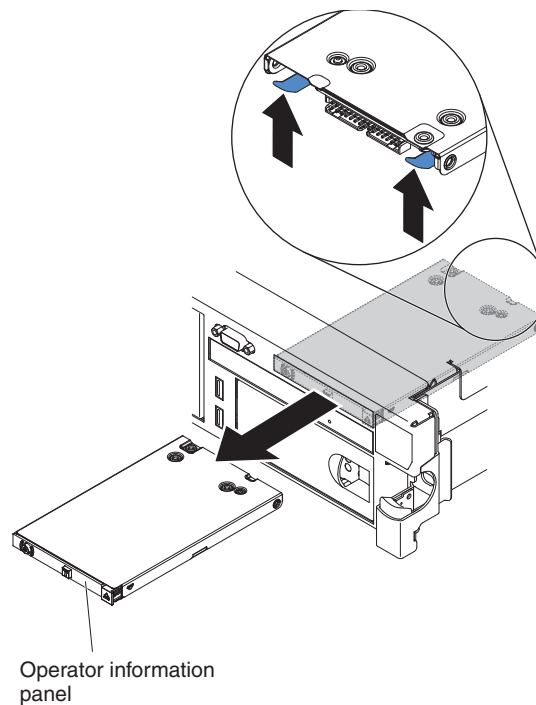
Removing the operator information panel assembly

The following procedure is for a Tier 1 customer replaceable unit (CRU).

Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

To remove the operator information panel assembly, complete the following steps.



Procedure

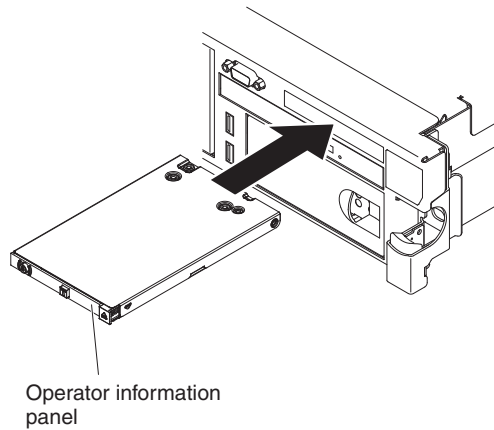
1. Read the Safety information and “Installation guidelines” on page 94.
2. Follow the procedure in “Removing a file module and disconnecting power” on page 92 to suspend the file module from the cluster and shut it down, and then disconnect all power cords and external cables.
3. Remove the cover (see “Removing the cover” on page 123).
4. Disconnect the cable from the back of the operator information panel assembly.
5. Reach inside the file module and press the release tab; then, while you hold the release tab down, push the assembly toward the front of the file module.
6. From the front of the file module, carefully pull the operator information panel assembly out of the file module.
7. If you are instructed to return the operator information panel assembly, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Installing the operator information panel assembly

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

To install the replacement operator information panel assembly, complete the following steps.



Procedure

1. Read the Safety information and “Installation guidelines” on page 94.
2. Position the operator information panel assembly so that the tabs face upward and slide it into the file module until it clicks into place.
3. Inside the file module, connect the cable to the rear of the operator information panel assembly.
4. Install the cover (see “Installing the cover” on page 124).
5. Slide the file module into the rack.
6. Follow the steps at the end of the procedure “Removing a file module and disconnecting power” on page 92 to reconnect the file module and resume its use in the cluster.

Removing the hot-swap drive backplane

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

To remove the hot-swap drive backplane, complete the following steps.

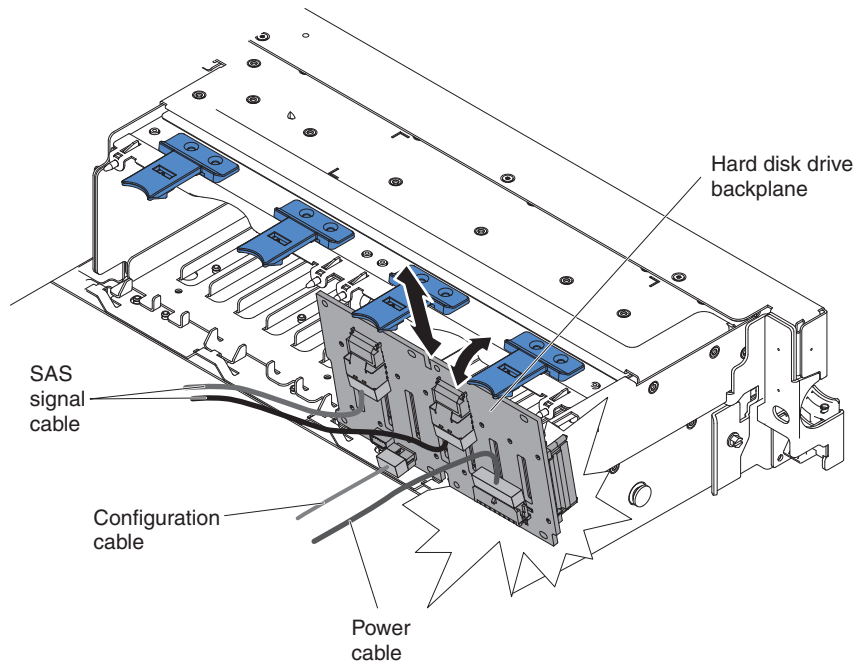


Figure 70. Removing the hot-swap drive backplane

Procedure

1. Read the safety information that begins on page Safety and “Installation guidelines” on page 94.
2. Follow the procedure in “Removing a file module and disconnecting power” on page 92 to suspend the file module from the cluster and shut it down and disconnect all power cords and external cables.
3. Slide the file module out of the rack.
4. Remove the cover. For more information, see Removing the cover.
5. Pull the hard disk drives or fillers out of the file module slightly to disengage them from the backplane.
6. To obtain more working room, remove the fans.
7. Lift the backplane out of the file module by pulling it toward the rear of the file module and then lifting it up.
8. Disconnect the backplane power cable, SAS signal cable, and configuration cable.
9. If you are instructed to return the backplane, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Installing the hot-swap drive backplane

The following procedure is for a Tier 1 customer replaceable unit (CRU). Replacement of Tier 1 CRUs is your responsibility. If IBM installs a Tier 1 CRU at your request, you will be charged for the installation. Service agreements can be purchased so that you can ask IBM to replace these units.

About this task

To install the replacement hot-swap drive backplane, complete the following steps.

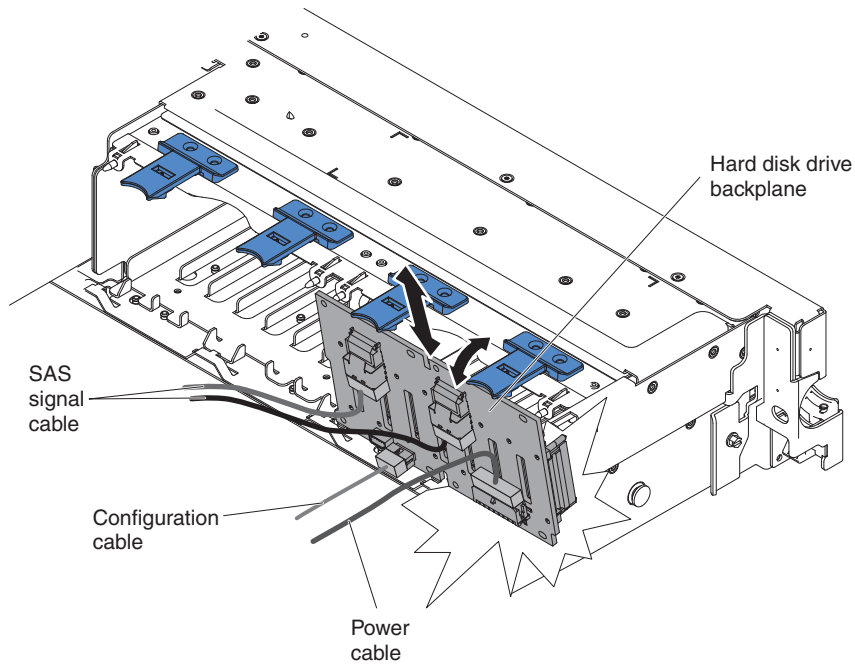


Figure 71. Installing the hot-swap drive backplane

Procedure

1. Connect the power and signal cables to the replacement backplane.
2. Align the backplane with the backplane slot in the chassis and the small slots on top of the hard disk drive cage.
3. Lower the backplane into the slots on the chassis.
4. Rotate the top of the backplane until the front tab clicks into place into the latches on the chassis.
5. Insert the hard disk drives and the fillers the rest of the way into the bays.
6. Replace the fan bracket and fans if you removed them.
7. Install the cover. For more information, see *Installing the cover*.
8. Slide the file module into the rack.
9. Follow the steps at the end of the procedure in “Removing a file module and disconnecting power” on page 92 to suspend the file module to reconnect the file module and resume its use in the cluster.

Removing a microprocessor and heat sink

IBM authorized service providers can remove and replace a microprocessor and heat sink in the file module. The following procedure is for a field replaceable unit (FRU). FRUs must be installed only by trained service technicians.

About this task

Attention:

- Always use the microprocessor installation tool to remove a microprocessor. Failing to use the microprocessor installation tool may damage the microprocessor sockets on the system board. Any damage to the microprocessor sockets may require replacing the system board.
- Microprocessors are to be removed only by trained service technicians.
- Do not allow the thermal grease on the microprocessor and heat sink to come in contact with anything. Contact with any surface can compromise the thermal grease and the microprocessor socket.
- Dropping the microprocessor during installation or removal can damage the contacts.
- Do not touch the microprocessor contacts; handle the microprocessor by the edges only. Contaminants on the microprocessor contacts, such as oil from your skin, can cause connection failures between the contacts and the socket.

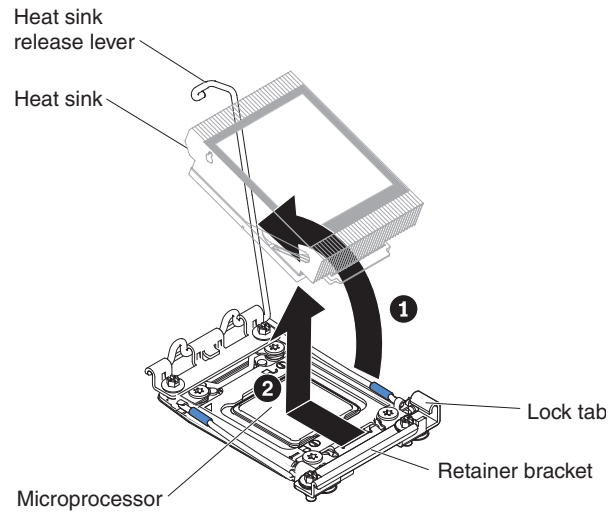
To remove a microprocessor and heat sink, complete the following steps:

Procedure

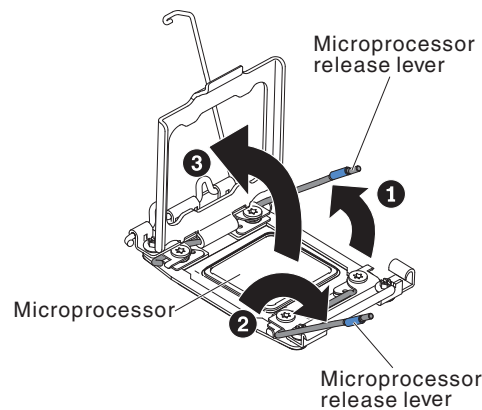
1. Read the Safety information, “Handling static-sensitive devices” on page 96, and “Installation guidelines” on page 94.
2. Follow the procedure in “Removing a file module and disconnecting power” on page 92 to suspend the file module from the cluster and shut it down, and then disconnect all power cords and external cables.
3. Remove the cover (see “Removing the cover” on page 123).
4. Remove the following components, if necessary:
 - PCI riser-card assembly 1 (see “Removing a PCI riser-card assembly” on page 142)
 - DIMM air baffle (see “Removing the air baffle” on page 137)
5. Disconnect any cables that impede access to the heat sink and the microprocessor.
6. Locate the microprocessor to be removed (see System-board internal connectors).
7. Remove the heat sink.

Attention: Do not touch the thermal material on the bottom of the heat sink. Touching the thermal material will contaminate it. If the thermal material on the microprocessor or heat sink becomes contaminated, you must wipe off the contaminated thermal material on the microprocessor or heat sink with the alcohol wipes and reapply clean thermal grease to the heat sink.

- a. Open the heat sink release lever to the fully open position.
- b. Lift the heat sink out of the file module. After removal, place the heat sink (with the thermal grease side up) on a clean, flat surface.



8. Open the microprocessor socket release levers and retainer:



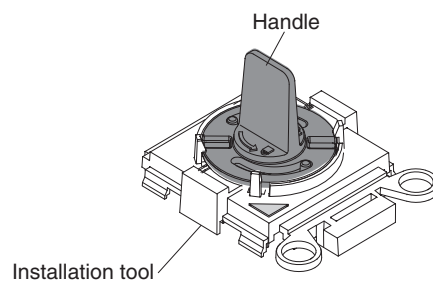
- a. Identify which release lever is labeled as the first release lever to open and open it.
- b. Open the second release lever on the microprocessor socket.
- c. Open the microprocessor retainer.

Attention: Do not touch the connectors on the microprocessor and the microprocessor socket.

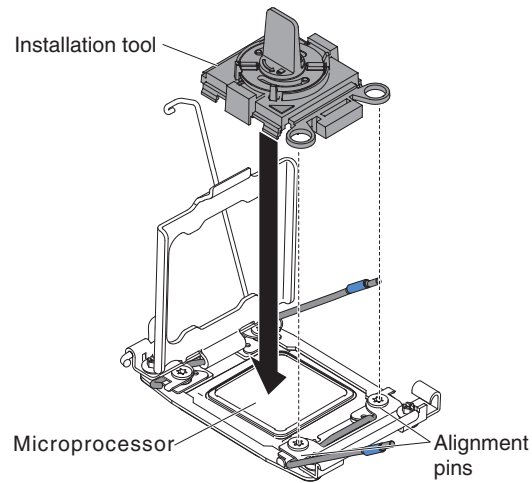
9. Install the microprocessor on the microprocessor installation tool.

Note: If you are replacing a microprocessor, use the empty installation tool that comes with the CRU to remove the microprocessor.

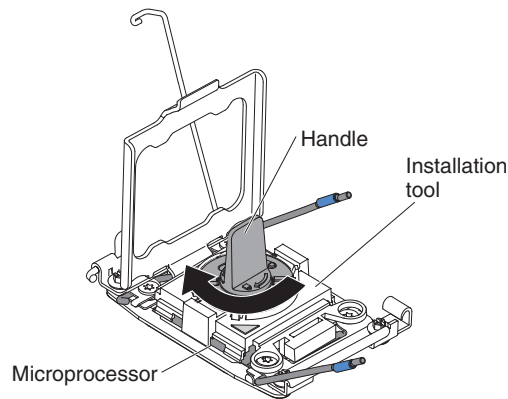
- a. Twist the handle on the microprocessor tool counterclockwise so that it is in the open position.



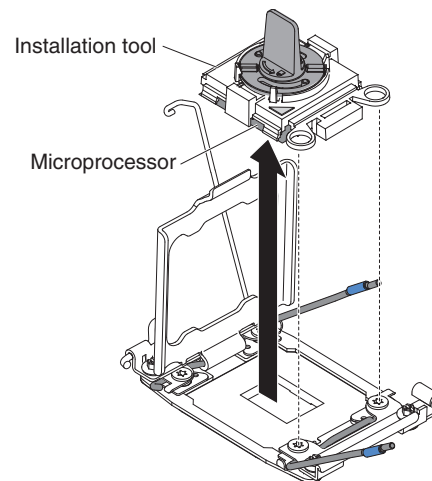
- b. Align the installation tool with the alignment pins on the microprocessor socket and lower the tool on the microprocessor. The installation tool rests flush on the socket only if aligned correctly.



- c. Twist the handle on the installation tool clockwise.



- d. Lift the microprocessor out of the socket.



10. If you do not intend to install a microprocessor on the socket, install the socket cover that you removed in step 8 on page 173 of "Installing a microprocessor and heat sink" on page 170 on the microprocessor socket.

Attention: The pins on the socket are fragile. Any damage to the pins may require replacing the system board.

11. If you are instructed to return the microprocessor, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Installing a microprocessor and heat sink

IBM authorized service providers can remove and replace a microprocessor and heat sink in the file module. The following procedure is for a field replaceable unit (FRU). FRUs must be installed only by trained service technicians.

About this task

The following notes describe the type of microprocessor that the file module supports and other information that you must consider when you install a microprocessor and heat sink:

- Microprocessors are to be installed only by trained service technicians.
- A 2073-720 file module supports one (1) microprocessor. See Parts listing for 2073-720 file modules.
- The microprocessor must always be installed in microprocessor socket 1 on the system board.
- The air baffle must be installed to provide proper system cooling.
- If you have to replace the microprocessor, call IBM Remote Technical Support for service.
- If the thermal-grease protective cover (for example, a plastic cap or tape liner) is removed from the heat sink, do not touch the thermal grease on the bottom of the heat sink or set down the heat sink. For more information about applying or working with thermal grease, see “Removing and replacing the thermal grease” on page 174.

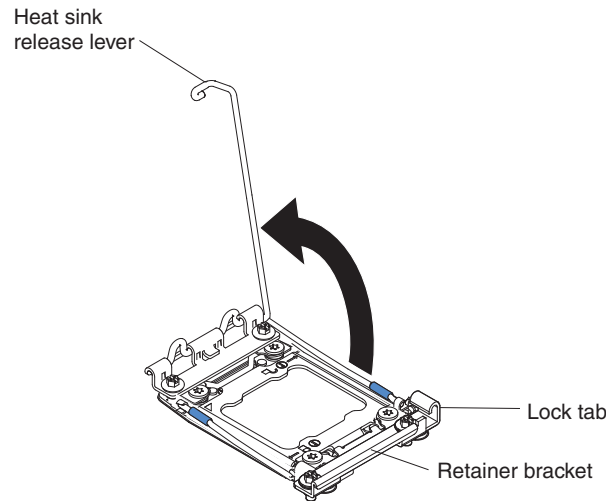
Note: Removing the heat sink from the microprocessor destroys the even distribution of the thermal grease and requires replacing the thermal grease.

To install an additional microprocessor and heat sink, complete the following steps:

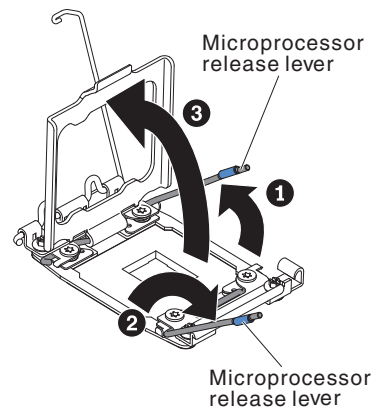
Procedure

1. Read the Safety information and “Installation guidelines” on page 94.
2. Follow the procedure in “Removing a file module and disconnecting power” on page 92 to suspend the file module from the cluster and shut it down, and then disconnect all power cords and external cables.

Attention: When you handle static-sensitive devices, take precautions to avoid damage from static electricity. For details about handling these devices, see “Handling static-sensitive devices” on page 96.
3. Remove the cover (see “Removing the cover” on page 123).
4. Remove the following components, if necessary:
 - PCI riser-card assembly 1 (see “Removing a PCI riser-card assembly” on page 142)
 - DIMM air baffle (see “Removing the air baffle” on page 137)
5. Rotate the heat sink release lever to the open position.



6. Open the microprocessor socket release levers and retainer:

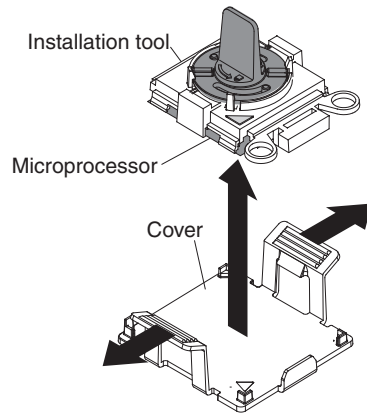


- a. Identify which release lever is labeled as the first release lever to open and open it.
- b. Open the second release lever on the microprocessor socket.
- c. Open the microprocessor retainer.

Attention: Do not touch the connectors on the microprocessor and the microprocessor socket.

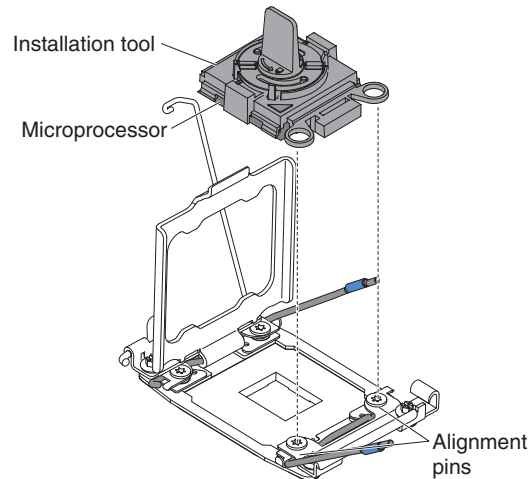
7. Install the microprocessor on the microprocessor socket:

- a. Touch the static-protective package that contains the new microprocessor to any *unpainted* on the chassis or any *unpainted* metal surface on any other grounded rack component; then, carefully remove the microprocessor from the package.
- b. Release the sides of the cover and remove the cover from the installation tool. The microprocessor is preinstalled on the installation tool.

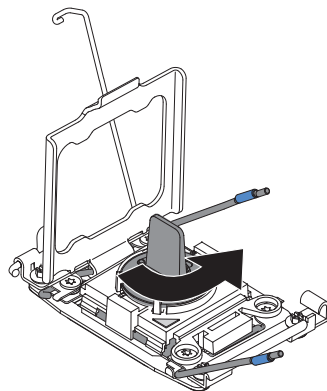


Note: Do not touch the microprocessor contacts. Contaminants on the microprocessor contacts, such as oil from your skin, can cause connection failures between the contacts and the socket.

- c. Align the installation tool with the microprocessor socket. The installation tool rests flush on the socket only if properly aligned.



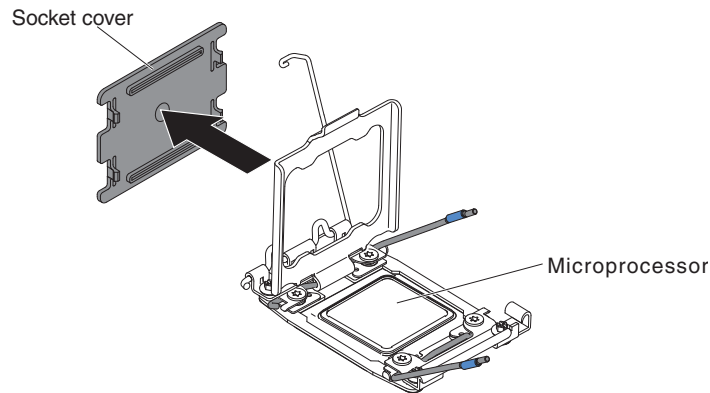
- d. Twist the handle on the microprocessor tool counterclockwise to insert the microprocessor into the socket. The microprocessor is keyed to ensure that the microprocessor is installed correctly. The microprocessor rests flush on the socket only if properly installed.



Attention:

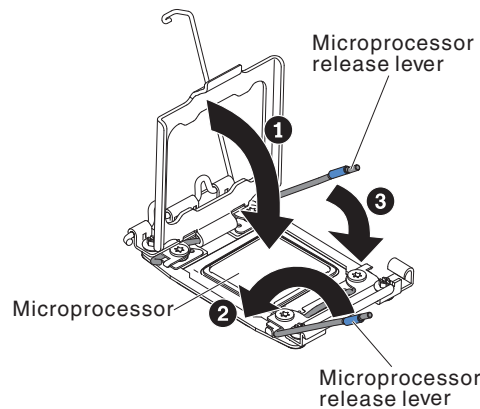
- Do not press the microprocessor into the socket.
- Make sure that the microprocessor is oriented and aligned correctly in the socket before you try to close the microprocessor retainer.
- Do not touch the thermal material on the bottom of the heat sink or on top of the microprocessor. Touching the thermal material will contaminate it.

8. Remove the microprocessor socket cover, tape, or label from the surface of the microprocessor socket, if one is present. Store the socket cover in a safe place.



Attention: When you handle static-sensitive devices, take precautions to avoid damage from static electricity. For details about handling these devices, see “Handling static-sensitive devices” on page 96.

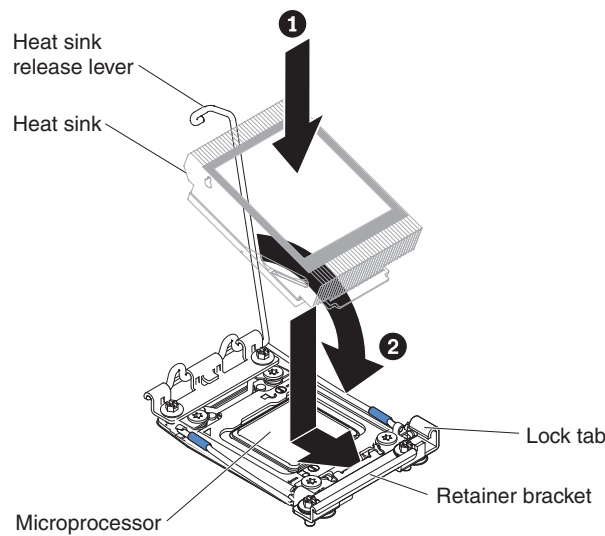
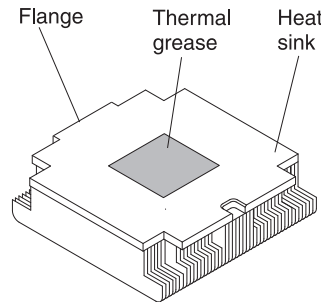
9. Close the microprocessor socket release levers and retainer:



- a. Close the microprocessor retainer on the microprocessor socket.
 - b. Identify which release lever is labeled as the first release lever to close and close it.
 - c. Close the second release lever on the microprocessor socket.
10. Install the heat sink.

Attention:

- Do not set down the heat sink after you remove the plastic cover.
- Do not touch the thermal grease on the bottom of the heat sink after you remove the plastic cover. Touching the thermal grease will contaminate it. See “Removing and replacing the thermal grease” for more information.



- a. Remove the plastic protective cover from the bottom of the heat sink.
 - b. Position the heat sink over the microprocessor. The heat sink is keyed to assist with proper alignment.
 - c. Align and place the heat sink on top of the microprocessor in the retention bracket, thermal material side down.
 - d. Press firmly on the heat sink.
 - e. Rotate the heat sink release lever to the closed position and hook it underneath the lock tab.
11. Reinstall the air baffle (see “Installing the air baffle” on page 138).
 12. Install the cover (see “Installing the cover” on page 124).
 13. Slide the file module into the rack.
 14. Follow the steps at the end of the procedure “Removing a file module and disconnecting power” on page 92 to reconnect the file module and resume its use in the cluster.

Removing and replacing the thermal grease

IBM authorized service providers must replace the thermal grease when the heat sink has been removed from the top of a microprocessor in the file module and the

heat sink is going to be reused or when debris is found in the grease. The following procedure is for a field replaceable unit (FRU). FRUs must be installed only by trained service technicians.

About this task

The thermal grease must be replaced whenever the heat sink has been removed from the top of the microprocessor and is going to be reused or when debris is found in the grease.

When you are installing the heat sink on the same microprocessor that it was removed from, make sure that the following requirements are met:

- The thermal grease on the heat sink and microprocessor is not contaminated.
- Additional thermal grease is not added to the existing thermal grease on the heat sink and microprocessor.

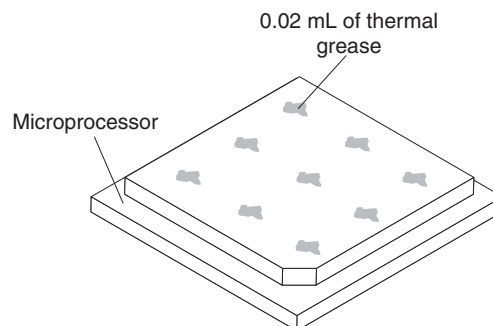
To replace damaged or contaminated thermal grease on the microprocessor and heat exchanger, complete the following steps:

Procedure

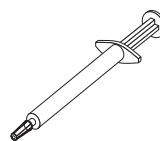
1. Read the Safety information, "Handling static-sensitive devices" on page 96, and "Installation guidelines" on page 94.
2. Place the heat-sink assembly on a clean work surface.
3. Remove the cleaning pad from its package and unfold it completely.
4. Use the cleaning pad to wipe the thermal grease from the bottom of the heat exchanger.

Note: Make sure that all of the thermal grease is removed.

5. Use a clean area of the cleaning pad to wipe the thermal grease from the microprocessor, and then dispose of the cleaning pad after all of the thermal grease is removed.



6. Use the thermal-grease syringe to place 9 uniformly spaced dots of 0.02 mL each on the top of the microprocessor. The outermost dots must be within approximately 5 mm of the edge of the microprocessor; this is to ensure uniform distribution of the grease.



Note: If the grease is properly applied, approximately half of the grease will remain in the syringe.

7. Install the heat sink onto the microprocessor as described in [Install the heat sink](#).

Removing a heat-sink retention module

IBM authorized service providers can remove and replace a heat-sink retention module in the file module. The following procedure is for a field replaceable unit (FRU). FRUs must be installed only by trained service technicians.

About this task

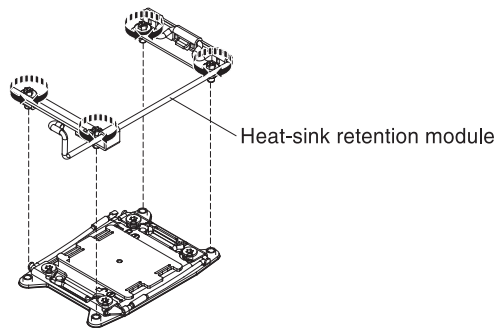
To remove a heat-sink retention module, complete the following steps:

Procedure

1. Read the Safety information and “Installation guidelines” on page 94.
2. Follow the procedure in “Removing a file module and disconnecting power” on page 92 to suspend the file module from the cluster and shut it down, and then disconnect all power cords and external cables.
3. Remove the cover (see “Removing the cover” on page 123).
4. Remove the applicable air baffle, and then remove the heat sink and microprocessor. See “Removing a microprocessor and heat sink” on page 166 for instructions, and then continue with step 5.

Attention: When you remove a microprocessor and heat sink, be sure to keep each heat sink with its microprocessor for reinstallation.

5. Use a screwdriver and remove the four screws that secure the retention module to the system board; then, lift the retention module from the system board.



6. If you are instructed to return the heat-sink retention module, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Installing a heat-sink retention module

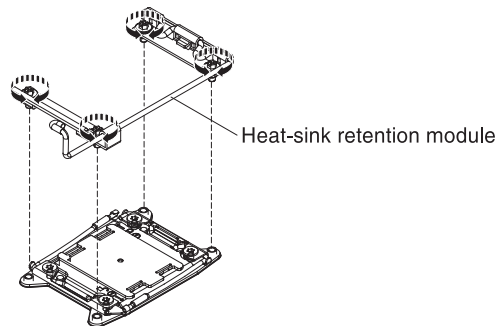
IBM authorized service providers can remove and replace a heat-sink retention module in the file module. The following procedure is for a field replaceable unit (FRU). FRUs must be installed only by trained service technicians.

About this task

To install a heat-sink retention module, complete the following steps:

Procedure

1. Read the Safety information and “Installation guidelines” on page 94.
2. Align the retention module with the holes on the system board.
3. Use a screwdriver to reinstall the four screws.



4. Reinstall the microprocessor and heat sink (see “Installing a microprocessor and heat sink” on page 170).
5. Reinstall the air baffle (see “Installing the air baffle” on page 138).
6. Install the cover (see “Installing the cover” on page 124).
7. Slide the file module into the rack.
8. Follow the steps at the end of the procedure “Removing a file module and disconnecting power” on page 92 to reconnect the file module and resume its use in the cluster.

Removing the system board

IBM authorized service providers can remove and replace the system board in the file module. The following procedure is for a field replaceable unit (FRU). FRUs must be installed only by trained service technicians.

About this task

To remove the system board, complete the following steps.

Procedure

1. Read the Safety information and “Installation guidelines” on page 94.
2. Follow the procedure in “Removing a file module and disconnecting power” on page 92 to suspend the file module from the cluster and shut it down, and then disconnect all power cords and external cables.
3. Pull the power supplies out of the rear of the file module, just enough to disengage them from the file module.
4. Remove the file module cover (see “Removing the cover” on page 123).
5. Remove the riser-card assemblies with adapters (see “Removing a PCI riser-card assembly” on page 142).

Attention: Place all removed components on a static-protective surface for reinstallation.

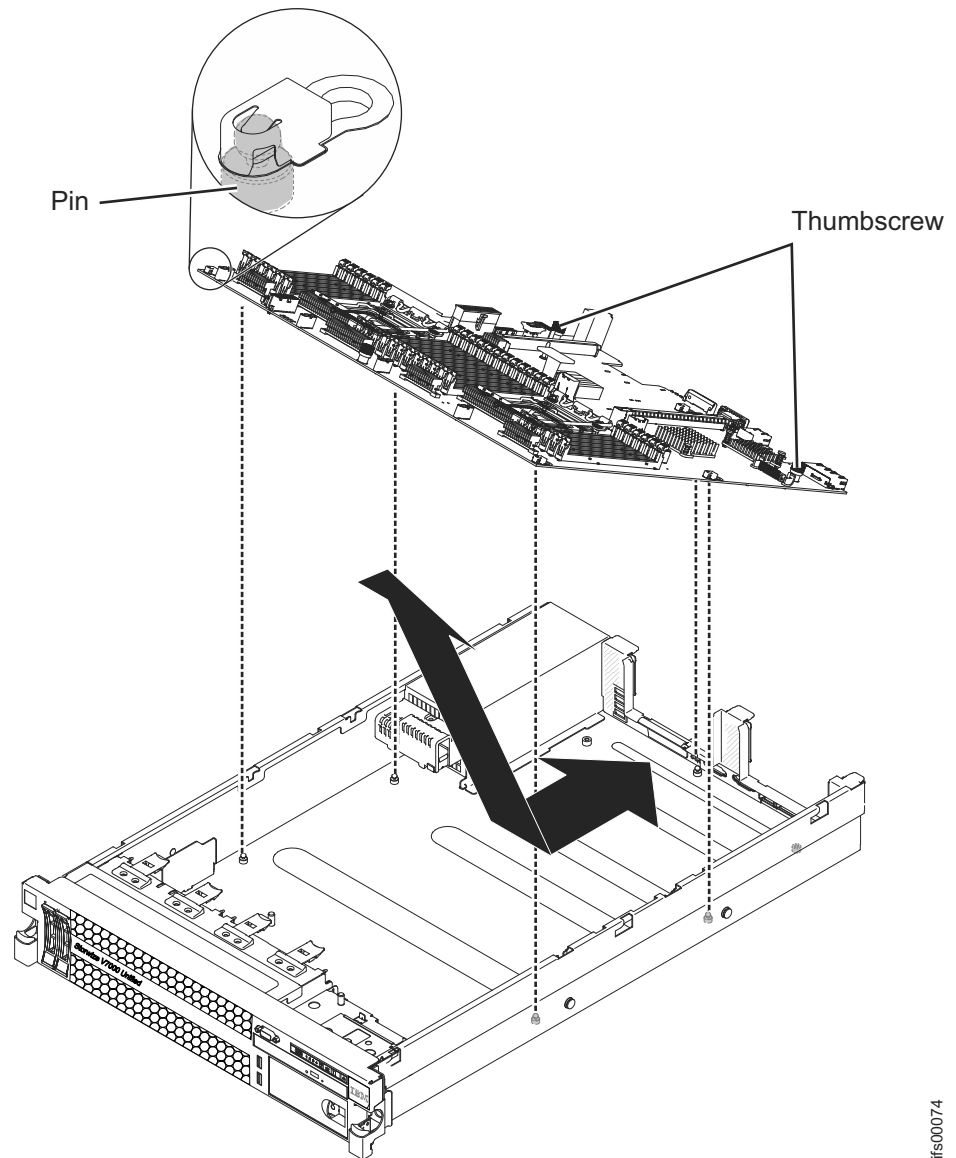
6. Remove the 10 Gbps Ethernet adapter (see Removing a 10-Gbps Ethernet adapter).
7. Remove the air baffle (see “Removing the air baffle” on page 137).
8. Remove all DIMMs, and place them on a static-protective surface for reinstallation (see “Removing a memory module” on page 152).

Important: Before you remove the DIMMs, note which DIMMs are in which connectors. You must install them in the same configuration on the replacement system board.

9. Remove the fans (see “Removing a hot-swap fan” on page 156).
10. Disconnect all cables from the system board.

Attention:

- In the following step, do not allow the thermal grease to come in contact with anything, and keep each heat sink paired with its microprocessor for reinstallation. Contact with any surface can compromise the thermal grease and the microprocessor socket; a mismatch between the microprocessor and its original heat sink can require the installation of a new heat sink.
 - Disengage all latches, release tabs or locks on cable connectors when you disconnect all cables from the system board. Failing to release them before removing the cables will damage the cable sockets on the system board. The cable sockets on the system board are fragile. Any damage to the cable sockets may require replacing the system board.
11. Remove the microprocessor heat sink and microprocessor, and then place them on a static-protective surface for reinstallation (see “Removing a microprocessor and heat sink” on page 166).
 12. Pull out and lift up the pin and the thumbscrews on each side of the system board.



ifs00074

13. Slide the system board forward and tilt it away from the power supplies. Using the two lift handles on the system board, pull the system board out of the file module.
14. If you are instructed to return the system board, follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.
15. Remove the socket covers from the microprocessor sockets on the new system board and place them on the microprocessor sockets of the system board you are removing.
Attention: Make sure to place the socket covers for the microprocessor sockets on the system board before you return the old system board.

Installing the system board

IBM authorized service providers can remove and replace the system board in the file module. The following procedure is for a field replaceable unit (FRU). FRUs must be installed only by trained service technicians.

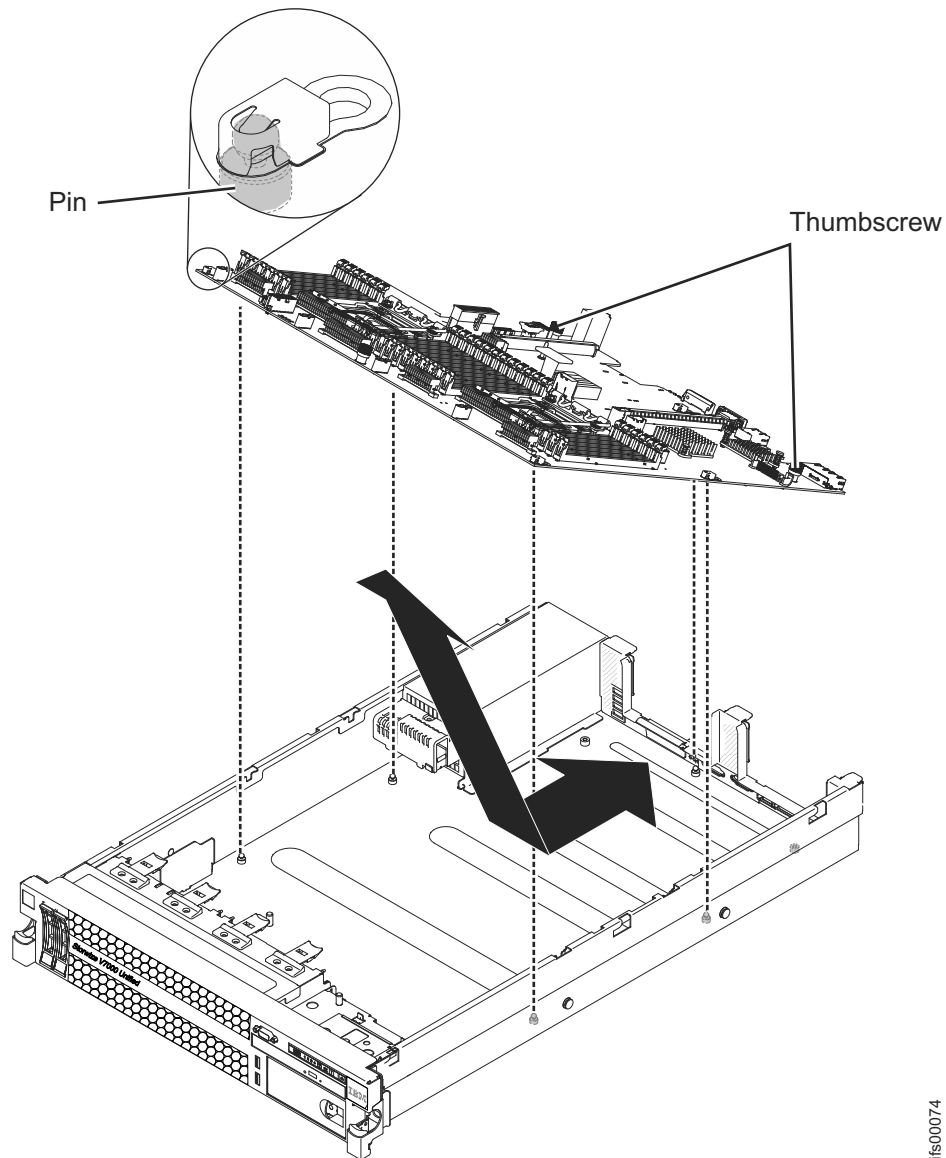
Before you begin

Notes:

1. When you reassemble the components in the file module, be sure to route all cables carefully so that they are not exposed to excessive pressure.
2. Connect a USB keyboard and VGA monitor to the file module before you power it back on in case it asks you to accept a system board configuration change before it will boot up completely.

About this task

To reinstall the system board, complete the following steps.



Procedure

1. Align the system board at an angle, as shown in the illustration; then, rotate and lower it flat and slide it back toward the rear of the file module. Make sure that the rear connectors extend through the rear of the chassis.

2. Reconnect to the system board the cables that you disconnected in step 10 on page 178 of “Removing the system board” on page 177.
3. Rotate the system-board thumbscrews toward the rear of the file module until the latch clicks into place.
4. Install the fans.
5. Install each microprocessor with its matching heat sink (see “Installing a microprocessor and heat sink” on page 170).
6. Install the DIMMs (see “Installing a memory module” on page 153).
7. Install the air baffle (see “Installing the air baffle” on page 138), making sure that all cables are out of the way.
8. Install the 10 Gbps Ethernet adapter (see “Installing a 10-Gbps Ethernet adapter” on page 147).
9. Install the PCI riser-card assemblies and all adapters (see “Installing a PCI riser-card assembly” on page 142).
10. Install the cover (see “Installing the cover” on page 124).
11. Push the power supplies back into the file module.
12. Slide the file module into the rack.

Note: After you replace the system board, the file module might reboot several times, and it might take up to an hour for the firmware to be updated on the new system board. Do not use the file module until the firmware update is completed.

13. Follow the steps at the end of the procedure “Removing a file module and disconnecting power” on page 92 to reconnect the file module and resume its use in the cluster.

Setting the machine serial number

This procedure is for IBM authorized service providers who, after replacing a system board in one of the file modules, must set the machine serial number vital product data stored on the system board using the Advanced Settings Utility (ASU). The following procedure is for a field replaceable unit (FRU). FRUs must be installed only by trained service technicians. The serial number must be set to the value that is visible on the MT-M SN label on the front of the file module.

About this task

The ASU package is part of the Storwize V7000 Unified code. ASU is available to authorized service personnel from the command-line interface (CLI) on the file module. Use ASU to modify selected settings in the integrated-management-module (IMM)-based Storwize V7000 Unified file modules.

Procedure

1. Access and log in to the file module with the new system board as root.
`ssh -p 1602 root@<file module IP>`

You are prompted for the file module root password.

2. To obtain the system type, run the following command on the other node:
`asu show SYSTEM_PROD_DATA.SysInfoProdName`
3. Issue the following command to view the current setting of the serial number:
`asu show SYSTEM_PROD_DATA.SysInfoSerialNum`
4. Issue the following ASU command on the Storwize V7000 Unified file module to set the serial number:

```
asu set SYSTEM_PROD_DATA.SysInfoSerialNum xxxxx
```

The variable `xxxxx` in the command stands for the serial number.

5. Issue the following command to verify that you set the serial number correctly:

```
asu show SYSTEM_PROD_DATA.SysInfoSerialNum
```

Check that the serial number is now set to the value that is visible on the MT-M SN label on the front of the file module.

How to reset/reboot server iMM interface

About this task

Use this procedure to initiate a reset/reboot of the iMM interface located on the file module. This action is not disruptive to the system operations and should only be used when directed to clear out fault conditions.

Note: This procedure requires root access to the file module node.

Procedure

1. Log in as a CLI user to the appropriate file module by using the service IP.
2. Type **`sc asu rebootimm --kcs`** and press **Enter**.

Note: If you are using a telnet connection, you can reboot using `resetsp`.

- a. Wait for the IMM reboot to complete (typically about 3 minutes). If the reboot is successful, the output of the previous command will be similar to the following:

```
IBM Advanced Settings Utility version 3.62.71B
Licensed Materials - Property of IBM
(C) Copyright IBM Corp. 2007-2010 All Rights Reserved
```

```
Try to connect to the primary node to get nodes number.
Connected via IPMI device driver (KCS interface)
Connected to primary node.
Nodes number is 1
Unable to locate a script required to set up LAN-over-USB device,
  tried location cdc_interface.sh
```

```
Connect to imm to reboot.
Issuing reset command to imm.
Checking if the imm has reset yet. (attempt 0)
imm has started the reset.
Disconnect from imm
```

- b. Wait for about 2 minutes to allow the iMM to completely reboot.

File module software problems

This section helps you to identify and resolve file module software problems.

About this task

Logical devices and physical port locations for a 2073-720 file module

Use this table to help identify logical devices, file module roles used, and physical locations on a 2073-720 file module.

Table 55. Default logical devices and physical port locations for a 2073-720 file module

Logical Ethernet device name	Device description	Physical location information
mgmtsl0_0	Internal connection between the file modules	Port 1 - Built-In xSeries Ethernet Port
mgmtsl0_1	Internal connection between the file modules	Port 2 - Built-In xSeries Ethernet Port
ethXsl0_0	1-Gbps Public Network	Port 3 - Built-In xSeries Ethernet Port
ethXsl0_1	1-Gbps Public Network	Port 4 - Built-In xSeries Ethernet Port
ethXsl1_0	1-Gbps Public Network	Port 7 - Quad Port 1-GB Ethernet adapter in PCI slot 2
ethXsl1_1	1-Gbps Public Network	Port 8 - Quad Port 1-GB Ethernet adapter in PCI slot 2
ethXsl1_2	1-Gbps Public Network	Port 9 - Quad Port 1-GB Ethernet adapter in PCI slot 2
ethXsl1_3	1-Gbps Public Network	Port 10 - Quad Port 1-GB Ethernet adapter in PCI slot 2
ethXsl2_0	10-Gbps Public network	Port 5 / 10-Gbps Ethernet adapter
ethXsl2_1	10-Gbps Public network	Port 6 / 10-Gbps Ethernet adapter

Note: The physical port locations on your system might differ from the port locations that are given in the preceding table if the port bonding has been changed.

Management node role failover procedures

The following procedures either restart the management service or initiate a management service failover from the file module hosting the active management node role to the file module hosting the passive management node role.

Once complete, the file module that previously hosted the active management node role now hosts the passive management node role. The file module that previously hosted the passive management node role now hosts the active management node role.

Note: All of these tasks require a user that is configured as a CLI admin. Other users cannot perform these tasks.

Determining the service IP for the management node roles

Use this procedure to identify the service IP addresses for the file modules that host the management node roles.

About this task

You need the service IP address of a file module that hosts a management node role to perform a management failover from the file module that hosts the active management node role to the file module that hosts the passive management node role, when the active management node fails and the current management IP does not respond.

Procedure

1. Connect to the CLI using SSH.

Note: Run the CLI command **lsnode**.

- If you get output from **lsnode** that shows the system configuration (as in Example 1), proceed to step 2.
- If you get a message that the management service is stopped or is not running (as in Example 2), attempt to log out and log in to the other file module hosting a management node role. If the other file module is not responding, refer to “Performing management node role failover procedures for failure conditions” on page 185.

Example 1: System configuration output from **lsnode** displays similar to the following example:

```
[admin@kq186wx.mgmt001st001 ~]# lsnode
Hostname      IP      Description      Role
mgmt001st001 172.31.8.2 active management node management,interface,storage
mgmt002st001 172.31.8.3 passive management node management,interface,storage

Product version Connection status GPFS status CTDB status Last updated
1.3.0.0-50a      OK      active      active      8/30/11 8:36 PM
1.3.0.0-50a      OK      active      active      8/30/11 8:36 PM

EFSSG1000I The command completed successfully.
[admin@kq186wx.mgmt001st001 ~]#
```

Example 2: Output for **lsnode** for a management service that is not running is similar to the following example:

```
[admin@kq186wx.mgmt002st001 ~]# lsnode
EFSSG0026I Cannot execute commands because Management Service is stopped.
Use startmgtsrv to restart the service.
```

2. Determine the service IP addresses for the file modules hosting a management node role by running the CLI command **lsnwmgt**. Output that is similar to the following example is displayed:

```
[admin@kq186wx.mgmt001st001 ~]# lsnwmgt
Interface Service IP Node1 Service IP Node2
ethX0      9.11.137.128      9.11.137.129

Management IP Network      Gateway      VLAN ID
9.11.137.127 255.255.254.0 9.11.136.1
EFSSG1000I The command completed successfully.
```

The following table describes the nodes that are identified by the command:

Table 56. Hostname and service IP reference

Host name	Corresponding Service IP reference
mgmt001st001	Service IP Node1
mgmt002st001	Service IP Node2

Performing management node role failover on a “good” system

Use this procedure to complete a failover process when both file modules appear to be operating correctly.

About this task

If both file modules are operating correctly with regard to management services, perform the following procedure to failover the active management node to the passive management node.

Procedure

1. Open an SSH connection to the service IP of the file module hosting the passive management node role.
Refer to “Determining the service IP for the management node roles” on page 183, if necessary.
2. To initiate the management services on the passive node and perform the switchover from the active management node, run the **startmgtsrv** command.

Note: If you run the **startmgtsrv** command from the node that is becoming active, you first need to run the **setcluster** command to set the cluster

environment variable. If you see the following error message when running the command, wait until the initialization has completed before running **setcluster** again:

```
IBM SONAS management service is starting up
EFSSG0654I The Management Service is starting up.
```

After you run the **startmgtsrv** command, the system displays information that is similar to the following example:

```
[yourlogon@yourmachine.mgmt002st001 ~]# startmgtsrv
Other node is reachable and its management state is active.
Are you sure? (Y/N)Y
EFSSG0717I Takeover initiated by root - this may take a few minutes
EFSSG0544I Takeover of the management functions from the active
node was successful
```

Results

Once complete, the file module that previously hosted the active management node role now hosts the passive management node role. The file module that previously hosted the passive management node role now hosts the active management node role.

Performing management node role failover procedures for failure conditions

Use this topic to isolate and perform file module failover for failed conditions.

About this task

“Failed conditions” exist when the active management node has failed and is not responding. This failure is exposed by the inability to access the file module, run CLI commands, and/or access the GUI.

Note: If the management IP is accessible and you can establish an SSH connection and run CLI tasks, do not perform a management failover. Refer to .

Complete the following procedure to address this issue.

Important: You need a CLI user with SystemAdmin privilege to perform this procedure. Performing this procedure does not repair a problem that caused the current system condition. This procedure provides for system access and troubleshooting to restart the management services or to failover the management service from a failed file module to the passive management node on the other file module. Once you complete this procedure, follow the appropriate troubleshooting documentation to isolate and repair the core problem that caused this condition.

Procedure

1. Attempt to open an SSH connection to the service IP of the file module with the active management node role. Refer to . Was the connection successful?
 - **Yes** - proceed to step 2
 - **No** - proceed to step 5 on page 186
2. If the connection is successful, verify that the management service is not running by executing the CLI command **lsnode** and then reviewing the output.
 - If the system responds with output for the **lsnode** command, then the management services are already running. If you still cannot access the GUI, refer to . If the GUI is accessible, then the management services are properly running on the active management node and no failover is needed. If you

want to initiate a failover, refer to “Performing management node role failover on a “good” system” on page 184.

- If the system responds that the management service is not running, proceed to the next step.

Note: For a management service that is not running, the system displays information similar to the following example:

```
[yourlogon@yourmachine.mgmt002st001 ~]# lsnode
EFSSG0026I Cannot execute commands because Management Service is stopped.
Use startmgtsrv to restart the service.
```

3. Attempt to stop and restart the management services. Wait for the commands to complete.
 - a. Run the CLI command **stopmgtsrv**.
 - b. Run the CLI command **startmgtsrv**. This restarts the management services.
4. Once command execution is complete:
 - a. Verify that the management service is running by again executing the CLI command **lsnode**. If the system responds that the management service is not running, proceed to step 5.
 - b. If the **lsnode** output provides system configuration information, verify that you can access and log in to the GUI. If you still have trouble with accessing the GUI, refer to .
 - c. If the problem appears to be resolved, DO NOT perform steps 5-9. Instead, using the GUI event log, follow the troubleshooting documentation to isolate the software or hardware problem that might have caused this issue.

Attention: Perform the following steps only if the active management node is not responding properly. These steps initiate a startup and failover of the management services on the file module hosting the passive management node role.

5. Open an SSH connection to the service IP and port of the file module with the passive management node role. Refer to “Determining the service IP for the management node roles” on page 183.
6. Verify the management service status by running the CLI command **lsnode**. If the file module responds that the management service is not running, proceed to the next step.
7. Run the CLI command **startmgtsrv**. This starts the management services on the passive node.
8. Once command execution is complete:
 - a. Verify that the management service is running by again executing the CLI command **lsnode**.
 - b. If the **lsnode** output provides system configuration information, verify that you can access and log in to the GUI. If you still have trouble with accessing the GUI, refer to .
 - c. If the **lsnode** output reports that the management service is still not running, contact IBM support.
9. Using the GUI event log, follow the troubleshooting documentation against the file module with the failed management node role to isolate the software or hardware problem that might have caused this issue.

Checking CTDB health

Use this information for checking the health of the clustered trivial database (CTDB) on each file module.

About this task

CTDB checks the health status of the Storwize V7000 Unified file modules, scanning elements such as storage access, General Parallel File System (GPFS), networking, Common Internet File System (CIFS) shares, and Network File System (NFS) exports.

An unhealthy file module cannot serve public Internet Protocol (IP) addresses and must be fixed. However, the high availability features of the Storwize V7000 Unified system can mask the unhealthy status from its clients by failing over IP addresses from an unhealthy file module to a healthy file module.

In the management graphical user interface (GUI), select **Monitoring > System** and check the health status for errors and degraded events.

Procedure

To check the status using either the GUI or the command-line interface (CLI), complete this procedure:

1. To check the CTDB status:
 - With the Storwize V7000 Unified GUI, use the following method:
Select **Monitoring > System Details > Interface Nodes > mgmt001st001 > NAS Services**. In the **CTDB state** row, a healthy status is displayed as “Active” and an unhealthy status is displayed as “unhealthy”.
A “disconnected” status is displayed when this CTDB node could not be reached through the network and is currently not participating in the cluster.

Review the status of both file modules, **1** mgmt001st001 and **2** mgmt002st001, as shown in Figure 72 on page 188.

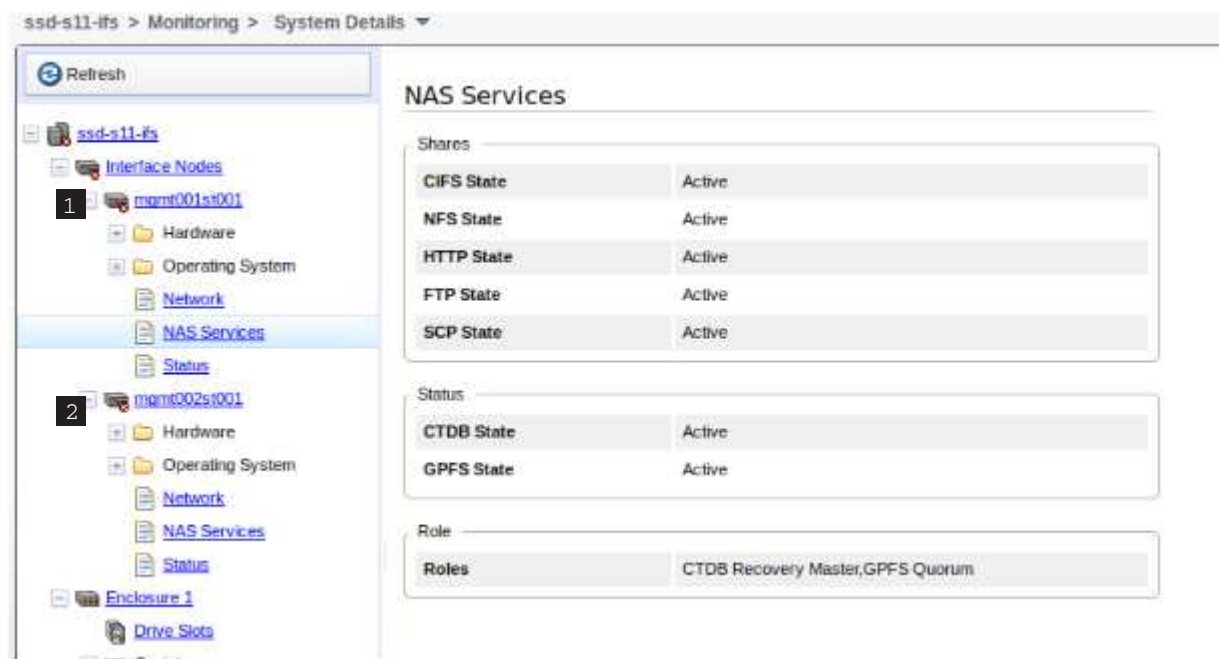


Figure 72. Management GUI showing CTDB status for both file modules

- With the CLI, log in as an admin user, then open the CLI and issue the **lsnode -r** command to determine whether CTDB is active on all nodes.

The system displays information similar to the following:

```
[your!ogon@yourmachine]$ lsnode -r
EFSSG0015I Refreshing data.
Hostname      IP           Description      Role
mgmt001st001  10.254.8.2   active management node management,interface,storage
mgmt002st001  10.254.8.3   passive management node management,interface,storage

Product version Connection status GPFS status CTDB status Last updated
1.3.0.0-55c      OK            active         active
1.3.0.0-55c      OK            active         unhealthy
EFSSG1000I The command completed successfully.
```

In the CTDB Status column, “active” indicates a healthy status and “unhealthy” indicates an error status. A “disconnected” status is displayed when this CTDB node could not be reached through the network and is currently not participating in the cluster.

2. If the CTDB status for a file module is not “active”, perform one or more of the following procedures:
 - Review the health status for any potential network problems. A network failure between a file module and the customer can result in an “UNHEALTHY” CTDB status. Follow the appropriate error code action plans to resolve the network problems. Refer to “Checking network interface availability” on page 195.
 - If the network has been examined without any problems being identified, perform the procedures in “Checking the GPFS file system mount on each file module” on page 189.
 - Refer to the information in “Troubleshooting the System x3650 server” topic in the *IBM Storwize V7000 Unified Information Center* to determine if any additional hardware problems might be causing the “unhealthy” CTDB status.

- Perform a reboot of the unhealthy file module. Refer to “Resuming services on a suspended file module” in the *IBM Storwize V7000 Unified Information Center*.
- If none of the above actions resolves the problem, contact IBM Remote Technical Support.

Checking the GPFS file system mount on each file module

Use this information to identify and resolve problems with General Parallel File System (GPFS) file system mounts on IBM Storwize V7000 Unified file modules.

About this task

A GPFS file system that is not mounted on an Storwize V7000 Unified file module can cause the clustered trivial database (CTDB) status to be 'UNHEALTHY'." The GPFS file system must be mounted on both file modules in the Storwize V7000 Unified product to support dual redundancy and to allow file input/output through all public IP addresses.

Note: You cannot change the cluster and file system configuration, when one of the nodes hosting the GPFS primary or secondary configuration server is not present in the cluster. You can identify the nodes hosting the primary or secondary configuration server in the output of the **lscluster** CLI command. This also applies to disk management operations such as creating or deleting a file system, adding or removing a disk using the CLI commands such as **mkfs**, **rmfs**, **chfs**, and **rpldisk**. It also includes the changes to the cluster configuration such as **addnode** and **delnode**. These CLI commands report an error and fail the operations when this condition is detected in the cluster.

Use the following procedure to get a file system that is not mounted on a file module to be mounted. Contact IBM Remote Technical Support if this procedure does not work.

Procedure

- To identify all of the currently created file systems on the Storwize V7000 Unified system, or on each file module, perform the procedure in “Identifying created and mounted file system mounts.”
- To resolve problems with mounted file systems that are missing, perform the procedure in “Resolving problems with missing mounted file systems” on page 190.
- To resolve problems with stale NFS file systems, perform the procedure in “Resolving stale NFS file systems” on page 191.
- To resolve problems that are not covered by the information that is presented in the previous topics, perform the procedure in “Recovering a GPFS file system” on page 196.

Identifying created and mounted file system mounts

You can identify and resolve problems in GPFS file system mounts on the Storwize V7000 Unified system and file modules.

About this task

Procedure

To identify and resolve problems in file system mounts, perform this procedure:

1. To identify all the currently created file systems on the Storwize V7000 Unified system, log in as the admin user, then enter the **lsfs -r** command from the command-line interface (CLI), as shown in the following example:

```
# lsfs -r

EFSSG0015I Refreshing data.
Cluster   Device name Quota           Def. quota Block size Inodes
kd18pz5.ibm gpfs1      user;group;fileset      256 kB    11373

Replication Dmapi Block allocation type Last update
none        yes  scatter                10/3/11 2:08 PM

EFSSG1000I The command completed successfully.
```

2. To verify the mount state of the currently created file systems on each Storwize V7000 Unified file module, enter the **lsmount -v** command from the CLI, as shown in the following example:

```
# onnode -n mgmt001st001 df | grep ibm

/dev/gpfs1      3221225472  4590080 3216635392  1% /ibm/gpfs1
```

Repeat the command for another file module by using the **onnode -n mgmt002st001 df | grep ibm** command, for example:

```
# onnode -n mgmt002st001 df | grep ibm

/dev/gpfs1      3221225472  4590080 3216635392  1% /ibm/gpfs1
```

Resolving problems with missing mounted file systems

You can resolve problems with mounted file systems that are missing on Storwize V7000 Unified file modules.

About this task

Display the file system by using the **lsfs -r** command. The **lsmount -r -v** command shows which file modules mount the file system. The Mounted status means that both file modules mount the file system. All other states, Partially, Internally, or Not mounted mean that a file system is not properly mounted.

Procedure

To resolve the problem with the missing mounted file system, perform the following procedure:

1. Log in to the Storwize V7000 Unified CLI as admin.
2. Identify on which file module the file system is missing, for example, mgmt001st001.
3. Mount the missing file system on the file module by using the **mountfs** command.

```
mountfs gpfs0
```
4. Issue the **lsmount** command to verify that all file systems are now mounted on file modules 0 and 1.
5. If the mounted file systems are not consistent across the file modules, reboot the file module on which a file system is missing, and then issue the **lsmount** command.

Reboot the file module by using the management GUI.

If they are not mounted on either file module, then reboot each file module.

6. Use the **lsnode** command to determine when the file modules are back up and when GPFS and CTDB are both active.

The file systems might take several minutes to get mounted after GPFS becomes active on both file modules. More than one reboot might be required to bring the file system back up. Allow time between reboots because the file system might take some time before it comes back up after a reboot.

7. Ensure that the CTDB status is now shown as **active** on both file modules, as described in “Checking CTDB health” on page 187.
8. If a GPFS file system fails to mount, complete the following steps:
 - a. Check the output log of the **lslog** command and look for the latest messages about mounting the file system.

If you find input/output errors and messages about a failure to read the super block, the problem is with the DMAPI clients of the IBM Spectrum Protect™ or HSM system.

Check for disk-related problems, such as errors reading from a disk or errors showing a non-existent disk. For these errors, check whether the path to the storage system is working. If it is, verify that the system itself is in working order.

- b. For additional information, refer to the “Diagnostics: Troubleshooting tables” information in “Troubleshooting the System x3650” in the *IBM Storwize V7000 Unified Information Center*.
 - c. If file systems remain unmounted, contact IBM support.

Resolving stale NFS file systems

You can resolve problems with stale NFS file systems on Storwize V7000 Unified file modules. A file module might have the file system mounted, but the file system remains inaccessible due to a stale NFS file handle.

About this task

Procedure

To identify and resolve stale file handle problems, complete this procedure:

1. To identify stale NFS file handle problems, log in to the CLI as privileged user and enter the **sc onnode all df | grep NFS** command:

```
# onnode all df | grep NFS

df: ~/ibm/gpfs0': Stale NFS file handle
```
2. If the command reports a stale NFS file handle on a particular file system, see “Working with file modules that report a stale NFS file handle” on page 403 for instructions on file system recovery.

Checking user and server authentication issues

Due the authentication failure, the users cannot log in to the system or a connection cannot be established between the servers.

About this task

For login issues, if you are sure that you have used the correct user ID and password, your user account might have been deleted or corrupted.

Refer to these topics in the *IBM Storwize V7000 Unified Information Center* “Planning for user authentication”, “Verifying the authentication configuration”, “Establishing user and group mapping for client access”, and “chkauth”.

If you cannot resolve the issue, contact the authentication server administrator to validate or reestablish your account.

Refer to “Managing authentication server integration” for more information about authentication and server configuration.

Resolving the “Missing SRV record in DNS” error

About this task

If the “Missing SRV record in DNS” error displays when you configure the active directory (AD) using the **cfgad** command, similar to the following example, verify that entries for DNS Domain Name, DNS Server, and DNS Search Domains are correct. Also, verify that the DNS server has valid SRV records for that domain.

```
$ cfgad -s 9.9.9.9 -u admin -p ****
(1/9) Fetching the list of cluster file modules.
(2/9) Check if cfgcluster has done the basic configuration successfully.
(3/9) Check whether file modules are
reachable from management file module.
(4/9) Detection of AD server and fetching domain information from AD server.
Missing SRV record in DNS : _ldap._tcp.xxxxx.COM
Missing SRV record in DNS : _ldap._tcp.dc._msdcs.xxxxx.COM
Missing SRV record in DNS : _kerberos._tcp.xxxxx.COM
Missing SRV record in DNS : _kerberos._tcp.dc._msdcs.xxxxx.COM
Necessary DNS entries are missing, the domain join step might fail.
(5/9) Check whether AD server is reachable
from file modules.
(6/9) Joining the domain of the specified ADS.
EFSSG0110C Configure AD failed on cluster. Cause: Error encountered while
executing netjoinAD.sh. Output till failure is :Join to Active Directory
domain with user Administrator
Failed to join domain: failed to find DC for domain SONAS
Error occurred due to reason : Join to Active Directory domain failed
```

If “netgroup” functionality with NIS or LDAP is not working

About this task

If “netgroup” functionality with Network Information Service (NIS) or Lightweight Directory Access Protocol (LDAP) is not working, ensure that you have included a “@” in front of the netgroup name, as shown in the following example:

```
$ mkexport testnetgrp5 /ibm/gpfs0/netgroup5 --nfs "@ng1(rw,no_root_squash)"
```

Do not create a netgroup with an IP address; instead, use a host name. The host name that is defined in a netgroup should resolve to a valid IP address that points back to the same host name when you query for it.

Possible misconfiguration on the Storwize V7000 Unified system

About this task

Authentication problems might be caused by a client-side NAS misconfiguration. To verify, issue the **lookupname** command on the active management file module, as shown in the following example, to verify that the file module can authenticate with the authentication server.

```
$ lookupname --user SONAS\\userr
USER          GROUP
SONAS\\userr  SONAS\\domain users
EFSSG1000I The command completed successfully.

$ chkauth -i -u SONAS\\userr
Command Output Data  UID      GID      Home_Directory      Template_Shell
FETCH USER INFO SUCCEED 12004360 12000513 /var/opt/IBM/sofs/scproot /usr/bin/rssh
EFSSG1000I The command completed successfully.
```

When the system is unable to authenticate against an external authentication server, you must ensure that it can obtain user information from the authentication server. For this user information, query commands can be run from the file modules. For example, in the case of the LDAP authentication server, you can issue a command as shown in the following example:

```
$ chkauth -a -u SONAS\\userr -p *****
AUTHENTICATE USER SUCCEED
EFSSG1000I The command completed successfully.
```

Trouble accessing exports when authentication server and clientStorwize V7000 Unified configurations are correct

About this task

If you cannot access an export and the server and Storwize V7000 Unified configurations are correct, it could be because of the following reasons.

- If Storwize V7000 Unified authentication is configured against an LDAP server, the user entries are case-sensitive when you access exports. If the server and client configurations are correct, ensure that the user entries have the correct case.
- If Storwize V7000 Unified authentication is configured against an Active Directory server, user entries are not case-sensitive when you access exports. When you access CIFS exports, ensure that you use the domain name and user name, separated by a backslash (\), for example, w2k3dom01\test1.

Resolving access failures on an Storwize V7000 Unified system with a subordinate ID map role

About this task

When two or more Storwize V7000 Unified systems are in asynchronous replication or remote caching relationship, and they both have authentication as **AD** and ID mapping as **auto**, then one system can be given the ID map role as master and another system is given the ID map role as subordinate. In this configuration, if a user cannot access data on the subordinate system (or if **chkauth -i** CLI command fails on the subordinate Storwize V7000 Unified system), then use the following steps to troubleshoot the issue.

Note:

- Use the following steps only when the systems are under asynchronous replication or remote caching relationship with **AD** authentication, **auto** ID mapping, and ID map role as master or subordinate. If systems are using LDAP authentication or AD authentication with SFU ID mapping, then these steps are not applicable.
- If there are multiple systems in the environment with all using **AD** authentication and **auto** ID mapping, and the systems share asynchronous replication or remote

caching relationship, then there must be only one system in the environment whose ID map role is master and rest all systems must have ID map roles as subordinate.

Procedure

1. Access the master Storwize V7000 Unified system with the same user (with whom the access was failing on the subordinate Storwize V7000 Unified system). You can also run the **chkauth -i** command on the master system.

Note: Accessing the master Storwize V7000 Unified system (or **chkauth -i** command) with a user creates the ID mapping information on the system. Although, any user from the domain can be used to create the ID maps, it is advisable (refer to idMapConfig option of **cfgad** CLI command) to use the same user.

2. Check whether the ID map is created. Use **lsidmap** CLI command to display the ID maps that also contains information about the domain name, SID, and the ID range.
3. Export the ID mapping information from the master system into an XML file, which is stored in the /ftdc folder of the active management node.

Note: Use the **cfgidmap** CLI command to export ID mapping information.

4. Import the ID mapping information in the Storwize V7000 Unified subordinate system from the XML file.

Note: You can use the **scp** command to transfer the XML file from master Storwize V7000 Unified system to the subordinate Storwize V7000 Unified system. Use the **cfgidmap** command to import the ID map XML file. The XML file must be at the /ftdc/files folder.

5. Try to access the data on the subordinate Storwize V7000 Unified system after the import operation is successfully completed.

Checking client access

Verify that your client workstation can successfully ping the full host name of the cluster and all of the IP addresses that are associated with it.

About this task

The following example shows how to ping an cluster. When the client connects to the host name of the cluster, the DNS server responds with IP addresses. You must then ping each IP address from the client machine.

If clients cannot successfully ping the IP addresses, then they are not able to access Storwize V7000 Unified whenever the DNS returns the IP address on name resolution requests. This can cause some clients have access while others do not.

Procedure

1. To obtain the IP addresses of your Storwize V7000 Unified cluster, issue the **nslookup** command. .

Information similar to the following example is displayed:

```
# nslookup yourdomainname
Server:          9.11.136.116
Address:         9.11.136.116#53
```

Non-authoritative answer:

```
Name: yourdomainname
Address: 129.42.16.103
Name: yourdomainname
Address: 129.42.17.103
Name: yourdomainname
Address: 129.42.18.103
```

The **nslookup** command returns the IP addresses (129.42.18.103 in the example above) that are configured on the DNS server for Storwize V7000 Unified. Ideally, these IP addresses should be the same as the addresses that are configured on the Storwize V7000 Unified cluster itself. To check this, issue the **lsnw** CLI command.

2. Ping each IP address that is listed in the output by issuing the **ping returned IP Address** command. A successful return indicates a working connection. The response Request timed out indicates a failed connection.

Note: If clients cannot ping the IP addresses, refer to “Checking network interface availability.”

3. If you have a failed connection, contact the system administrator or IBM Remote Technical Support.

Checking network interface availability

You have several options for checking network availability by using the Storwize V7000 Unified GUI or the CLI.

Procedure

1. In the GUI, select **Monitoring > System Details > mgmt00xst001 > Network**.
2. In the CLI, check the status of the interface “ethX0” (the interface of file modules to the customer net).
 - a. Open the CLI.
 - b. Issue the **lsnwinterface** command to display the status for the desired IP addresses.

```
# lsnwinterface
```

The system displays information similar to the following example:

Node	Interface	MAC	Master/Subordinate	Bonding mode
mgmt001st001	ethX0	e4:1f:13:d6:ae:ac	MASTER	balance-alb (6)
mgmt001st001	ethX1	00:c0:dd:17:bc:ac	MASTER	active-backup (1)
mgmt002st001	ethX0	e4:1f:13:d6:ae:94	MASTER	balance-alb (6)
mgmt002st001	ethX1	00:c0:dd:17:c5:50	MASTER	active-backup (1)

Up/Down	Speed	IP-Addresses	MTU
UP	1000		1500
UP	10000	9.11.84.84,9.11.84.85	1500
UP	1000		1500
UP	10000	9.11.84.82,9.11.84.83	1500

EFSSG1000I The command completed successfully.

In the **Up/Down** column, the value UP indicates a connection.

3. If the network interface is not available, check the cables and ensure that the cables are plugged in. For instance, if you have no machine connectivity between file modules and switches, check the external Ethernet cabling. If all cables are correctly connected, check intranet and external Internet availability. If none of these checks indicate a problem, contact the next level of support.

Recovering a GPFS file system

Use this procedure to recover a GPFS file system after a storage system failure has been fully addressed. You should use this procedure only under the supervision of IBM support.

Before you begin

Prerequisites:

- You are running this procedure on a file module.
- You are logged on to the CLI as member of the “SystemAdmin” group.
- GPFS must be active on both nodes.

For storage system recovery, see the procedure for recovering a storage system.

Restriction:

To recover the filesystem for reasons other than storage node failure, for example, to recover from multiple disks or storage controller failures, do not use GPFS root commands without consulting with the next level of IBM support. Using GPFS root commands to recover filesystems might result in further filesystem damage.

About this task

This procedure provides steps to recover a GPFS file system after a failure of the block storage system. The file volumes were offline and are now back online after a repair or recovery action. The disks referred to in this procedure are the volumes that are provided by the block storage system.

Note: Because no I/O can be done by GPFS, it is assumed for these procedures that the storage unit failure caused the GPFS file system to unmount. After satisfying the prerequisites above, take the following steps:

Procedure

1. Verify that GPFS is running on both file modules by using the **lsnode -r** command.

The column **GPFS status** shows active.

2. With GPFS functioning normally on both file modules, ensure that all disks in the file system are available by running the **lsdisk -r** command. The **Availability** column shows Up.

3. Issue the **chkfs file_system_name -v | tee /ftdc/chkfs_fs_name.log1** command to capture the output to a file.

Review the output file for errors and save it for IBM support to investigate any problems.

If the file contains a TSM ERROR message, perform the following steps:

- a. Issue the **stopbackup -d file_system_name** command and the **stoprestore -d file_system_name** command to stop any backup or restore operation.
- b. Validate that no error occurred while stopping any IBM Spectrum Protect service.
- c. Issue the **chkfs file_system_name -v | tee /ftdc/chkfs_fs_name.log2** command to recapture the output to a file.
- d. Issue the **startrestore** command and the **startbackup** command to enable IBM Spectrum Protect.

If you receive an error message (the number of mounted or used file modules does not matter) at step 5 of the command internal execution steps like the following,

```
(5/9) Performing mmfsck call for the file system check stderr:
Cannot check. "gpfs0" is mounted on 1 node(s) and in use on 1 node(s).
mmfsck: Command failed.
Examine previous error messages to determine cause.
```

perform the following steps:

- a. Monitor the **lsmount -r** command until the mount status changes to not mounted.
- b. Issue the **chkfs file_system_name** command again.

Review the new output file for errors and save it for IBM support to investigate any problems. It is expected that the file contains Lost blocks were found messages. It is normal to have some missing file system blocks. If the only errors that are reported are missing blocks, no further repair is needed. However, if the **chkfs** command reports more severe errors, contact IBM support to assist with repairing the file system.

Resolving an ANS1267E error

An ANS1267E error might indicate an incorrect setting in the IBM Spectrum Protect server configuration.

About this task

An ANS1267E error can result from the IBM Spectrum Protect server not being set up to handle hierarchical storage management (HSM) migrated files and that the management class is not accepting files from HSM.

To correct this error, set the **spacemgtech** value to "auto".

Resolving issues reported by lshhealth

Use this information to resolve **lshhealth** reported issues, specifically for "MGMTNODE_REPL_STATE ERROR DATABASE_REPLICATION_FAILED" and "The mount state of the file system /ibm/Filesystem_Name changed to error level" errors.

About this task

These errors might be transient and can clear automatically at any time.

Error for "MGMTNODE_REPL_STATE ERROR DATABASE_REPLICATION_FAILED"

About this task

To resolve the "MGMTNODE_REPL_STATE ERROR DATABASE_REPLICATION_FAILED" error, complete the following steps.

Procedure

1. Verify that the other file module role displays Host State OK. Repair the host state if necessary.

2. Allow fifteen minutes for the error to disappear. If the error does not disappear, attempt to reboot the passive management node. The issue should be resolved after the reboot and within five minutes after the file module displays Host State OK again.

Error for “The mount state of the file system /ibm/ Filesystem_Name changed to error level”

About this task

If the command `lshealth -i gpfs_fs -r` returns “The mount state of the file system /ibm/Filesystem_Name changed to error level”, complete the following steps to resolve the issue.

Procedure

1. Verify that the other file module role displays Host State OK. Repair the host state if necessary.
2. Issue the command `mountfs fileSystem`.
3. Issue the command `lsfs -r`.
4. Issue the command `lshealth -i gpfs_fs -r`.
The command output should display The mount state of the file system /ibm/gpfs1 was set back to normal level.
5. If the error persists, refer to the GPFS documentation to debug or repair the error.

Resolving network errors

Use the following information and examples to resolve network errors that are identified by the health system.

If a network is not attached to an interface, the health center monitors all ports and logs and displays any port failures.

To stop monitoring an unused port, use the `attachnw` command to attach the interface that corresponds to that port. After this command is issued, you must manually mark the displayed error events for the unused port as *resolved*. You can use the GUI system details panel to manually mark the events.

Whether a port is monitored depends on the attached network configuration and the interface that is used for the management network. If an interface is not in use by an attached network, the interface is monitored if it is used by the management network.

Important:

It is highly recommended, that for a given VLAN subnet definition, you must also consistently provide the same VLAN ID (tag) for that subnet on any network definition that you create within a common interface bond on the clustered system. If you define a VLAN on your management network, and you have the same VLAN subnet on the same interface bond for data connectivity, ensure that you provide the exact same VLAN ID. Additionally, the switch configurations supporting this VLAN should be same for all connections that the VLAN uses. Ensure that you use trunk, access link, or native VLAN consistently on the switch ports that are connected to all ports on the clustered system. The clustered system shares the same VLAN subnet. Avoid defining the same VLAN with a tag ID

either different or missing, when providing the VLAN ID to a management network bond, as well as to a shared data network bond.

Unless you have intentionally configured your switching network to support this unique case, where the VLAN ID for the management network and data network are not the same and you are confident on how this will be routed from the clustered system to your switch, you might incur unpredictable routing path behavior and even network connectivity loss.

See the following examples.

Scenario 1

```
lsnwgrouper returns:
[admin@kd66t4v.mgmt001st001 ~]# lsnwgrouper
Network Group Nodes          Interfaces
DEFAULT      mgmt001st001,mgmt002st001
EFSSG1000I The command completed successfully.
and lsnwmgt returns:
[admin@kd52v6k.mgmt001st001 ~]# lsnwmgt
Interface Service IP Node1 Service IP Node2 Management IP Network Gateway VLAN ID
ethX0      . . .      . . .      . . .      . . .      . . .
EFSSG1000I The command completed successfully.
```

None of the interfaces is attached and the management network uses interface ethX0. If any ethX1 port cable is unplugged, the health center displays a failure because no network is attached to an interface, which causes the system to monitor all ports.

Scenario 2

```
lsnwgrouper returns:
[admin@kd52v6k.mgmt001st001 ~]# lsnwgrouper
Network Group Nodes          Interfaces
DEFAULT      mgmt001st001,mgmt002st001 ethX0
EFSSG1000I The command completed successfully.
and lsnwmgt returns:
[admin@kd52v6k.mgmt001st001 ~]# lsnwmgt
Interface Service IP Node1 Service IP Node2 Management IP Network Gateway VLAN ID
ethX0      . . .      . . .      . . .      . . .      . . .
EFSSG1000I The command completed successfully.
```

The interface ethX0 is attached to a network and the management network uses ethX0. If any ethX1 port cable is unplugged, the health center does not display a failure, because ethX1 is not used by either an attached network or the management network.

Scenario 3

```
lsnwgrouper returns:
[admin@kd52v6k.mgmt001st001 ~]# lsnwgrouper
Network Group Nodes          Interfaces
DEFAULT      mgmt001st001,mgmt002st001 ethX0
EFSSG1000I The command completed successfully.
and lsnwmgt returns:
[admin@kd52v6k.mgmt001st001 ~]# lsnwmgt
Interface Service IP Node1 Service IP Node2 Management IP Network Gateway VLAN ID
ethX1      . . .      . . .      . . .      . . .      . . .
EFSSG1000I The command completed successfully.
```

If any ethX1 port cable is unplugged, the health center displays a failure because ethX1 is used by the management network.

Resolving full condition for GPFS file system

Use this procedure when the management GUI reports a critical error when a file system is 100% full.

About this task

You must have root access to perform this procedure.

Note: If you use GPFS snapshots, the file system locks when it reaches 100% utilization.

Procedure

To resolve the full condition for the file system, perform the following steps:

1. Review the contents of the GPFS file system.
 - If the file system has snapshots, remove the oldest snapshot after verifying that it is no longer needed. Continue to remove the snapshots from oldest to newest until the level of free space that you want is achieved.
 - If no snapshots exist, perform the following steps:
 - a. Run the **lspool** command to determine what storage pool is out of space.
 - b. Remove files to free up storage.
 - c. If the **mmdf** command output shows that there is space in free fragments, run the **sc mmdefragfs** command to combine the fragments into full blocks.

Note: You can run the GPFS **defrag** command while the file systems are mounted. However, for better results, unmount the GPFS file system before performing the defragmentation operation.

2. If there is no space in fragments or if the **sc mmdefragfs** command does not free up space, add disks (NSDs) to the file system to create space.
 - a. Add disks to the file system.

Note: If free space exists in the **mdiskgroup** then you can modify the file system by editing it in the GUI or simply running the command: **mkdisk fileSystem size mdiskgroup**

For example:

```
[root@kd01gln.mgmt002st001 ~]# mkdisk gpfs0 10GB 0
(1/4) Creating Storage System volumes
(2/4) Scanning for new devices
(3/4) Creating NSDs
(4/4) Adding disks to filesystem
Successfully created disk
```

- b. If there is no storage space available, contact IBM support.

Analyzing GPFS logs

Use this procedure when reviewing GPFS log entries.

About this task

Note: Contact IBM support if you want to analyze GPFS log entries.

Procedure

1. Log in as a CLI user to the appropriate file module using the service IP.
2. Review the log file `/var/adm/ras/mmfs.log.latest`. The details in the log are listed from oldest to newest, so you can find the latest GPFS information at the end.

Note: The GPFS log is a complex raw log file for GPFS. If you do not understand the conditions listed in the log, contact IBM support for assistance.

Synchronizing time on the file modules

Use this information to synchronize the time on all Storwize V7000 Unified file module.

About this task

Synchronizing the time on all the file module can help as you start troubleshooting because the timestamps on the logs then indicate whether you have concurrent, legitimate results.

You can ensure that the Storwize V7000 Unified, Active Directory (AD), Kerberos, and other servers are synchronized with a valid Network Time Protocol (NTP) source. This is important both for log checking and because if the cluster falls behind the correct time, Kerberos tickets, for example, can expire and then no one can access the cluster. For the Storwize V7000 Unified file module, the **ntpq -p** command shows you which server is used for synchronization and any peers and a set of data about their status. The * in the first column indicates that the local clock is used for synchronization.

```
# ntpq -p
      remote           refid      st t when poll reach  delay  offset  jitter
=====
*machine.domain.i 9.19.0.220  2 u  269 1024  377    0.659   -0.115   0.164
+machine.domain.i 9.19.0.220  2 u  992 1024  377    1.380    0.337   0.564
LOCAL(0)         .LOCL.      10 l   50   64  377    0.000    0.000   0.001
```

As NTP is drift based, large time differences can prevent NTP from synchronizing, or cause synchronization to take a long time. It can be helpful to synchronize time manually once and to verify that the time is picked up correctly afterward. Use the separate commands of **service ntpd stop**, **ntpdate *your IP***, and **service ntpd start**. The following example shows the sequence:

```
[root@domain.node ~]# service ntpd stop
Shutting down ntpd: [ OK ]
[root@domain.node ~]# ntpdate 9.19.0.220
14 Jan 12:06:46 ntpdate[25360]: adjust time server 9.19.0.220 offset 0.003277 sec
[root@domain.node ~]# service ntpd start
Starting ntpd: [ OK ]
[root@domain.node ~]#
```

After the time on all of the servers is synchronized, you can verify that the logs apply to your troubleshooting situation.

Chapter 5. Control enclosure

Find out how to troubleshoot the control enclosure, which includes the use of error codes, problem scenarios, software, and removal and replacement instructions.

About this task

Storwize V7000 system interfaces

The Storwize V7000 system provides a number of user interfaces to troubleshoot, recover, or maintain your system. The interfaces provide various sets of facilities to help resolve situations that you might encounter. The interfaces for servicing your system connect through the 1 Gbps Ethernet ports that are accessible from port 1 of each canister or by directly connecting to the technician port of a node canister. You cannot manage a system by using the 10 Gbps Ethernet ports.

You can perform almost all of the configuration, troubleshooting, recovery, and maintenance of the storage system from within the Storwize V7000 Unified management GUI or the CLI commands that are running on the Storwize V7000 file modules.

Attention: Do not use the Storwize V7000 system interfaces directly unless you are directed to do so by a service procedure.

Use the initialization tool to do the initial setup of your system. Use the Storwize V7000 Unified management GUI or the Storwize V7000 system management GUI to monitor and maintain the configuration of storage that is associated with your systems. Perform service procedures from the service assistant. Use the command-line interface (CLI) to manage your system.

Service assistant interface

The service assistant interface is a browser-based GUI that is used to service individual node canisters in the control enclosures.

You connect to the service assistant on one node canister through the service IP address. If there is a working communications path between the node canisters, you can view status information and perform service tasks on the other node canister by making the other node canister the current node. You do not have to reconnect to the other node.

When to use the service assistant

The primary use of the service assistant is when a node canister in the control enclosure is in service state. The node canister cannot be active as part of a system while it is in service state.

Attention: Complete service actions on node canisters only when directed to do so by the fix procedures. If used inappropriately, the service actions that are available through the service assistant can cause loss of access to data or even data loss.

The node canister might be in a service state because it has a hardware issue, has corrupted data, or has lost its configuration data.

Use the service assistant in the following situations:

- When you cannot access the system from the management GUI and you cannot access the Storwize V7000 Unified to run the recommended actions
- When the recommended action directs you to use the service assistant.

The storage system management GUI operates only when there is an online system. Use the service assistant if you are unable to create a system or if both node canisters in a control enclosure are in service state.

The service assistant does not provide any facilities to help you service expansion enclosures. Always service the expansion enclosures by using the management GUI.

The service assistant provides detailed status and error summaries, and the ability to modify the World Wide Node Name (WWNN) for each node.

You can also complete the following service-related actions:

- Collect logs to create and download a package of files to send to support personnel.
- Remove the data for the system from a node.
- Recover a system if it fails.
- Install a code package from the support site or rescue the code from another node.
- Update code on node canisters manually versus completing a standard update procedure.
- Configure a control enclosure chassis after replacement.
- Change the service IP address that is assigned to Ethernet port 1 for the current node canister.
- Install a temporary SSH key if a key is not installed and CLI access is required.
- Restart the services used by the system.

The service assistant completes a number of tasks that cause the node canister to restart. It is not possible to maintain the service assistant connection to the node canister when it restarts. If the current node canister on which the tasks are completed is also the node canister that the browser is connected to and you lose your connection, reconnect and log on to the service assistant again after running the tasks.

Accessing the service assistant

The service assistant is a web application that helps troubleshoot and resolve problems on a node canister in a control enclosure.

About this task

You must use a supported web browser. For a list of supported browsers, refer to the topic Web browser requirements to access the management GUI.

Procedure

To start the application, complete the following steps.

1. Start a supported web browser and point your web browser to *serviceaddress/service* for the node canister that you want to work on.

For example, if you set a service address of 11.22.33.44 for a node canister, point your browser to 11.22.33.44/service. If you are unable to connect to the service assistant, see “Problem: Cannot connect to the service assistant” on page 245.

2. Log on to the service assistant using the superuser password.

If you are accessing a new node canister, the default password is `passwd`. If the node canister is a member of a system or was a member of a system, use the password for the superuser password.

If you do not know the current superuser password, try to find out. If you cannot find out what the password is, reset the password. Go to “Procedure: Resetting the Storwize V7000 Gen1 superuser password” on page 251.

Results

Complete the service assistant actions on the correct node canister. If you did not connect to the node canister that you wanted to work on, access the **Change Node** panel from the home page to select a different current node.

Commands are run on the current node. The current node might not be the node canister that you connected to. The current node identification is shown on the left at the top of the service assistant screen. The identification includes the enclosure serial number, the slot location, and if it has one, the node name of the current node.

The Storwize V7000 Gen2 technician port:

The technician port provides a convenient, direct connection to a node canister for servicing.

On uninitialized systems, the technician port provides access to the system initialization wizard instead of the service assistant. An uninitialized system is one where all node canisters have the green power LED on, the green status LED blinking, and amber fault LED is off.

Once a system has been initialized, the technician port provides access to:

- The service assistant
- The password reset facility (if enabled)

Storage system command-line interface

Use the storage system command-line interface (CLI) to manage a storage system by using the task commands and information commands.

You can also access most of the storage system CLI commands from the Storwize V7000 Unified CLI that runs in the file system on one of the file modules.

For a full description of the storage system commands and how to start an SSH command-line session, see the “Command-line interface” topic in the “Reference” section of the Storwize V7000 Unified Information Center.

When to use the storage system CLI

The storage system CLI is intended for use by advanced users who are confident at using a command-line interface.

Nearly all of the flexibility that is offered by the CLI is available through the management GUI. However, the CLI does not provide the fix procedures that are available in the management GUI. Therefore, use the fix procedures in the management GUI to resolve the problems. Use the CLI when you require a configuration setting that is unavailable in the management GUI.

You might also find it useful to create command scripts by using the CLI commands to monitor for certain conditions or to automate configuration changes that you make on a regular basis.

Accessing the storage system CLI

Follow the steps that are described in the “Command-line interface” topic in the “Reference” section of the Storwize V7000 Unified Information Center to initialize and use a CLI session.

Service command-line interface

Use the service command-line interface (CLI) to manage a node canister in a control enclosure by using the task commands and information commands.

Note: The service command line interface can also be accessed by using the technician port.

For a full description of the commands and how to start an SSH command line session, see the “Command line interface” topic in the “Reference” section of this product information.

When to use the service CLI

The service CLI is intended for use by advanced users who are confident at using a command-line interface.

To access a node canister directly, it is normally easier to use the service assistant with its graphical interface and extensive help facilities.

Accessing the service CLI

To initialize and use a CLI session, complete the steps in the Command-line interface topic in the Reference section of this product information.

USB flash drive and Initialization tool interface

Use a USB flash drive to initialize a system and also to help service the node canisters in a control enclosure.

The initialization tool is a Windows application. Use the initialization tool to set up the USB flash drive to perform the most common tasks.

When a USB flash drive is inserted into one of the USB ports on a node canister in a control enclosure, the software searches for a control file on the USB flash drive and runs the command that is specified in the file. When the command completes, the command results and node status information are written to the USB flash drive.

When to use the USB flash drive

The USB flash drive is normally used to initialize the configuration after installing a new system; however, it can be used for other functions.

Using the USB flash drive is required in the following situations:

- When you cannot connect to a node canister in a control enclosure using the service assistant and you want to see the status of the node.
- When you do not know, or cannot use, the service IP address for the node canister in the control enclosure and must set the address.
- When you have forgotten the superuser password and must reset the password.

Using a USB flash drive

Use any USB flash drive that is formatted with a FAT32 file system on its first partition.

About this task

When a USB flash drive is plugged into a node canister, the node canister code searches for a text file named `satask.txt` in the root directory. If the code finds the file, it attempts to run a command that is specified in the file. When the command completes, a file called `satask_result.html` is written to the root directory of the USB flash drive. If this file does not exist, it is created. If it exists, the data is inserted at the start of the file. The file contains the details and results of the command that was run and the status and the configuration information from the node canister. The status and configuration information matches the detail that is shown on the service assistant home page panels.

The `satask.txt` file can be created on any workstation by using a text editor. If a Microsoft Windows workstation is being used, the initialization tool can be used to create the commands that are most often used.

The fault light-emitting diode (LED) on the node canister flashes when the USB service action is being performed. When the fault LED stops flashing, it is safe to remove the USB flash drive.

Results

The USB flash drive can then be plugged into a workstation and the `satask_result.html` file viewed in a web browser.

To protect from accidentally running the same command again, the `satask.txt` file is deleted after it has been read.

If no `satask.txt` file is found on the USB flash drive, the result file is still created, if necessary, and the status and configuration data is written to it.

Using the initialization tool

The initialization tool is a graphical user interface (GUI) wizard that creates the initial configuration on the control enclosure.

Before you begin

Access the initialization tool through a USB flash drive.

Verify that you are using a supported operating system. The initialization tool is valid for the following operating systems:

- Microsoft Windows 8.1 (64-bit), or Microsoft Windows 7 (64-bit)

About this task

By using the initialization tool, you can set the USB flash drive to run one of the following tasks:

- Initialize a new system.
- Reset the superuser password.
- Set or reset the service assistant IP address on a node canister on the control enclosure.
- Set the management IP addresses.

For any other tasks that you want to perform on a node canister on the control enclosure, you must create the `satask.txt` file using a text editor.

The initialization tool is available on the USB flash drive that is shipped with the control enclosures. The name of the application file is `InitTool.exe`. If you cannot locate the USB flash drive, you can download the application from the following support website. (Search for “initialization tool.”)

www.ibm.com/storage/support/storwize/v7000/unified

Procedure

To use the initialization tool, complete the following steps.

1. If you downloaded the initialization tool, copy the file onto the USB flash drive that you are going to use.
2. To start the initialization tool, insert the USB flash drive that contains the program into a USB slot on a suitable personal computer.
3. Run the `InitTool.exe` program from the USB drive.
 - **Windows:** Open the USB flash drive and double-click `InitTool.bat`.

The initialization tool prompts you for the task that you want to perform and for the parameters that are relevant to that task. It prompts you when to put it in the node canister on the control enclosure or the file module.

4. After the `satask.txt` file is created, follow the instructions in “Using a USB flash drive” on page 207 to run the commands on the node.
5. When the commands have run, return the USB flash drive to your personal computer and start the tool again to see the results.

USB memory key has incorrect gateway address information

If the link on the `InitTool` panel to the management GUI does not work, the USB key may have an incorrect gateway address.

About this task

The `InitTool.exe` may indicate that the initial setup was successful, however, the link on the `InitTool` panel to the management GUI may not work. Given this scenario, it is possible that you have entered a management gateway IP address that is in the same subnet as the management IP address but is not the IP address of the gateway for this subnet. To check this, look inside the `satask.exe` file on the USB flash drive and note the IP address after the `-gw` switch. Make sure this IP address is the gateway for this subnet. If an IP address is needed then check this with your 1 Gbps Ethernet administrator.

If you did enter the wrong IP address for the gateway of this subnet and you have the correct gateway IP address ready, then it is possible to re-configure the control enclosure and file module to use the correct management gateway IP address.

If you have access to a computer that is plugged into the same Ethernet switch as the 1 Gbps Ethernet port 3 of each file module and the 1 Gbps Ethernet port 1 of each node canister in the control enclosure, then you may be able to ssh from it to the management IP address and log on as admin.

In this example, the default password is **admin001**:

```
ssh admin@<management IP address>
```

Use the **lssystemip** CLI command to show you the current management IP address setting on the control enclosure:

```
[kd52v6h.ibm]$ lssystemip
cluster_id cluster_name location port_id IP_address subnet_mask gateway
IP_address_6 prefix_6 gateway_6
00000200A9E0089E ifsc1uster-svt2 local 1 9.71.16.208 255.255.255.0 9.71.16.2
00000200A9E0089E ifsc1uster-svt2 local 2
```

If this command fails because the file module could not ssh to the control enclosure then refer to **Troubleshooting > Getting started troubleshooting > Installation troubleshooting > Problems with initial configuration** from the Problem Determination Guide.

Use the **chsystemip** CLI command to change the managed gateway IP address setting on the control enclosure. (This must be done first before you change the management gateway IP address setting on the file modules):

```
[kd52v6h.ibm]$ chsystemip -gw 9.71.16.1 -port 1
```

The active management node on the file module is not able to ssh CLI commands to the control enclosure until you change the management gateway setting to match the setting on the control enclosure. Use the **lsnwmgt** CLI command to show you the current management IP address setting on the file modules.

```
[kd52v6h.ibm]$ lsnwmgt
Interface Service IP Node1 Service IP Node2 Management IP Network Gateway LAN ID
ethX0 9.71.16.204 9.71.16.205 9.71.16.216 255.255.255.0 9.71.16.2
EFSSG1000I The command completed successfully
```

Use the **chnwmgt** CLI command to change the managed gateway IP address setting on the file modules.

```
[kd52v6h.ibm]$ chnwmgt --gateway 9.71.16.1
EFSSG0015I Refreshing data.
EFSSG1000I The command completed successfully
```

The active management node on the file module should now be able to ssh CLI commands to the control enclosure again. You should be able to access the management GUI or CLI from a computer, which is on a different subnet or different Ethernet switch to the Storwize V7000 Unified system. The link to the management GUI from the InitTool.exe panel should now work.

satask.txt commands

If you are creating the **satask.txt** command file by using a text editor, the file must contain a single command on a single line in the file.

The commands that you use are the same as the service CLI commands except where noted. Not all service CLI commands can be run from the USB flash drive. The **satask.txt** commands always run on the node that the USB flash drive is plugged into.

Reset service IP address and superuser password command:

Use this command to obtain service assistant access to a node canister even if the current state of the node canister is unknown. The physical access to the node canister is required and is used to authenticate the action.

Syntax

```

>> satask — chserviceip — --serviceip—ipv4— [—gw—ipv4—] [—mask—ipv4—] [—resetpassword—]
>> satask — chserviceip — --serviceip_6—ipv6— [—gw_6—ipv6—] [—prefix_6—int—]
>> satask — chserviceip — --default— [—resetpassword—]

```

Parameters

- serviceip *ipv4***
(Optional) The IPv4 address for the service assistant.
- gw *ipv4***
(Optional) The IPv4 gateway for the service assistant.
- mask *ipv4***
(Optional) The IPv4 subnet for the service assistant.
- serviceip_6 *ipv6***
(Optional) The IPv6 address for the service assistant.
- gw_6 *ipv6***
(Optional) The IPv6 gateway for the service assistant.
- default**
(Optional) Resets to the default IPv4 address.
- prefix_6 *int***
(Optional) The IPv6 prefix for the service assistant.
- resetpassword**
(Optional) Sets the service assistant password to the default value.

Description

This command resets the service assistant IP address to the default value. If the command is run on the upper canister, the default value is 192.168.70.121 subnet mask: 255.255.255.0. If the command is run on the lower canister, the default value is 192.168.70.122 subnet mask: 255.255.255.0. If the node canister is active in a system, the superuser password for the system is reset; otherwise, the superuser password is reset on the node canister.

If the node canister becomes active in a system, the superuser password is reset to that of the system. You can configure the system to disable resetting the superuser password. If you disable that function, this action fails.

This action calls the **satask chserviceip** command and the **satask resetpassword** command.

Reset service assistant password command:

Use this command when you are unable to logon to the system because you have forgotten the superuser password, and you wish to reset it.

Attention: Run this command only when instructed by IBM support. Running this command directly on a Storwize V7000 can affect your I/O operations on the file modules.

Syntax

▶▶ satask — resetpassword —————▶▶

Parameters

None.

Description

This command resets the service assistant password to the default value `passwd0rd`. If the node canister is active in a system, the superuser password for the system is reset; otherwise, the superuser password is reset on the node canister.

If the node canister becomes active in a system, the superuser password is reset to that of the system. You can configure the system to disable resetting the superuser password. If you disable that function, this action fails.

This command calls the **satask resetpassword** command.

Snap command:

Use the **snap** command to collect diagnostic information from the node canister and to write the output to a USB flash drive.

Attention: Run this command only when instructed by IBM support. Running this command directly on a Storwize V7000 can affect your I/O operations on the file modules.

Syntax

▶▶ satask — snap — [-dump] [-noimm] [*panel_name*]————▶▶

Parameters

-dump

(Optional) Indicates the most recent dump file in the output.

-noimm

(Optional) Indicates the /dumps/imm.ffdc file should not be included in the output.

panel_name

(Optional) Indicates the node on which to execute the **snap** command.

Description

This command moves a snap file to a USB flash drive.

This command calls the **satask snap** command.

If collected, the IMM FFDC file is present in the **snap** archive in /dumps/imm.ffdc.<node.dumptime>.<date>.<time>.tgz. The system waits for up to five minutes for the IMM to generate its FFDC. The status of the IMM FFDC is located in the **snap** archive in /dumps/imm.ffdc.log. These two files are not left on the node.

An invocation example

```
satask snap -dump 111584
```

The resulting output:

No feedback

Install software command:

Use this command to install a specific update package on the node canister.

Attention: Run this command only when instructed by IBM support. Running this command directly on a Storwize V7000 can affect your I/O operations on the file modules.

Syntax

```

▶▶ satask — installsoftware — — -file —filename— —————▶
                                     |
                                     | — -ignore —————▶
                                     | — -pacedccu —————▶

```

Parameters**-file filename**

(Required) The *filename* designates the name of the update package .

-ignore | -pacedccu

(Optional) Overrides prerequisite checking and forces installation of the update package.

Description

This command copies the file from the USB flash drive to the update directory on the node canister and then installs the update package.

This command calls the **satask installsoftware** command.

Create system command:

Use this command to create a storage system.

Note: The reference to cluster is not the same as the file system cluster on the Storwize V7000 file modules.

Attention: Run this command only when instructed by IBM support. Running this command directly on a Storwize V7000 can affect your I/O operations on the file modules.

Syntax

```
►► satask mkcluster -- -clusterip ipv4 [ -gw ipv4 ] [ -mask ipv4 ] [ -name cluster_name ]
►► satask mkcluster -- -clusterip_6 ipv6 [ -gw_6 ipv6 ] [ -prefix_6 int ] [ -name cluster_name ]
```

Parameters

-clusterip *ipv4*

(Optional) The IPv4 address for Ethernet port 1 on the system.

-gw *ipv4*

(Optional) The IPv4 gateway for Ethernet port 1 on the system.

-mask *ipv4*

(Optional) The IPv4 subnet for Ethernet port 1 on the system.

-clusterip_6 *ipv6*

(Optional) The IPv6 address for Ethernet port 1 on the system.

-gw_6 *ipv6*

(Optional) The IPv6 gateway for Ethernet port 1 on the system.

-prefix_6 *int*

(Optional) The IPv6 prefix for Ethernet port 1 on the system.

-name *cluster_name*

(Optional) The name of the new system.

Description

This command creates a storage system.

This command calls the **satask mkcluster** command.

Change system IP address:

Use this command to change the system IP address of the storage system.

It is best to use the initialization tool to create this command in satask.txt together with the associated cltask.txt file that changes the file modules management IP addresses.

Syntax

➡ satask — setssystemip — — -systemip —ipv4 — — -gw —ipv4 — — -mask —ipv4 — — -consoleip — ipv4➡

Parameters

-systemip

The IPv4 address for Ethernet port 1 on the system.

-gw

The IPv4 gateway for Ethernet port 1 on the system.

-mask

The IPv4 subnet for Ethernet port 1 on the system.

-consoleip

The management IPv4 address of Storwize V7000 Unified system.

Description

This command is only supported in the satask.txt file on a USB flash drive.

It calls the svctask chssystemip command if the USB flash drive is inserted in the configuration node canister, Otherwise it will blink the amber identify LED of the node canister that is the configuration node.

If the amber identify LED for a different node canister starts to blink then move the USB flash drive over to that node canister because it is the configuration node.

When the amber LED turns off you can move the USB flash drive to one of the file modules so that it will use the clitask.txt file to change the file module management IP addresses.

Leave the USB flash drive in the file module for at least two minutes before you remove it. Use a workstation to check the clitask_results.txt and satask.txt results files on the USB flash drive.

If the IP address change was successful then you must run the startmgtsrv -r command to restart the management service so that it will not continue to ssh commands to the old system IP address of the volume storage system.

For example, on a Linux workstation with network access to the new management IP address:

```
satask setssystemip -systemip 123.123.123.20 -gw 123.123.123.1 -mask 255.255.255.0  
-consoleip 123.123.123.10
```

You can now access the management GUI, which you can use to change any other IP address that needs to be changed.

Here is an example of what could be in the clitask.txt file:

```
chnwmgmt --serviceip1 123.123.123.11 --serviceip2 123.123.123.12  
--mgtip 123.123.123.10 --gateway 123.123.123.1 --netmask 255.255.255.0 --force  
chstoragesystem --ip1 123.123.123.20
```

Here is an example of what could be in the satask.txt file:

```
satask setssystemip -systemip 123.123.123.20 -gw 123.123.123.1 -mask 255.255.255.0  
-consoleip 123.123.123.10
```

Query status command:

Use this command to determine the current service state of the node canister.

Syntax

►— `sainfo — getstatus —` —————►◄

Parameters

None.

Description

This command writes the output from each node canister to the USB flash drive.

This command calls the **sainfo lsservicenodes** command, the **sainfo lsservicestatus** command, and the **sainfo lsservicerecommendation** command.

Starting statistics collection

You can start the collection of cluster statistics from the Starting the Collection of Statistics panel in the management GUI.

Introduction

For each collection interval, the management GUI creates four statistics files: one for managed disks (MDisks), named **Nm_stat**; one for volumes and volume copies, named **Nv_stat**; one for nodes, named **Nn_stat**; and one for drives, named **Nd_stat**. The files are written to the `/dumps/iostats` directory on the node. To retrieve the statistics files from the non-configuration nodes onto the configuration node, **svctask cpdumps** command must be used.

A maximum of 16 files of each type can be created for the node. When the 17th file is created, the oldest file for the node is overwritten.

Fields

The following fields are available for user definition:

Interval

Specify the interval in minutes between the collection of statistics. You can specify 1 - 60 minutes in increments of 1 minute.

Tables

The following tables describe the information that is reported for individual nodes and volumes.

Table 57 describes the statistics collection for MDisks, for individual nodes.

Table 57. Statistics collection for individual nodes

Statistic name	Description
id	Indicates the name of the MDisk for which the statistics apply.

Table 57. Statistics collection for individual nodes (continued)

idx	Indicates the identifier of the MDisk for which the statistics apply.
rb	Indicates the cumulative number of blocks of data that is read (since the node has been running).
re	Indicates the cumulative read external response time in milliseconds for each MDisk. The cumulative response time for disk reads is calculated by starting a timer when a SCSI read command is issued and stopped when the command completes successfully. The elapsed time is added to the cumulative counter.
ro	Indicates the cumulative number of MDisk read operations that are processed (since the node has been running).
rq	Indicates the cumulative read queued response time in milliseconds for each MDisk. This response is measured from above the queue of commands to be sent to an MDisk because the queue depth is already full. This calculation includes the elapsed time that is taken for read commands to complete from the time they join the queue.
wb	Indicates the cumulative number of blocks of data written (since the node has been running).
we	Indicates the cumulative write external response time in milliseconds for each MDisk. The cumulative response time for disk writes is calculated by starting a timer when a SCSI write command is issued and stopped when the command completes successfully. The elapsed time is added to the cumulative counter.
wo	Indicates the cumulative number of MDisk write operations processed (since the node has been running).
wq	Indicates the cumulative write queued response time in milliseconds for each MDisk. This is measured from above the queue of commands to be sent to an MDisk because the queue depth is already full. This calculation includes the elapsed time taken for write commands to complete from the time they join the queue.

Table 58 describes the VDisk (volume) information that is reported for individual nodes.

Note: MDisk statistics files for nodes are written to the /dumps/iostats directory on the individual node.

Table 58. Statistic collection for volumes for individual nodes

Statistic name	Description
id	Indicates the volume name for which the statistics apply.
idx	Indicates the volume for which the statistics apply.
rb	Indicates the cumulative number of blocks of data read (since the node has been running).
rl	Indicates the cumulative read response time in milliseconds for each volume. The cumulative response time for volume reads is calculated by starting a timer when a SCSI read command is received and stopped when the command completes successfully. The elapsed time is added to the cumulative counter.
rlw	Indicates the worst read response time in microseconds for each volume since the last time statistics were collected. This value is reset to zero after each statistics collection sample.

Table 58. Statistic collection for volumes for individual nodes (continued)

ro	Indicates the cumulative number of volume read operations processed (since the node has been running).
wb	Indicates the cumulative number of blocks of data written (since the node has been running).
wl	Indicates the cumulative write response time in milliseconds for each volume. The cumulative response time for volume writes is calculated by starting a timer when a SCSI write command is received and stopped when the command completes successfully. The elapsed time is added to the cumulative counter.
wlw	Indicates the worst write response time in microseconds for each volume since the last time statistics were collected. This value is reset to zero after each statistics collection sample.
wo	Indicates the cumulative number of volume write operations processed (since the node has been running).
wou	Indicates the cumulative number of volume write operations that are not aligned on a 4K boundary.
xl	Indicates the cumulative read and write data transfer response time in milliseconds for each volume since the last time the node was reset. When this statistic is viewed for multiple volumes and with other statistics, it can indicate if the latency is caused by the host, fabric, or the Storwize V7000 Unified.

Table 59 describes the VDisk information related to Metro Mirror or Global Mirror relationships that is reported for individual nodes.

Table 59. Statistic collection for volumes that are used in Metro Mirror and Global Mirror relationships for individual nodes

Statistic name	Description
gwl	Indicates cumulative secondary write latency in milliseconds. This statistic accumulates the cumulative secondary write latency for each volume. You can calculate the amount of time to recovery from a failure based on this statistic and the gws statistics.
gwo	Indicates the total number of overlapping volume writes. An overlapping write is when the logical block address (LBA) range of write request collides with another outstanding request to the same LBA range and the write request is still outstanding to the secondary site.
gwot	Indicates the total number of fixed or unfixed overlapping writes. When all nodes in all clusters are running Storwize V7000 Unified version 4.3.1, this records the total number of write I/O requests received by the Global Mirror feature on the primary that have overlapped. When any nodes in either cluster are running Storwize V7000 Unified versions earlier than 4.3.1, this value does not increment.
gws	Indicates the total number of write requests that have been issued to the secondary site.

Table 60 describes the port information that is reported for individual nodes

Table 60. Statistic collection for node ports

Statistic name	Description
----------------	-------------

Table 60. Statistic collection for node ports (continued)

bbcz	Indicates the total time in microseconds for which the port had data to send but was prevented from doing so by a lack of buffer credit from the switch.
cbr	Indicates the bytes received from controllers.
cbt	Indicates the bytes transmitted to disk controllers.
cer	Indicates the commands received from disk controllers.
cet	Indicates the commands initiated to disk controllers.
hbr	Indicates the bytes received from hosts.
hbt	Indicates the bytes transmitted to hosts.
her	Indicates the commands received from hosts.
het	Indicates the commands initiated to hosts.
icrc	Indicates the number of CRC that are not valid.
id	Indicates the port identifier for the node.
itw	Indicates the number of transmission word counts that are not valid.
lf	Indicates a link failure count.
lnbr	Indicates the bytes received to other nodes in the same cluster.
lnbt	Indicates the bytes transmitted to other nodes in the same cluster.
lner	Indicates the commands received from other nodes in the same cluster.
lnet	Indicates the commands initiated to other nodes in the same cluster.
lsi	Indicates the lost-of-signal count.
lsy	Indicates the loss-of-synchronization count.
pspe	Indicates the primitive sequence-protocol error count.
rmbr	Indicates the bytes received to other nodes in the other clusters.
rmbt	Indicates the bytes transmitted to other nodes in the other clusters.
rmer	Indicates the commands received from other nodes in the other clusters.
rmet	Indicates the commands initiated to other nodes in the other clusters.
wwpn	Indicates the worldwide port name for the node.

Table 61 describes the node information that is reported for each nodes.

Table 61. Statistic collection for nodes

Statistic name	Description
cluster_id	Indicates the name of the cluster.
cluster	Indicates the name of the cluster.

Table 61. Statistic collection for nodes (continued)

cpu	busy - Indicates the total CPU average core busy milliseconds since the node was reset. This statistic reports the amount of the time the processor has spent polling while waiting for work versus actually doing work. This statistic accumulates from zero.
	comp - Indicates the total CPU average core busy milliseconds for compression process cores since the node was reset.
	system - Indicates the total CPU average core busy milliseconds since the node was reset. This statistic reports the amount of the time the processor has spent polling while waiting for work versus actually doing work. This statistic accumulates from zero. This is the same information as the information provided with the cpu busy statistic and will eventually replace the cpu busy statistic.
cpu_core	id - Indicates the CPU core id.
	comp - Indicates the per-core CPU average core busy milliseconds for compression process cores since node was reset.
	system - Indicates the per-core CPU average core busy milliseconds for system process cores since node was reset.
id	Indicates the name of the node.
node_id	Indicates the unique identifier for the node.
rb	Indicates the number of bytes received.
re	Indicates the accumulated receive latency, excluding inbound queue time. This statistic is the latency that is experienced by the node communication layer from the time that an I/O is queued to cache until the time that the cache gives completion for it.
ro	Indicates the number of messages or bulk data received.
rq	Indicates the accumulated receive latency, including inbound queue time. This statistic is the latency from the time that a command arrives at the node communication layer to the time that the cache completes the command.
wb	Indicates the bytes sent.
we	Indicates the accumulated send latency, excluding outbound queue time. This statistic is the time from when the node communication layer issues a message out onto the Fibre Channel until the node communication layer receives notification that the message has arrived.
wo	Indicates the number of messages or bulk data sent.
wq	Indicates the accumulated send latency, including outbound queue time. This statistic includes the entire time that data is sent. This time includes the time from when the node communication layer receives a message and waits for resources, the time to send the message to the remote node, and the time taken for the remote node to respond.

Table 62 describes the statistics collection for volumes.

Table 62. Cache statistics collection for volumes and volume copies

Statistic	Acronym	Statistics for volume cache	Statistics for volume copy cache	Statistics for volume cache partition	Statistics for volume copy cache partition	Statistics for the Node Overall Cache	Cache statistics for mdisks	Units and state
read ios	ri	Yes	Yes					ios, cumulative
write ios	wi	Yes	Yes					ios, cumulative

Table 62. Cache statistics collection for volumes and volume copies (continued)

Statistic	Acronym	Statistics for volume cache	Statistics for volume copy cache	Statistics for volume cache partition	Statistics for volume copy cache partition	Statistics for the Node Overall Cache	Cache statistics for mdisks	Units and state
read misses	r	Yes	Yes					sectors, cumulative
read hits	rh	Yes	Yes					sectors, cumulative
flush_through writes	ft	Yes	Yes					sectors, cumulative
fast_write writes	fw	Yes	Yes					sectors, cumulative
write_through writes	wt	Yes	Yes					sectors, cumulative
write hits	wh	Yes	Yes					sectors, cumulative
prefetches	p		Yes					sectors, cumulative
prefetch hits (prefetch data that is read)	ph		Yes					sectors, cumulative
prefetch misses (prefetch pages that are discarded without any sectors read)	pm		Yes					pages, cumulative
modified data	m	Yes	Yes					sectors, snapshot, non-cumulative
read and write cache data	v	Yes	Yes					sectors snapshot, non-cumulative
destages	d	Yes	Yes					sectors, cumulative
fullness Average	fav			Yes	Yes			%, non-cumulative
fullness Max	fmx			Yes	Yes			%, non-cumulative
fullness Min	fmn			Yes	Yes			%, non-cumulative
Destage Target Average	dtav				Yes		Yes	IOs capped 9999, non-cumulative
Destage Target Max	dtmx				Yes			IOs, non-cumulative
Destage Target Min	dtmn				Yes			IOs, non-cumulative

Table 62. Cache statistics collection for volumes and volume copies (continued)

Statistic	Acronym	Statistics for volume cache	Statistics for volume copy cache	Statistics for volume cache partition	Statistics for volume copy cache partition	Statistics for the Node Overall Cache	Cache statistics for mdisks	Units and state
Destage In Flight Average	dfav				Yes		Yes	IOs capped 9999, non-cumulative
Destage In Flight Max	dfmx				Yes			IOs, non-cumulative
Destage In Flight Min	dfmn				Yes			IOs, non-cumulative
destage latency average	dav	Yes	Yes	Yes	Yes	Yes	Yes	µs capped 9999999, non-cumulative
destage latency max	dmx			Yes	Yes	Yes		µs capped 9999999, non-cumulative
destage latency min	dmn			Yes	Yes	Yes		µs capped 9999999, non-cumulative
destage count	dcn	Yes	Yes	Yes	Yes	Yes		ios, non-cumulative
stage latency average	sav	Yes	Yes			Yes		µs capped 9999999, non-cumulative
stage latency max	smx					Yes		µs capped 9999999, non-cumulative
stage latency min	smn					Yes		µs capped 9999999, non-cumulative
stage count	scn	Yes	Yes			Yes		ios, non-cumulative
prestage latency average	pav		Yes			Yes		µs capped 9999999, non-cumulative
prestage latency max	pmx					Yes		µs capped 9999999, non-cumulative
prestage latency min	pmn					Yes		µs capped 9999999, non-cumulative
prestage count	pcn		Yes			Yes		ios, non-cumulative

Table 62. Cache statistics collection for volumes and volume copies (continued)

Statistic	Acronym	Statistics for volume cache	Statistics for volume copy cache	Statistics for volume cache partition	Statistics for volume copy cache partition	Statistics for the Node Overall Cache	Cache statistics for mdisks	Units and state
Write Cache Fullness Average	wfav					Yes		%, non-cumulative
Write Cache Fullness Max	wfmx					Yes		%, non-cumulative
Write Cache Fullness Min	wfmn					Yes		%, non-cumulative
Read Cache Fullness Average	rfav					Yes		%, non-cumulative
Read Cache Fullness Max	rfmx					Yes		%, non-cumulative
Read Cache Fullness Min	rfmn					Yes		%, non-cumulative
Pinned Percent	pp	Yes	Yes	Yes	Yes	Yes		% of total cache snapshot, non-cumulative
data transfer latency average	tav	Yes	Yes					µs capped 9999999, non-cumulative
Track Lock Latency (Exclusive) Average	teav	Yes	Yes					µs capped 9999999, non-cumulative
Track Lock Latency (Shared) Average	tsav	Yes	Yes					µs capped 9999999, non-cumulative
Cache I/O Control Block Queue Time	hpt					Yes		Average µs, non-cumulative
Cache Track Control Block Queue Time	ppt					Yes		Average µs, non-cumulative
Owner Remote Credit Queue Time	opt					Yes		Average µs, non-cumulative
Non-Owner Remote Credit Queue Time	npt					Yes		Average µs, non-cumulative
Admin Remote Credit Queue Time	apt					Yes		Average µs, non-cumulative
Cddb Queue Time	cpt					Yes		Average µs, non-cumulative

Table 62. Cache statistics collection for volumes and volume copies (continued)

Statistic	Acronym	Statistics for volume cache	Statistics for volume copy cache	Statistics for volume cache partition	Statistics for volume copy cache partition	Statistics for the Node Overall Cache	Cache statistics for mdisks	Units and state
Buffer Queue Time	bpt					Yes		Average μ s, non-cumulative
Hardening Rights Queue Time	hrpt					Yes		Average μ s, non-cumulative

Note: Any statistic with a name **av**, **mx**, **mn**, and **cn** is not cumulative. These statistics reset every statistics interval. For example, if the statistic does not have a name with name **av**, **mx**, **mn**, and **cn**, and it is an Ios or count, it will be a field containing a total number.

- The term *pages* means in units of 4096 bytes per page.
- The term *sectors* means in units of 512 bytes per sector.
- The term μ s means microseconds.
- Non-cumulative means totals since the previous statistics collection interval.
- Snapshot means the value at the end of the statistics interval (rather than an average across the interval or a peak within the interval).

Table 63 describes the statistic collection for volume cache per individual nodes.

Table 63. Statistic collection for volume cache per individual nodes. This table describes the volume cache information that is reported for individual nodes.

Statistic name	Description
cm	Indicates the number of sectors of modified or dirty data that are held in the cache.
ctd	Indicates the total number of cache destages that were initiated writes, submitted to other components as a result of a volume cache flush or destage operation.
ctds	Indicates the total number of sectors that are written for cache-initiated track writes.
ctp	Indicates the number of track stages that are initiated by the cache that are prestage reads.
ctps	Indicates the total number of staged sectors that are initiated by the cache.
ctrh	Indicates the number of total track read-cache hits on prestage or non-prestage data. For example, a single read that spans two tracks where only one of the tracks obtained a total cache hit, is counted as one track read-cache hit.
ctrhps	Indicates the number of track reads received from other components, treated as cache hits on any prestaged data. For example, if a single read spans two tracks where only one of the tracks obtained a total cache hit on prestaged data, it is counted as one track read for the prestaged data. A cache hit that obtains a partial hit on prestage and non-prestage data still contributes to this value.
ctrhps	Indicates the total number of sectors that are read for reads received from other components that obtained cache hits on any prestaged data.

Table 63. *Statistic collection for volume cache per individual nodes (continued).* This table describes the volume cache information that is reported for individual nodes.

ctrhs	Indicates the total number of sectors that are read for reads received from other components that obtained total cache hits on prestige or non-prestige data.
ctr	Indicates the total number of track reads received. For example, if a single read spans two tracks, it is counted as two total track reads.
ctrs	Indicates the total number of sectors that are read for reads received.
ctwft	Indicates the number of track writes received from other components and processed in flush through write mode.
ctwfts	Indicates the total number of sectors that are written for writes that are received from other components and processed in flush through write mode.
ctwfw	Indicates the number of track writes received from other components and processed in fast-write mode.
ctwfwsh	Indicates the track writes in fast-write mode that were written in write-through mode because of the lack of memory.
ctwfwshs	Indicates the track writes in fast-write mode that were written in write through due to the lack of memory.
ctwfws	Indicates the total number of sectors that are written for writes that are received from other components and processed in fast-write mode.
ctwh	Indicates the number of track writes received from other components where every sector in the track obtained a write hit on already dirty data in the cache. For a write to count as a total cache hit, the entire track write data must already be marked in the write cache as dirty.
ctwhs	Indicates the total number of sectors that are received from other components where every sector in the track obtained a write hit on already dirty data in the cache.
ctw	Indicates the total number of track writes received. For example, if a single write spans two tracks, it is counted as two total track writes.
ctws	Indicates the total number of sectors that are written for writes that are received from components.
ctwwt	Indicates the number of track writes received from other components and processed in write through write mode.
ctwwts	Indicates the total number of sectors that are written for writes that are received from other components and processed in write through write mode.
cv	Indicates the number of sectors of read and write cache data that is held in the cache.

Table 64 describes the XML statistics specific to an IP Partnership port.

Table 64. *XML statistics for an IP Partnership port*

Statistic name	Description
ipbz	Indicates the average size (in bytes) of data that are being submitted to the IP partnership driver since the last statistics collection period.
ipre	Indicates the bytes retransmitted to other nodes in other clusters by the IP partnership driver.
iprt	Indicates the average round-trip time in microseconds for the IP partnership link since the last statistics collection period.

Table 64. XML statistics for an IP Partnership port (continued)

Statistic name	Description
iprx	Indicates the bytes received from other nodes in other clusters by the IP partnership driver.
ipsz	Indicates the average size (in bytes) of data that are being transmitted by the IP partnership driver since the last statistics collection period.
iptx	Indicates the bytes transmitted to other nodes in other clusters by the IP partnership driver.

Actions

The following actions are available to the user:

OK Click this button to change statistic collection.

Cancel

Click this button to exit the panel without changing statistic collection.

XML formatting information

The XML is more complicated now, as seen in this raw XML from the volume (Nv_statistics) statistics. Notice how the names are similar but because they are in a different section of the XML, they refer to a different part of the VDisk.

```
<vdisk idx="0"
ctrs="213694394" ctps="0" ctrhs="2416029" ctrhps="0"
ctds="152474234" ctwfts="9635" ctwwts="0" ctwfwts="152468611"
ctwhs="9117" ctws="152478246" ctr="1628296" ctw="3241448"
ctp="0" ctrh="123056" ctrhp="0" ctd="1172772"
ctwft="200" ctwwt="0" ctwfw="3241248" ctwfwsh="0"
ctwfwshs="0" ctwh="538" cm="13768758912876544" cv="13874234719731712"
gwot="0" gwo="0" gws="0" gwl="0"

id="Master_iogrp0_1"
ro="0" wo="0" rb="0" wb="0"
rl="0" wl="0" rlw="0" wlw="0" xl="0">
Vdisk/Volume statistics
<ca r="0" rh="0" d="0" ft="0"
wt="0" fw="0" wh="0" ri="0"
wi="0" dav="0" dcn="0" pav="0" pcn="0" teav="0" tsav="0" tav="0"
pp="0"/>
<cpy idx="0">
volume copy statistics
<ca r="0" p="0" rh="0" ph="0"
d="0" ft="0" wt="0" fw="0"
wh="0" pm="0" ri="0" wi="0"
dav="0" dcn="0" sav="0" scn="0"
pav="0" pcn="0" teav="0" tsav="0"
tav="0" pp="0"/>
</cpy>
</vdisk>
```

The <cpy idx="0"> means its in the volume copy section of the VDisk, whereas the statistics shown under Vdisk/Volume statistics are outside of the cpy idx section and therefore refer to a VDisk/volume.

Similarly for the volume cache statistics for node and partitions:

```

<uca><ca dav="18726" dcn="1502531" dmx="749846" dmn="89"
sav="20868" scn="2833391" smx="980941" smn="3"
pav="0" pcn="0" pmx="0" pmn="0"
wfav="0" wfm="2" wfmn="0"
rfav="0" rfm="1" rfmn="0"
pp="0"
hpt="0" ppt="0" opt="0" npt="0"
apt="0" cpt="0" bpt="0" hrpt="0"
/><partition id="0"><ca dav="18726" dcn="1502531" dmx="749846" dmn="89"
fav="0" fmx="2" fmn="0"
dfav="0" dfm="0" dfmn="0"
dtav="0" dtm="0" dtmn="0"
pp="0"/></partition>

```

This output describes the volume cache node statistics where `<partition id="0">` the statistics are described for partition 0.

Replacing `<uca>` with `<lca>` means that the statistics are for volume copy cache partition 0.

Event reporting

Events that are detected are saved in an event log. As soon as an entry is made in this event log, the condition is analyzed. If any service activity is required, a notification is sent, if you have set up notifications.

Event reporting process

The following methods are used to notify you and the IBM Support Center of a new event:

- If you enabled Simple Network Management Protocol (SNMP), an SNMP trap is sent to an SNMP manager that is configured by the customer.
- If enabled, log messages can be forwarded on an IP network by using the syslog protocol.
- If enabled, event notifications can be forwarded by email by using Simple Mail Transfer Protocol (SMTP).
- Call Home can be enabled so that critical faults generate a problem management record (PMR) that is then sent directly to the appropriate IBM Support Center by using email.

Understanding events

When a significant change in status is detected, an event is logged in the event log.

Error data

Events are classified as either alerts or messages:

- An alert is logged when the event requires some action. Some alerts have an associated error code that defines the service action that is required. The service actions are automated through the fix procedures. If the alert does not have an error code, the alert represents an unexpected change in state. This situation must be investigated to see if it is expected or represents a failure. Investigate an alert and resolve it as soon as it is reported.
- A message is logged when a change that is expected is reported, for instance, an IBM FlashCopy[®] operation completes.

Viewing the event log

You can view the event log by using the management GUI or the command-line interface (CLI).

About this task

You can view the event log by using the **Monitoring > Events** options in the management GUI. The event log contains many entries. You can, however, select only the type of information that you need.

You can also view the event log by using the command-line interface (**lseventlog**). See the “Command-line interface” topic for the command details.

Managing the event log

The event log has a limited size. After it is full, newer entries replace entries that are no longer required.

To avoid having a repeated event that fills the event log, some records in the event log refer to multiple occurrences of the same event. When event log entries are coalesced in this way, the time stamp of the first occurrence and the last occurrence of the problem is saved in the log entry. A count of the number of times that the error condition has occurred is also saved in the log entry. Other data refers to the last occurrence of the event.

Describing the fields in the event log

The event log includes fields with information that you can use to diagnose problems.

Table 65 describes some of the fields that are available to assist you in diagnosing problems.

Table 65. Description of data fields for the event log

Data field	Description
Event ID	This number precisely identifies why the event was logged.
Description	A short description of the event.
Status	Indicates whether the event requires some attention. Alert: if a red icon with a cross is shown, follow the fix procedure or service action to resolve the event and turn the status green. Monitoring: the event is not yet of concern. Expired: the event no longer represents a concern. Message: provide useful information about system activity.
Error code	Indicates that the event represents an error in the system that can be fixed by following the fix procedure or service action that is identified by the error code. Not all events have an error code. Different events have the same error code if the same service action is required for each.
Sequence number	Identifies the event within the system.
Event count	The number of events that are coalesced into this event log record.
Object type	The object type to which the event relates.
Object ID	Uniquely identifies the object within the system to which the event relates.

Table 65. Description of data fields for the event log (continued)

Data field	Description
Object name	The name of the object in the system to which the event relates.
Copy ID	If the object is a volume and the event refers to a specific copy of the volume, this field is the number of the copy to which the event relates.
Reporting node ID	Typically identifies the node responsible for the object to which the event relates. For events that relate to nodes, it identifies the node that logged the event, which can be different from the node that is identified by the object ID.
Reporting node name	Typically identifies the node that contains the object to which the event relates. For events that relate to nodes, it identifies the node that logged the event, which can be different from the node that is identified by the object name.
Fixed	Where an alert is shown for an error or warning condition, it indicates that the user marked the event as fixed, completed the fix procedure, or that the condition was resolved automatically. For a message event, this field can be used to acknowledge the message.
First time stamp	The time when this error event was reported. If events of a similar type are being coalesced together, so that one event log record represents more than one event, this field is the time the first error event was logged.
Last time stamp	The time when the last instance of this error event was recorded into this event log record.
Root sequence number	If set, it is the sequence number of an event that represents an error that probably caused this event to be reported. Resolve the root event first.
Sense data	Additional data that gives the details of the condition that caused the event to be logged.

Event notifications

Storwize V7000 Unified can use Simple Network Management Protocol (SNMP) traps, syslog messages, emails and Call Homes to notify you and IBM(r) Remote Technical Support when significant events are detected. Any combination of these notification methods can be used simultaneously. Notifications are normally sent immediately after an event is raised. However, there are some events that might occur because of service actions that are being performed. If a recommended service action is active, these events are notified only if they are still unfixed when the service action completes.

Only events recorded in the event log can be notified. Most CLI messages in response to some CLI commands are not recorded in the event log so do not cause an event notification.

Table 66 on page 229 describes the levels of event notifications.

Table 66. Notification levels

Notification level	Description
Error	<p>Error notification is sent to indicate a problem that must be corrected as soon as possible.</p> <p>This notification indicates a serious problem with the system. For example, the event that is being reported could indicate a loss of redundancy in the system, and it is possible that another failure could result in loss of access to data. The most typical reason that this type of notification is sent is because of a hardware failure, but some configuration errors or fabric errors also are included in this notification level. Error notifications can be configured to be sent as a call home message to your support center.</p>
Warning	<p>A warning notification is sent to indicate a problem or unexpected condition with the system. Always immediately investigate this type of notification to determine the effect that it might have on your operation, and make any necessary corrections.</p> <p>A warning notification does not require any replacement parts and therefore should not require involvement from your support center. The allocation of notification type Warning does not imply that the event is less serious than one that has notification level Error.</p>
Information	<p>An informational notification is sent to indicate that an expected event has occurred: for example, a FlashCopy operation has completed. No remedial action is required when these notifications are sent.</p>

Power-on self-test

When you turn on the system, the file modules and the control enclosure node canisters complete self-tests.

A series of tests is completed to check the operation of components and some of the options that have been installed when the units are first turned on. This series of tests is called the power-on self-test (POST).

If a critical failure is detected during the POST, the software is not loaded and the fault LED is illuminated. To determine if there is a POST error on a file module or a node canister, go to “Procedure: Understanding the Storwize V7000 Gen1 system status using the LEDs” on page 262.

When the code is loaded, additional testing takes place, which ensures that all of the required hardware and code components are installed and functioning correctly.

Understanding event codes

Informational events provide information on the status of an operation. Information events are recorded in the error event log, and depending on the configuration, you can be notified through email, SNMP, and syslog.

Understanding the error codes

Error codes are generated by the event-log analysis and system configuration code.

Error codes help you to identify the cause of a problem, a failing component, and the service actions that might be needed to solve the problem.

Viewing logs and traces

The Storwize V7000 Unified clustered system maintains log files and trace files that can be used to manage your system and diagnose problems.

You can view information about collecting log files or you can view examples of a configuration dump, error log, or featurization log. To do this, click **Reference** in the left pane of the IBM online information, and then expand the **Logs and traces** section.

Battery operation for the control enclosure

Storwize V7000 Unified node canisters cache volume data and hold state information in volatile memory. Battery operations differ, depending on the generation of your control enclosure model.

Storwize V7000 Unified *Gen1* refers to the enclosure models in the following table:

Table 67. Storwize V7000 Unified Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 Unified control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 Unified control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 Unified control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 Unified control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 Unified expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Storwize V7000 Unified *Gen2* refers to the newer generation of enclosures in the following table:

Table 68. Storwize V7000 Unified Gen2 model numbers

Machine type/model	Description
2076-524	Storwize V7000 Unified control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-12F	Storwize V7000 Unified expansion enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Battery operation for Storwize V7000 Gen2 control enclosures

Each node canister in the control enclosure caches critical data and holds state information in volatile memory.

If power to a node canister fails, the node canister uses battery power to write cache and state data to its boot drive.

Note: Storwize V7000 Gen2 expansion canisters do not cache volume data or store state information in volatile memory. Therefore, expansion canisters do not require battery power. If ac power to both power supplies in an expansion enclosure fails, the enclosure powers off. When ac power is restored to at least one power supply, the enclosure restarts without operator intervention.

The battery is maintained in a fully charged state by the battery subsystem. At maximum power, the battery can save critical data and state information in two back-to-back power failures. If power to a node canister is lost, saving critical data starts after a five-second wait. (If the outage is shorter than five seconds, the battery continues to support the node and critical data is not saved.) The node canister stops handling I/O requests from host applications. The saving of critical data runs to completion, even if power is restored during this time. The loss of power might be because the input power to the enclosure is lost, or because the node canister is removed from the enclosure.

When power is restored to the node canister, the system restarts without operator intervention. How quickly it restarts depends on whether there is a history of previous power failures. The system restarts only when the battery has sufficient charge for the node canister to save the cache and state data again. A node canister with multiple power failures might not have sufficient charge to save critical data. In such a case, the system starts in service state and waits to start I/O operations until the battery has sufficient charge.

Two light-emitting diode (LED) indicators indicate the state of the battery:

- Status LED - Green
- Fault LED - Amber

See “Procedure: Understanding the system status from the Storwize V7000 Gen2 LEDs” for a complete description of the battery LEDs.

Important: Although Storwize V7000 Gen2 is resilient to power failures and brown outs, always install Storwize V7000 Gen2 in an environment with reliable, consistent, and required ac power. Consider uninterruptible power supply units to avoid extended interruptions to data access.

Design parameters

Consider the following important design parameters:

- The design life of the battery in the Storwize V7000 Gen2 is five years service after one year on the shelf.
- Each battery is automatically reconditioned every three months to measure the battery capacity. Batteries in the same enclosure are not reconditioned within two days of each other. If a battery has a lower capacity than required (below the planned threshold), it is marked as “End Of Life” and should be replaced.
- Each battery provides power only for the canister in which it is installed. If a battery fails, the canister goes offline and reports a node error. The single running canister destages its cache and runs the I/O group in “write-through” mode until its partner canister is repaired and online.

Reconditioning the Storwize V7000 Gen2 battery

Reconditioning the battery ensures that the system can accurately determine the charge in the battery.

As a battery ages, it loses capacity. When a battery no longer has capacity to protect against two power loss events, it reports the battery end of life event and it should be replaced.

A reconditioning cycle is automatically scheduled to occur approximately once every three months, but reconditioning is rescheduled or canceled if the system loses redundancy. In addition, a two-day delay is imposed between the recondition cycles of the two batteries in one enclosure.

Battery operation for Storwize V7000 Unified Gen1 control enclosures

Storwize V7000 Unified node canisters cache volume data and hold state information in volatile memory.

If the power fails, cache and state data is written to a local flash drive in the canister. The batteries within the control enclosure provide the power to write the cache and state data to a local drive.

Note: Storwize V7000 Unified expansion canisters do not cache volume data or store state information in volatile memory. Therefore, expansion canisters do not require battery power. If ac power to both power supplies in an expansion enclosure fails, the enclosure powers off. When ac power is restored to at least one power supply, the expansion enclosure restarts without operator intervention.

Two power supply units exist in the control enclosure. Each one contains an integrated battery. Both power supply units and batteries provide power to both control canisters. Each battery has sufficient charge to save critical data from both node canisters to the local drive. A fully redundant system has two batteries and two canisters. In such a system, the batteries can save critical data and state information from both canisters to a local drive twice. In a system with a failed battery, there is enough charge in the remaining battery to support saving critical data from both canisters to a local drive once.

Without ac power, a canister starts saving critical data to a local drive approximately 10 seconds after it detects the loss. If power is restored within the 10 seconds, the system continues to operate. This loss in power is called a *brown out*. When it saves critical data, the system stops handling I/O requests from host applications. At the same time, Metro Mirror and Global Mirror relationships go offline. The system powers off after it saves the critical data.

If both node canisters shut down without writing the cache and state data to the local drive, the system is unable to restart without an extended service action. The system configuration must be restored. If any cache write data is lost, volumes must be restored from a backup. It is, therefore, important not to remove the canisters or the power supply units from the control enclosures unless directed to do so by the service procedures. Removing either of these components might prevent the node canister from writing its cache and state data to the local drive.

When the ac power is restored to the control enclosure, the system restarts without operator intervention. How quickly it restarts depends on whether there is a history of previous power failures.

When ac power is restored after both canisters save critical data, the system restarts without operator intervention. How quickly it restarts depends on whether there is a history of previous power failures. The system restarts only when the batteries have sufficient charge to save critical data for both node canisters again. A fully redundant system can restart as soon as power is restored after it saves critical data once. If a second ac power outage occurs before batteries charge completely, the system starts in service state. The system does not begin I/O operations until the batteries are half charged. Recharging takes approximately 30 minutes.

In a system with a failed battery, an ac power failure causes both canisters to save critical data and completely discharges the remaining battery. When the ac power is restored, the system starts in service state and does not start I/O operations until the remaining battery is fully charged. The recharging takes approximately 1 hour.

A battery is considered failed for the following conditions:

- The system can communicate with the battery, but the battery reports an error.
- The system is unable to communicate with the battery. The failure exists because the power supply is absent or because the power supply failed without communicating with the system.

Other conditions can cause critical data to be saved and the nodes to go into service state and stop I/O operations. Each node canister saves critical data if it detects there is no longer sufficient battery charge to support saving critical data. This situation happens, for example, when both batteries have two-thirds of a charge. The total charge is sufficient to support saving critical data once. Therefore, both canisters are in active state and I/O operations occur. If one battery fails, the remaining battery has only two-thirds of a charge. The total charge in the enclosure is now insufficient to save critical data if the ac power fails. Data protection is not guaranteed in this case. The nodes use ac power to save critical data in such a situation and enter service state. The nodes do not handle I/O operations until the remaining battery has sufficient charge to support the saving of critical data. When the battery has sufficient charge, the system automatically restarts.

Important: Although Storwize V7000 Unified is resilient to power failures and brown outs, always install in an environment with reliable, consistent, and required AC power. Consider uninterruptible power supply units to avoid extended interruptions to data access.

Design parameters

Consider the following important design parameters:

- The design life of the battery in the Storwize V7000 Unified is five years service after one year on the shelf.
- No periodic *learning mode* or reconditioning cycle occurs in the battery of this product.
- Each battery provides power only for the canister in which it is installed. If a battery fails, the canister goes offline and reports a node error. The single running canister destages its cache and runs the I/O group in “write-through” mode until its partner is repaired and online.

Maintenance discharge cycles

Maintenance discharge cycles extend the lifetime of the batteries and ensure that the system can accurately measure the charge in the batteries. Discharge cycles guarantee that the batteries have sufficient charge to protect the Storwize V7000 Unified system.

Maintenance discharge cycles are scheduled automatically by the system and involve fully discharging a battery and then recharging it again. Maintenance discharges are normally scheduled only when the system has two fully charged batteries. This condition ensures that for the duration of the maintenance cycle, the system still has sufficient charge to complete a save of the critical data if the ac power fails. This condition also ensures that I/O operations continue while the maintenance cycle is completed. It is usual for both batteries to require a maintenance discharge at the same time. In these circumstances, the system automatically schedules the maintenance of one battery. When the maintenance on that battery completes, the maintenance on the other battery starts.

Maintenance discharges are scheduled for the following situations:

- A battery has been powered on for three months without a maintenance discharge.
- A battery has provided protection for saving critical data at least twice.
- A battery has provided protection for at least 10 brown outs, which lasted up to 10 seconds each.

A maintenance discharge takes approximately 10 hours to complete. If the ac power outage occurs during the maintenance cycle, the cycle must be restarted. The cycle is scheduled automatically when the battery is fully charged.

Under the following conditions, a battery is not considered when calculating whether there is sufficient charge to protect the system. This condition persists until a maintenance discharge cycle is completed.

- A battery is completing a maintenance discharge.
- A battery has provided protection for saving critical data at least four times without any intervening maintenance discharge.
- A battery has provided protection for at least 20 brown outs, which lasted up to 10 seconds each.
- A battery must restart a maintenance discharge because the previous maintenance cycle was disrupted by an ac power outage.

If a system suffers repeated ac power failures without a sufficient time interval in between the ac failures to complete battery conditioning, then neither battery is considered when calculating whether there is sufficient charge to protect the system. In these circumstances, the system enters service state and does not permit I/O operations to be restarted until the batteries have charged and one of the batteries has completed a maintenance discharge. This activity takes approximately 10 hours.

If one of the batteries in a system fails and is not replaced, it prevents the other battery from completing a maintenance discharge. Not only does this condition reduce the lifetime of the remaining battery, but it also prevents a maintenance discharge cycle from occurring after the battery has provided protection for at least 2 critical saves or 10 brown outs. Preventing this maintenance cycle from occurring increases the risk that the system accumulates a sufficient number of power outages to cause the remaining battery to be discounted when calculating whether

there is sufficient charge to protect the system. This condition results in the system entering service state while the one remaining battery completes a maintenance discharge. I/O operations are not permitted during this process. This activity takes approximately 10 hours.

Understanding the medium errors and bad blocks

A storage system returns a medium error response to a host when it is unable to successfully read a block. The Storwize V7000 Unified response to a host read follows this behavior.

The volume virtualization that is provided extends the time when a medium error is returned to a host. Because of this difference to non-virtualized systems, the Storwize V7000 Unified uses the term *bad blocks* rather than medium errors.

The Storwize V7000 Unified allocates volumes from the extents that are on the managed disks (MDisks). The MDisk can be a volume on an external storage controller or a RAID array that is created from internal drives. In either case, depending on the RAID level used, there is normally protection against a read error on a single drive. However, it is still possible to get a medium error on a read request if multiple drives have errors or if the drives are rebuilding or are offline due to other issues.

The Storwize V7000 Unified provides migration facilities to move a volume from one underlying set of physical storage to another or to replicate a volume that uses FlashCopy or Metro Mirror or Global Mirror. In all these cases, the migrated volume or the replicated volume returns a medium error to the host when the logical block address on the original volume is read. The system maintains tables of bad blocks to record where the logical block addresses that cannot be read are. These tables are associated with the MDisks that are providing storage for the volumes.

The **dumpmdiskbadblocks** command and the **dumpallmdiskbadblocks** command are available to query the location of bad blocks.

Important: The **dumpmdiskbadblocks** only outputs the virtual medium errors that have been created, and not a list of the actual medium errors on MDisks or drives.

It is possible that the tables that are used to record bad block locations can fill up. The table can fill either on an MDisk or on the system as a whole. If a table does fill up, the migration or replication that was creating the bad block fails because it was not possible to create an exact image of the source volume.

The system creates alerts in the event log for the following situations:

- When it detects medium errors and creates a bad block
- When the bad block tables fill up

Table 69 lists the bad block error codes.

Table 69. Bad block errors

Error code	Description
1840	The managed disk has bad blocks. On an external controller, this can only be a copied medium error.

Table 69. Bad block errors (continued)

Error code	Description
1226	The system has failed to create a bad block because the MDisk already has the maximum number of allowed bad blocks.
1225	The system has failed to create a bad block because the system already has the maximum number of allowed bad blocks.

The recommended actions for these alerts guide you in correcting the situation.

Clear bad blocks by deallocating the volume disk extent, by deleting the volume or by issuing write I/O to the block. It is good practice to correct bad blocks as soon as they are detected. This action prevents the bad block from being propagated when the volume is replicated or migrated. It is possible, however, for the bad block to be on part of the volume that is not used by the application. For example, it can be in part of a database that has not been initialized. These bad blocks are corrected when the application writes data to these areas. Before the correction happens, the bad block records continue to use up the available bad block space.

Resolving a problem

Described here are some procedures to help resolve fault conditions that might exist on your system and which assume a basic understanding of the Storwize V7000 Unified system concepts.

The following procedures are often used to find and resolve problems:

- Procedures that involve data collection and system configuration
- Procedures that are used for hardware replacement.

Always use the recommended actions on the Events panel of the management GUI as the starting point to diagnose and resolve a problem.

The following topics describe a type of problem that you might experience, that is not resolved by using the management GUI. In those situations, review the symptoms and follow the actions that are provided here.

The “Start here: Use the management GUI recommended actions” topic gives the starting point for any service action. The situations covered in this section are the cases where you cannot start the management GUI or the node canisters in the control enclosure are unable to run the system software.

Note: After you have created your clustered system, remove hardware components only when directed to do so by the fix procedures. Failure to follow the procedures can result in loss of access to data or loss of data. Follow the fix procedures when servicing a control enclosure.

Start here: Use the management GUI recommended actions

The management GUI provides extensive facilities to help you troubleshoot and correct problems on your system.

You can connect to and manage a Storwize V7000 Unified system as soon as you have completed the USB initialization.

When you have logged on, select **Monitoring > Events**. You can work with two separate event logs:

- To work with events for the file modules, select the **File** tab. No fix procedures are available to be run. From the Storwize V7000 Unified Information Center, look up the errors.
- To work with events for the storage system, select the **Block** tab.

For the Storwize V7000 storage system, depending on how you choose to filter alerts, you might see only the alerts that require attention, alerts and messages that are not fixed, or all event types whether they are fixed or unfixed.

Select the recommended alert, or any other alert, and run the fix procedure. The fix procedure steps you through the process of troubleshooting and correcting the problem. The fix procedure displays information that is relevant to the problem and provides various options to correct the problem. Where it is possible, the fix procedure runs the commands that are required to reconfigure the system.

Always use the recommended action for an alert because these actions ensure that all required steps are taken. Use the recommended actions even in cases where the service action seems obvious, such as a drive showing a fault. In this case, the drive must be replaced and reconfiguration must be performed. The fix procedure performs the reconfiguration for you.

The fix procedure also checks that another existing problem does not result in a fix procedure that causes volume data to be lost. For example, if a power supply unit in a node enclosure must be replaced, the fix procedure checks and warns you if the integrated battery in the other power supply unit is not sufficiently charged to protect the system.

If possible, fix the alerts in the order shown to resolve the most serious issues first. Often, other alerts are fixed automatically because they were the result of a more serious issue.

After all the alerts are fixed, go to “Procedure: Checking the status of your system” on page 254.

Problem: Management IP address unknown

If you are not able to run the management GUI because you do not know the management IP address, you can retrieve the management IP address.

This topic also helps if the configuration communication between the file system (file modules) and the control enclosure is not working because the wrong IP address is being used.

The management IP address is set when the USB initialization is completed. An address for port 2 can be added later.

Problem: Unable to connect to the management GUI

If you are unable to connect to the management GUI from your web browser and received a Page not found or similar error, this information might help you resolve the issue. The connection information differs, depending on the generation of your control enclosure model.

Storwize V7000 Unified *Gen1* refers to the enclosure models in the following table:

Table 70. Storwize V7000 Unified Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 Unified control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 Unified control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 Unified control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 Unified control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 Unified expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Storwize V7000 Unified Gen2 refers to the newer generation of enclosures in the following table:

Table 71. Storwize V7000 Unified Gen2 model numbers

Machine type/model	Description
2076-524	Storwize V7000 Unified control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-12F	Storwize V7000 Unified expansion enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Problem: Unable to connect to the Storwize V7000 Gen2 management GUI

If you are unable to connect to the management GUI from your web browser and received a Page not found or similar error, this information might help you resolve the issue.

If you are unable to connect to the management GUI from your web browser and received a Certificate expired or similar error, see Resolving a problem with SSL certificates.

If you are unable to connect to the management GUI from your web browser and received a cipher error, SSL error, TLS error, or handshake error or similar error, see .

Consider the following possibilities if you are unable to connect to the Storwize V7000 Gen2 management GUI:

- You cannot connect if the system is not operational with at least one node online. If you know the service address of a node canister, either use the service assistant to verify that the state of at least one node canister is active, or if the node canister is not active, use the LEDs to see if any node canister state is active.

If there is not a node canister with a state of active, resolve the reason why it is not in active state. If the state of all node canisters is candidate, then there is not a clustered system to connect to. If the node state is service, see the topic that contains information about fixing node errors.

- Ensure that you are using the correct system IP address. If you can access the service assistant using the service address or the technician port of a node canister, log in to find the node and system addresses on the **Access** tab of the Node Detail table.
- Ensure that all node canisters have an Ethernet cable that is connected to port 1 and that the port is working. Use the Ethernet LEDs to understand the port status.
- Ping the management address to see if the Ethernet network permits the connection. If the ping fails, check the Ethernet network configuration to see if there is a routing or a firewall issue. Ensure that the Ethernet network configuration is compatible with the gateway and subnet or prefix settings. Ensure that you did not use the Ethernet address of another device as the management address. If necessary, modify your network settings to establish a connection.
- If the system IP address settings are incorrect for your environment, take these steps:
 1. You can determine this if you can access the service assistant on any node canister. Access the service assistant using the technician port on the rear of a node canister if it cannot be accessed over your network. Alternatively use the summary data returned, when a USB flash drive is plugged into a node canister.
 2. You can temporarily run the management GUI on the service address of the configuration node. Point your browser to service address/gui. For example, if the service address of the configuration node is 11.22.33.44, point your browser to 11.22.33.44/gui.
 3. In the Management GUI, use the options in the **Settings > Network** to change the management IP settings.
 4. As an alternative to using the management GUI, you can use the **chsystemip** CLI command to correct the system IP address settings by using ssh to the service IP of the configuration node.

Problem: Unable to connect to the Storwize V7000 Gen1 management GUI

If you are unable to connect to the management GUI from your web browser and received a Page not found or similar error, this information might help you resolve the issue.

If you are unable to connect to the management GUI from your web browser and received a Certificate expired or similar error, see Resolving a problem with SSL certificates.

If you are unable to connect to the management GUI from your web browser and received a cipher error, SSL error, TLS error, or handshake error or similar error, see .

Consider the following possibilities if you are unable to connect to the management GUI:

- You cannot connect if the system is not operational with at least one node online. If you know the service address of a node canister, go to “Procedure: Getting node canister and system information using the service assistant” on page 254; otherwise, go to “Procedure: Getting node canister and system information using a USB flash drive” on page 255 and obtain the state of each of the node canisters from the data that is returned. If there is not a node canister with a state of active, resolve the reason why it is not in active state. If the state

of all node canisters is candidate, then there is not a clustered system to connect to. If all nodes are in a service state, go to “Procedure: Fixing node errors” on page 271.

- Ensure that you are using the correct system IP address. If you know the service address of a node canister, go to “Procedure: Getting node canister and system information using the service assistant” on page 254; otherwise, go to “Procedure: Getting node canister and system information using a USB flash drive” on page 255 and obtain the management IP address from the data that is returned.
- Ensure that all node canisters have an Ethernet cable that is connected to port 1 and that the port is working. To understand the port status, go to “Procedure: Finding the status of Storwize V7000 Gen1 Ethernet connections” on page 268.
- Ping the management address to see if the Ethernet network permits the connection. If the ping fails, check the Ethernet network configuration to see if there is a routing or a firewall issue. Ensure that the Ethernet network configuration is compatible with the gateway and subnet or prefix settings. Ensure that you have not used the Ethernet address of another device as the management address. If necessary, modify your network settings to establish a connection.
- If the system IP address settings are incorrect for your environment, take these steps:
 1. Determine the service address of the configuration node canister. You can determine this if you can access the service assistant on any node canister, alternatively use the summary data returned, when a USB flash drive is plugged into a node canister.
 2. You can temporarily run the management GUI on the service address of the configuration node. Point your browser to *service address/gui*. For example, if the service address of the configuration node is 11.22.33.44, point your browser to 11.22.33.44/gui.
 3. Use the options in the **Settings > Network** panel to change the management IP settings.
 4. As an alternative to using the management GUI, you can use the **chsystemip** CLI command to correct the system IP address settings by using ssh to the service IP of the configuration node.

Problem: Unable to log on to the management GUI

If you can see the management GUI login screen but cannot log on, you have several options for correcting the problem.

Log on using your user name and password. Complete the suggested actions when you encounter a specific situation.

- If you are not logging on as superuser, contact your system administrator to verify your user name and reset your account password.
- If the user name that you are using is authenticated through a remote authentication server, verify that the server is available. If the authentication server is unavailable, you can log on as user name superuser. This user is always authenticated locally.
- If you do not know the password for superuser, go to “Procedure: Resetting superuser password” on page 250.

Problem: Cannot initialize or create a clustered system

Use this information if your attempt to create a clustered system has failed. This information varies depending on the generation of your control enclosure model.

Storwize V7000 Unified *Gen1* refers to the enclosure models in the following table:

Table 72. Storwize V7000 Unified Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 Unified control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 Unified control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 Unified control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 Unified control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 Unified expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Storwize V7000 Unified *Gen2* refers to the newer generation of enclosures in the following table:

Table 73. Storwize V7000 Unified Gen2 model numbers

Machine type/model	Description
2076-524	Storwize V7000 Unified control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-12F	Storwize V7000 Unified expansion enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Problem: Cannot initialize or create a Storwize V7000 Gen2 clustered system

Use this information if your attempt to create a Storwize V7000 Gen2 clustered system failed.

The failure is reported regardless of the method that you used to create a system:

- System initialization wizard (using the technician port of a node canister)
- USB flash drive
- Management console
- Service assistant
- Service command line

To prevent accidental loss of volumes, a system cannot be initialized on an enclosure that is already configured in an existing system. Check the following details using the service assistant to confirm that the enclosure is not already configured in a system:

- The node canister that you are attempting to create a clustered system on must be in candidate state.

- The partner node canister in the control enclosure must not be in active state.
- The latest system ID of the control enclosure must be 0. If the partner node is in active state, see [Start here: Use the management GUI recommended actions](#). If the partner code is not in the active state and the node canister on which you are attempting to create the system is in service state, see [“Procedure: Fixing node errors”](#) on page 271.

Problem: Cannot initialize or create a Storwize V7000 Gen1 clustered system

If your attempt to create a system has failed, there are several possible solutions depending on the current state of the storage system.

Note: This clustered storage system is different from the file system cluster on the file modules.

The failure is reported regardless of the method that you used to create a clustered storage system:

- USB flash drive
- Management console
- Service assistant
- Service command line

The create clustered-system function protects the system from loss of volume data. If you create a clustered system on a control enclosure that was previously used, you lose all of the volumes that you previously had. To determine if there is an existing system, use data that is returned by [“Procedure: Getting node canister and system information using the service assistant”](#) on page 254 or [“Procedure: Getting node canister and system information using a USB flash drive”](#) on page 255.

- The node canister that you are attempting to create a clustered system on is in candidate state. The node canister is in candidate state if it is a new canister.
- The partner node canister in the control enclosure is not in active state.
- The latest system ID of the control enclosure is 0.

If the create function failed because there is an existing system, fix the existing clustered system; do not re-create a new clustered system. If you want to create a clustered system and do not want to use any data from the volumes used in the previous clustered system, go to [“Procedure: Deleting a system completely”](#) on page 270, and then run the create function again.

You might not be able to create a cluster if the node canister (the one on which you are attempting to create the clustered system) is in service state. Check whether the node canister is in service state by using the data returned by [“Procedure: Getting node canister and system information using the service assistant”](#) on page 254 or [“Procedure: Getting node canister and system information using a USB flash drive”](#) on page 255. If the node is in service state, fix the reported node errors. For more information, go to [“Procedure: Fixing node errors”](#) on page 271. After the node error is corrected, attempt to create a clustered storage system again.

Problem: Node canister service IP address unknown

You can use several methods to determine the service address of a node canister. The methods of determining the service address differ, depending on the generation of your control enclosure model.

Storwize V7000 Unified *Gen1* refers to the enclosure models in the following table:

Table 74. Storwize V7000 Unified Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 Unified control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 Unified control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 Unified control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 Unified control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 Unified expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Storwize V7000 Unified *Gen2* refers to the newer generation of enclosures in the following table:

Table 75. Storwize V7000 Unified Gen2 model numbers

Machine type/model	Description
2076-524	Storwize V7000 Unified control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-12F	Storwize V7000 Unified expansion enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Problem: Storwize V7000 Gen2 node canister service IP address unknown

You can use several methods to determine the service address of a node canister.

A default service address is initially assigned to each node canister, as shown in Table 76 on page 244. Try using these addresses if the node has not been reconfigured, and the addresses are valid on your network.

If you are able to access the management GUI, the service IP addresses of the node canisters are shown by selecting a node and port at **Settings > Network > Service IP Addresses**.

If you are unable to access the management GUI but you know the management IP address of the system, you can use the address to log in to the service assistant that is running on the configuration node:

1. Point your browser at the /service directory of the management IP address of the system. If your management IP address is 11.22.33.44, point your web browser to 11.22.33.44/service.
2. Log in to the service assistant.
3. The service assistant home page lists the node canister that can communicate with the node.

4. If the service address of the node canister that you are looking for is listed in the Change Node window, make the node the current node. Its service address is listed under the Access tab of the node details.

If you know the service IP address of any node canister in the system, you can log in to the service assistant of that node. Follow the previous instructions for using the service assistant, but at step 1, point your browser at the /service directory of the service IP address you know. If you know that a service IP address is 11.22.33.56, point your web browser to 11.22.33.56/service

Some types of errors can prevent nodes from communicating with each other; in that event, it might be necessary to point your browser directly at the service assistant of the node that requires administering, rather than change the current node in the service assistant.

If you cannot find the service address of the Storwize V7000 Gen2 node using the management GUI or service assistant of a different node, two options remain:

- You can connect directly to the service assistant of a node using the technician port of the node, as described in “Procedure: Initializing the Storwize V7000 Gen2 system using the technician port” on page 274.
- You can also use a USB flash drive to find the service address of the node. For more information, see “Procedure: Getting node canister and system information using a USB flash drive” on page 255.

Table 76. Default service IP addresses

Canister and port	IPv4 address	IPv4 subnet mask
Canister 1 (left) port 1 (left)	192.168.70.121	255.255.255.0
Canister 2 (right) port 1 (left)	192.168.70.122	255.255.255.0

Problem: node canister service IP address unknown

You can use several methods to determine the service address of a node canister.

A default service address is initially assigned to each node canister, as shown in Table 77 on page 245. Try using these addresses if the node has not been reconfigured, and the addresses are valid on your network.

If you are able to access the management GUI, the service IP addresses of the node canisters are shown by selecting a node and port at **Settings > Network > Service IP Addresses**.

If you are unable to access the management GUI but you know the management IP address of the system, you can use the address to log into the service assistant that is running on the configuration node.

1. Point your browser at the /service directory of the management IP address of the system. If your management IP address is 11.22.33.44, point your web browser to 11.22.33.44/service.
2. Log into the service assistant.
The service assistant home page lists the node canister that can communicate with the node.
3. If the service address of the node canister that you are looking for is listed in the Change Node window, make the node the current node. Its service address is listed under the Access tab of the node details.

If you know the service IP address of any node canister in the system, you can log into the service assistant of that node. Follow the previous instructions for using the service assistant, but at step 1, point your browser at the /service directory of the service IP address you know. If you know a service IP address is 11.22.33.56, point your web browser to 11.22.33.56/service.

Some types of errors can prevent nodes from communicating with each other; in that event, it might be necessary to point your browser directly at the service assistant of the node that requires administering, rather than change the current node in the service assistant.

If you are unable to find the service address of the node using the management GUI or service assistant, you can also use a USB flash drive to find it. For more information, see “Procedure: Getting node canister and system information using a USB flash drive” on page 255.

Table 77. Default service IP addresses

Canister and port	IPv4 address	IPv4 subnet mask
Canister 1 (left) port 1 (left)	192.168.70.121	255.255.255.0
Canister 2 (right) port 1 (left)	192.168.70.122	255.255.255.0

Problem: Cannot connect to the service assistant

Use this information if you are unable to display the service assistant on your browser via the service IP address. Note that you can use the Technician port to access the service assistant GUI on a Storwize V7000 Gen2 node canister.

You might encounter a number of situations when you cannot connect to the service assistant.

- Check that you have entered the /service path after the service IP address. Point your web browser to *service IP address/service* for the node that you want to work on. For example, if you set a service address of 11.22.33.44 for a node canister, point your browser to 11.22.33.44/service.
- Check that you are using the correct service address for the node canister. To find the IPv4 and IPv6 addresses that are configured on the node, go to “Problem: node canister service IP address unknown” on page 244. Try accessing the service assistant through these addresses. Verify that the IP address, subnet, and gateway are specified correctly for IPv4 addresses. Verify that the IP address, prefix, and gateway are specified for the IPv6 addresses. If any of the values are incorrect, see “Procedure: Changing the service IP address of a node canister” on page 271.
- You cannot connect to the service assistant if the node canister is not able to start the code. To verify that the LEDs indicate that the code is active, see “Procedure: Understanding the Storwize V7000 Gen1 system status using the LEDs” on page 262.
- The service assistant is configured on Ethernet port 1 of a node canister. Verify that an Ethernet cable is connected to this port and to an active port on your Ethernet network. See “Procedure: Finding the status of Storwize V7000 Gen1 Ethernet connections” on page 268 for details.
- Ping the service address to see if the Ethernet network permits the connection. If the ping fails, check the Ethernet network configuration to see if there is a routing or a firewall issue. Check that the Ethernet network configuration is compatible with the gateway and subnet or prefix settings. Check that you have not used an address that is used by another device on your Ethernet network. If

necessary, change the network configuration or see “Procedure: Changing the service IP address of a node canister” on page 271 to change the service IP address of a node.

- A default service address is initially assigned to each node canister. The service IP address 192.168.70.121 subnet mask 255.255.255.0 is preconfigured on Ethernet port 1 of the upper canister, canister 1. The service IP address 192.168.70.122 subnet mask 255.255.255.0 is preconfigured on Ethernet port 1 of the lower canister, canister 2.

You might not be able to access these addresses because of the following conditions:

- These addresses are the same as the addresses that are used by other devices on the network.
- These addresses cannot be accessed on your network.
- There are other reasons why they are not suitable for use on your network.

If the previous conditions apply, see “Procedure: Changing the service IP address of a node canister” on page 271 to change the service IP address to one that works in your environment.

If you are unable to change the service address, for example, because you cannot use a USB flash drive in the environment, see “Procedure: Accessing a Storwize V7000 Gen1 canister using a directly attached Ethernet cable” on page 276.

Problem: Management GUI or service assistant does not display correctly

If the Management GUI or the service assistant does not display correctly, verify that you are using a supported web browser.

For a list of supported browsers, see http://pic.dhe.ibm.com/infocenter/storwize/ic/topic/com.ibm.storwize.v7000.730.doc/svc_configuringbrowser_1obg15.html.

Problem: A node canister has a location node error

The node error listed on the service assistant home page or in the event log can indicate a location error.

A location error means that the node canister or the enclosure midplane has been moved or changed. This is normally due to a service action not being completed or not being implemented correctly.

A number of different conditions are reported as location errors. Each condition is indicated by different node error. To find out how to resolve the node error, go to “Procedure: Fixing node errors” on page 271.

Be aware that after a node canister has been used in a system, the node canister must not be moved to a different location, either within the same enclosure or in a different enclosure because this might compromise its access to storage, or a host application's access to volumes. Do not move the canister from its original location unless directed to do so by a service action.

Problem: SAS cabling not valid

Use this procedure if you receive errors to determine if your SAS cabling is valid. The procedure differs, depending on the generation of your control enclosure model.

About this task

Storwize V7000 Unified *Gen1* refers to the enclosure models in the following table:

Table 78. Storwize V7000 Unified Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 Unified control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 Unified control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 Unified control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 Unified control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 Unified expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Storwize V7000 Unified *Gen2* refers to the newer generation of enclosures in the following table:

Table 79. Storwize V7000 Unified Gen2 model numbers

Machine type/model	Description
2076-524	Storwize V7000 Unified control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-12F	Storwize V7000 Unified expansion enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Problem: Storwize V7000 Gen2 SAS cabling not valid

This topic provides information to be aware of if you receive errors that indicate the SAS cabling is not valid.

Check the following items:

- No more than 10 expansion enclosures can be chained to port 1 (below the control enclosure). The connecting sequence from SAS port 1 of the node canister is called chain 1.
- No more than 10 expansion enclosures can be chained to port 2 (above the control enclosure). The connecting sequence from SAS port 2 of the node canister is called chain 2.
- Do not connect a SAS cable between a port on a left canister and a port on a right canister.
- In any enclosure, the same ports must be used on both canisters.
- Do not connect a SAS cable between ports in the same enclosure.
- In each enclosure, where a cable connects from SAS port 1 of the left canister to another enclosure, a cable must also connect from SAS port 1 of the right canister to the other enclosure. Similarly, where a cable connects from SAS port 2 of the left canister to another enclosure, a cable must also connect from SAS port 2 of the right canister to the other enclosure.

- For cables connected between expansion enclosures, one end is connected to port 1 while the other end is connected to port 2.
- For cables that are connected between a control enclosure and expansion enclosures, port 1 must be used on the expansion enclosures.
- The last enclosure in a chain must not have cables in port 2 of the left canister, nor port 2 of the right canister.
- Ensure that each SAS cable is fully inserted.

See the topic about installing SAS cables in the *IBM Storwize V7000 Gen2 Quick Installation Guide*.

Problem: Storwize V7000 Gen1 SAS cabling not valid

This topic provides information to be aware of if you receive errors that indicate the SAS cabling is not valid.

Check the following items:

- No more than five expansion enclosures can be chained to port 1 (below the control enclosure). The connecting sequence from port 1 of the node canister is called chain 1.
- No more than four expansion enclosures can be chained to port 2 (above the control enclosure). The connecting sequence from port 2 of the node canister is called chain 2.
- Do not connect a SAS cable between a port on an upper canister and a port on a lower canister.
- In any enclosure, the same ports must be used on both canisters.
- No SAS cable can be connected between ports in the same enclosure.
- For any enclosure, the cables that are connected to SAS port 1 on each canister must attach to the same enclosure. Similarly, for any enclosure, the cables that are connected to SAS port 2 on each canister must attach to the same enclosure. Cable attachments for SAS port 1 and cable attachments for SAS port 2 do not go to the same enclosure.
- For cables connected between expansion enclosures, one end is connected to port 1 while the other end is connected to port 2.
- For cables that are connected between a control enclosure and expansion enclosures, port 1 must be used on the expansion enclosures.
- The last enclosure in a chain must not have cables in port 2 of canister 1 and port 2 of canister 2.
- Ensure that each SAS cable is fully inserted.

Problem: New expansion enclosure not detected

Determine why a newly installed expansion enclosure was not detected by the system.

When you install a new expansion enclosure, follow the management GUI Add Enclosure wizard. Select **Monitoring > System**. From the **Actions** menu, select **Add Enclosures**.

If the expansion enclosure is not detected, complete the following verifications:

- Verify the status of the LEDs at the back of the expansion enclosure. At least one power supply unit must be on with no faults shown. At least one canister must be active, with no fault LED on, and all the serial-attached SCSI (SAS) port 1

LEDs must be on. For details about the LED status, see “Procedure: Understanding the Storwize V7000 Gen1 system status using the LEDs” on page 262.

- Verify that the SAS cabling to the expansion enclosure is correctly installed. To review the requirements, see “Problem: Storwize V7000 Gen1 SAS cabling not valid” on page 248.

Problem: Control enclosure not detected

If a control enclosure is not detected by the system, this procedure can help you resolve the problem.

When installing a new control enclosure, use the **Add Enclosures** wizard in the management GUI. To access this wizard, select **Monitoring > System**. On the **Systems** page, select **Actions > Add Enclosures**.

If the control enclosure is not detected, check the following items:

- The enclosure is powered on.
- The enclosure is not part of another system.
- At least one node is in candidate state.
- The Fibre Channel cables are connected and zoning is set up according to the zoning rules defined in the *SAN configuration and zoning rules summary* topic. There must be a zone that includes all ports from all node canisters.
- The existing system and the nodes in the enclosure that are not detected have version 6.3 or later installed.

Problem: Mirrored volume copies no longer identical

The management GUI provides options to either check copies that are identical or to check that the copies are identical and to process any differences that are found.

To confirm that the two copies of a mirrored volume are still identical, choose the volume view that works best for you. Select one of the volume copies in the volume that you want to check. From the **Actions** menu, select the **Validate Volume Copies** option.

You have the following choices:

- Validate that the volume copies are identical.
- Validate that the volume copies are identical, mark, and repair any differences that are found.

If you want to resolve any differences, you have the following options:

- Consider that the primary volume copy is correct and make the other volume copy match the primary volume copy if any differences are found. The primary volume copy is the copy that is considered correct.
- Do not assume that either volume copy is correct. If a difference is found, the sector is marked. A media error is returned if the volume is read by a host application.

Problem: Command file not processed from USB flash drive

This information assists you in determining why the command file is not being processed, when using a USB flash drive.

You might encounter this problem during initial setup or when running commands if you are using your own USB flash drive rather than the USB flash drive that was packaged with your order.

If you encounter this situation, verify the following items:

- That an `satask_result.html` file is in the root directory on the USB flash drive. If the file does not exist, then the following problems are possible:
 - The USB flash drive is not formatted with the correct file system type. Use any USB flash drive that is formatted with FAT32 file system on its first partition; for example, NTFS is not a supported type. Reformat the key or use a different key.
 - The USB port is not working. Try the key in the other USB port.
 - The node is not operational. Check the status using the light-emitting diodes (LEDs). See “Procedure: Understanding the Storwize V7000 Gen1 system status using the LEDs” on page 262.
- If there is a `satask_result.html` file, check the first entry in the file. If there is no entry that matches the time the USB flash drive was used, it is possible that the USB port is not working or the node is not operational. Check the node status using the LEDs. See “Procedure: Understanding the Storwize V7000 Gen1 system status using the LEDs” on page 262.
- If there is a status output for the time the USB flash drive was used, then the `satask.txt` file was not found. Check that the file was named correctly. The `satask.txt` file is automatically deleted after it has been processed.

Procedure: Resetting superuser password

You can reset the superuser password to the default password of `passwd0rd` by using a special command action. The password procedure differs, depending on the generation of your control enclosure model.

About this task

Storwize V7000 Unified *Gen1* refers to the enclosure models in the following table:

Table 80. Storwize V7000 Unified Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 Unified control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 Unified control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 Unified control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 Unified control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 Unified expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Storwize V7000 Unified *Gen2* refers to the newer generation of enclosures in the following table:

Table 81. Storwize V7000 Unified Gen2 model numbers

Machine type/model	Description
2076-524	Storwize V7000 Unified control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-12F	Storwize V7000 Unified expansion enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Procedure: Resetting the superuser password for Storwize V7000 Gen2

The primary method for resetting the superuser password is to change the password as you log in, with the link on the log-in page. You can also access the service assistant from the technician port to change the password.

If the password reset function is enabled, the log-in page displays a link for resetting the password. You can also use the technician port to access the Storwize V7000 Gen2 service assistant.

If the password reset function is not enabled, there is no work-around. You must contact the person who knows the password. The USB flash drive interface also supports resetting the password.

Procedure: Resetting the Storwize V7000 Gen1 superuser password

You can reset the superuser password to the default password of `passw0rd` by using a USB flash drive command action.

About this task

You can use this procedure to reset the superuser password if you have forgotten the password. This command runs differently depending on whether you run it on a node canister that is active in a clustered system.

Note: If a node canister is not in active state, the superuser password is still required to log on to the service assistant.

It is possible to configure your system so that resetting the superuser password with the USB flash drive command action is not permitted. If your system is configured this way, there is no work-around. Contact the person who knows the password.

To use a USB flash drive to reset the superuser password, see “USB flash drive and Initialization tool interface” on page 206.

See also “Problem: Unable to log on to the management GUI” on page 240.

Results

If the node canister is active in a clustered system, the password for superuser is changed on the clustered system. If the node canister is not in active state, the superuser password for the node canister is changed. If the node canister joins a clustered system later, the superuser password is reset to that of the clustered system.

Procedure: Identifying which enclosure or canister to service

Use this procedure to identify which enclosure or canister must be serviced. Identification procedures differ, depending on the generation of your control enclosure model.

About this task

Storwize V7000 Unified *Gen1* refers to the enclosure models in the following table:

Table 82. Storwize V7000 Unified Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 Unified control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 Unified control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 Unified control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 Unified control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 Unified expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Storwize V7000 Unified *Gen2* refers to the newer generation of enclosures in the following table:

Table 83. Storwize V7000 Unified Gen2 model numbers

Machine type/model	Description
2076-524	Storwize V7000 Unified control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-12F	Storwize V7000 Unified expansion enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Procedure: Identifying which Storwize V7000 Gen2 enclosure or canister to service

Use this procedure to identify which Storwize V7000 Gen2 enclosure or canister must be serviced.

About this task

To prevent loss of access to data or loss of data when servicing your system, it is important to be able to identify the correct enclosure or canister when completing a service action. Each enclosure is identified by its model type and serial number. Model type and serial number are indicated by labels on the enclosure front and rear.

Each canister is identified by its enclosure and slot location. Viewing an enclosure from the rear, slot 1 is on the left and slot 2 is on the right. There are physical differences between Storwize V7000 Gen2 control enclosures and expansion enclosures.

Looking at the front of a rack:

- The type of the enclosure, either Control or Expansion, is labeled on the left bezel.
- The model type and serial number of the enclosure are found at the bottom of the left bezel.

Storwize V7000 Unified Gen2 refers to the newer generation of enclosures in the following table:

Table 84. Storwize V7000 Unified Gen2 model numbers

Machine type/model	Description
2076-524	Storwize V7000 Unified control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-12F	Storwize V7000 Unified expansion enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Looking at the rear of a rack:

- Control enclosures contain tall and narrow power supplies at the far left and right of the enclosure, with the node canisters installed in between, side by side.
- Expansion enclosures contain short node canisters installed side by side, above short power supply units also installed side by side.
- Control enclosures contain node canisters that have Ethernet ports, and USB ports. Expansion enclosures do not have any of these ports. The model type is shown on a label.
- Each canister is identified by its slot location. Viewing an enclosure from the rear, slot 1 is on the left and slot 2 is on the right.

A canister is uniquely identified by the enclosure that it is in and its slot location. This ID is shown as E-C or E|C where E is the enclosure ID and C is the canister location. On the service assistant, the ID is known as the Panel.

Notes:

- When a node canister is added to a clustered system as a node, it is given a node name and a node ID. The default node name is nodeN, where N is the node ID. The node ID does not represent the slot location of the node. On the **Monitoring > System** page use the dynamic graphic to show the back of the system. Hover over the canister to display the node name and canister location. The service assistant home page also shows both the node name and the canister location. If you have only the node name, use these panels to determine the node canister location.
- It is good practice to use this procedure to identify which enclosure or canister must be serviced, as completing a service action on the wrong canister can lead to loss of access to data or loss of data.

To control the identify LED of an enclosure or online canister, use the management GUI:

1. Log in to the management GUI for the system.
2. Select the **Monitoring > System** panel.
3. Select the canister or enclosure to be identified.
4. Select **Action > Identify** to control the identify LEDs for the component.

Alternatively, if a node canister is not online to the system, use the service assistant to control the identify LED.

1. Log in to the service assistant of the node canister to be identified.
2. Click **Identify** at the upper left of the page to control the identify LEDs.

Procedure: Identifying which Storwize V7000 Gen1 enclosure or canister to service

Before completing service actions on an enclosure or canister, verify that you are working on the correct enclosure or canister.

About this task

Procedure

Use the following options to identify an enclosure. An enclosure is identified by its ID and serial number.

Procedure: Checking the status of your system

Use this procedure to verify the status of objects in your system using the management GUI. If the status of the object is not online, view the alerts and run the recommended fix procedures.

About this task

Volumes normally show offline because another object is offline. A volume is offline if one of the MDisk that makes up the storage pool that it is in is offline. You do not see an alert that relates to the volume; instead, the alert relates to the MDisk. Performing the fix procedures for the MDisk enables the volume to go online.

Procedure

Use the following management GUI functions to find a more detailed status:

- **Monitoring > System**
- **Pools > MDisk by Pools**
- **Volumes > Volumes**
- **Monitoring > Events**, and then use the filtering options to display alerts, messages, or event types.

Procedure: Getting node canister and system information using the service assistant

This procedure explains how to view information about the node canisters and system using the service assistant.

About this task

To obtain the information:

1. Log on to the service assistant.
2. View the information about the node canister to which you connected or the other node canisters in the enclosure. To change which node's information is shown, select the node in the **Change Node** table of the Home page.

The Home page shows a table of node errors that exist on the node canister and a table of node details for the current node. The node errors are shown in priority order.

The node details are divided into several sections. Each section has a tab. Examine the data that is reported in each tab for the information that you want.

- The Node tab shows general information about the node that includes the node state and whether it is a configuration node.
- The Hardware tab shows information about the hardware.
- The Access tab shows the management IP addresses and the service addresses for this node.
- The Location tab identifies the enclosure in which the node canister is located.
- The Ports tab shows information about the I/O ports.

Procedure: Getting node canister and system information using a USB flash drive

You can view information about the node canister and system using a USB flash drive.

About this task

Use any USB flash drive with a FAT32 file system on its first partition.

1. Ensure that the USB flash drive does not contain a file named `satask.txt` in the root directory.

If `satask.txt` does exist in the directory, the node attempts to run the command that is specified in the file. The information that is returned is appended to the `satask_result.html` file. Delete this file if you no longer want the previous output.

Procedure

1. Insert the USB flash drive in the USB port of the node from which you want to collect data. The node fault light-emitting diode (LED) flashes while information is collected and written to the USB flash drive.
2. Wait until the LED stops flashing before removing the USB flash drive. Because the LED is a fault indicator, it might remain permanently on or off.
3. View the results in file `satask_result.html` in a web browser. The file contains the details and results of the command that was run and the status and the configuration information from the node canister.

Procedure: Understanding the system status using the LEDs

Use this procedure to determine the system status using the LED indicators on the system. The procedure differs, depending on the generation of your control enclosure model.

About this task

Storwize V7000 Unified *Gen1* refers to the enclosure models in the following table:

Table 85. Storwize V7000 Unified Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 Unified control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 Unified control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 Unified control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 Unified control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 Unified expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Storwize V7000 Unified *Gen2* refers to the newer generation of enclosures in the following table:

Table 86. Storwize V7000 Unified Gen2 model numbers

Machine type/model	Description
2076-524	Storwize V7000 Unified control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-12F	Storwize V7000 Unified expansion enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Procedure: Understanding the Storwize V7000 Gen2 system status from the LEDs

To determine the Storwize V7000 2076-524 system status using the LED indicators on a control enclosure, use this procedure.

About this task

To understand the status of the I/O port at the rear of a control enclosure, refer to the topic about Storwize V7000 2076-524 node canister ports and indicators that is linked at the end of this topic.

Status indicators on the front of a control enclosure are described in the topic about components in the front of the enclosure that is linked at the end of this topic.

A detailed view of the system state is provided in the Monitoring sections of the management GUI and by the service assistant. If neither the management GUI nor the service assistant is accessible, use this procedure to determine the system status using the LED indicators on the control enclosures.

The system status LEDs visible at the rear of each control enclosure can show one of several states, as described in Table 87 on page 257.

Table 87. LED state descriptions used in the Storwize V7000 2076-524 enclosure

State description	Detail
Off	The LED is continuously not lit.
Flashing slowly	The LED turns on and off at a frequency of 1 Hz: It is on for 500 ms, then off for 500 ms, then repeats.
Flashing	The LED turns on and off at a frequency of 2 Hz: It is on for 250 ms, then off for 250 ms, then repeats.
Flashing fast	The LED turns on and off at a frequency of 4 Hz: It is on for 125 ms, then off for 125 ms, then repeats.
On	The LED is continuously lit.
Flashing	The LED is lit to indicate some activity, then turns off. The rate and duration that the LED is lit depends on the rate and duration of the activity.

Procedure

Complete the following steps to understand when a node canister is not participating in the system, and what remedial action to take.

1. Identify the control enclosures for the system you are troubleshooting.
To understand which are the control enclosures, refer to the topic about identifying the Storwize V7000 2076-524 enclosure or canister to service.
2. Use Table 88 to check the status of each power supply unit (PSU) in a control enclosure.
3. If at least one PSU is providing power to a control enclosure, use Table 89 on page 258 to check the status of each node canister in that enclosure.

Table 88. Understanding the power supply unit LEDs



LED state			Action
! Fault (amber)	 ac power (green)	 dc power (green)	
On	(any)	(any)	Replace the power supply unit, as described in “Replacing a Storwize V7000 Gen2 power supply unit for a control enclosure” on page 306.
Off	On	On	The power supply is functioning normally, providing power to the enclosure.
		Off	Replace the power supply unit, as described in “Replacing a Storwize V7000 Gen2 power supply unit for a control enclosure” on page 306.
	Off	On	Replace the power supply unit, as described in “Replacing a Storwize V7000 Gen2 power supply unit for a control enclosure” on page 306.
		Off	The power supply is not receiving power from the power outlet through the power cord. To power on the system, connect the power supply to an outlet in use for the power cord and turn on the power outlet.
Note: The fourth status LED on the power supply unit is not used.			

Table 89. Understanding the node canister status LEDs




LED state			Description
 Power (green)	 Fault (amber)	 System status (green)	
Off	(any)	(any)	<p>Power is not being provided by any power supply unit (PSU) in the enclosure, or power is coming from the battery in the canister.</p> <p>If at least one PSU is powering the enclosure, the canister or the enclosure midplane might be faulty.</p> <p>If both node canisters in an enclosure indicate this state at the same time, it is more likely that the enclosure midplane is faulty. Reseat any node canister in this state, as described in “Procedure: Reseating a Storwize V7000 Gen2 node canister” on page 278. If the condition persists, replace the node canister. If just one canister is affected, see “Replacing a Storwize V7000 Gen2 node canister” on page 296. If both canisters are affected, see “Replacing a Storwize V7000 Gen2 control enclosure midplane assembly” on page 353.</p>
Flashing slowly			<p>Power is available, but the canister is powered off. (The main CPU is not running.)</p> <p>Restart the node canister, as described in “Procedure: Reseating a Storwize V7000 Gen2 node canister” on page 278.</p>
Flashing fast			<p>The node canister is doing a self test during start-up.</p> <p>Wait for the canister to complete its start-up sequence.</p>

Table 89. Understanding the node canister status LEDs (continued)


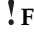

LED state			Description
 Power (green)	 Fault (amber)	 System status (green)	
On	Off	Off	<p>The node canister is in standby mode. (The Storwize V7000 Unified software is not running.)</p> <p>It is safe to remove or reseal the canister.</p> <p>Restart the node canister, as described in “Procedure: Reseating a Storwize V7000 Gen2 node canister” on page 278.</p>
		Flashing	<p>The node canister is in candidate state. The node is not doing I/O in the system.</p> <p>Unless indicated by the battery status LED, it is safe to remove or reseal the canister. See Table 90 on page 261.</p> <p>If the node state is candidate and the system is running on the other node canister, the candidate node is automatically added to the system.</p> <p>If both node canisters are in candidate state, determine whether you must recover the system or create a new system.</p>
		Flashing fast	<p>The node is doing an emergency shutdown operation, using the battery for power, after detecting a loss of power from the power supply units.</p> <p>Wait for the emergency shutdown operation to complete. If the partner node has also done an emergency shutdown operation, there was most likely a loss of input power to both enclosure power supply units and the system restarts when the input power is restored. Otherwise, there might be a fault the node canister, enclosure midplane, or power supply units.</p>
		On	<p>The Storwize V7000 Unified software is running, and the node canister is participating in the system.</p> <p>The canister must not be removed.</p> <p>This is the normal operational LED state.</p>
On	Flashing	Off	<p>The node canister is in standby mode. (The Storwize V7000 Unified software is not running.)</p> <p>It is safe to remove or reseal the canister.</p>
		Flashing	<p>The Identify function for this canister has been activated. To determine if it is safe to remove the canister, use the management GUI or service assistant to turn off the Identify function, then check the node canister status LEDs again.</p>

Table 89. Understanding the node canister status LEDs (continued)




LED state			Description
 Power (green)	 Fault (amber)	 System status (green)	
On	Flashing	Flashing fast	<p>The Identify function for this canister has been activated.</p> <p>The node is doing an emergency shutdown operation, using the battery for power, after detecting a loss of power from the power supply units.</p> <p>Wait for the emergency shutdown operation to complete. If the partner node has also done an emergency shutdown operation, there was most likely a loss of input power to both enclosure power supply units and the system restarts when the input power is restored. Otherwise, there might be a fault in the power supply units or the node canister.</p>
On	Flashing	On	<p>The Identify function for this canister was activated.</p> <p>To determine if it is safe to remove the canister, use the management GUI or service assistant to turn off the Identify function, then check the node canister status LEDs, again.</p>
On	On	Off	<p>The Storwize V7000 Unified software is not running. The BIOS might have detected a fault.</p> <p>It is safe to remove or reseal the canister.</p> <p>Try reseating the canister, as described in “Procedure: Reseating a Storwize V7000 Gen2 node canister” on page 278. If the canister still shows this fault, replace the node canister, as described in “Replacing a Storwize V7000 Gen2 node canister” on page 296.</p>
On	On	Flashing	<p>The node is in service state.</p> <p>It is safe to remove or reseal the canister.</p> <p>See Table 90 on page 261 to determine whether the battery charge is insufficient.</p> <p>Use the service assistant to identify any node errors or to determine what to do, see “Procedure: Fixing node errors” on page 271.</p> <p>If the node is able to communicate with the configuration node, there might also be an error alert in the system event log, in which case you should run the associated fix procedure.</p>
On	On	Flashing fast	<p>The node is in service state.</p> <p>There is a code update in progress.</p> <p>The canister must not be removed.</p> <p>No action is required until the code update completes.</p>

Table 89. Understanding the node canister status LEDs (continued)


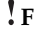

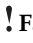

LED state			Description
 Power (green)	 Fault (amber)	 System status (green)	
On	On	On	<p>The Storwize V7000 Unified software is running but there might be an error alert in the event log, such as error code 550.</p> <p>The canister must not be removed.</p> <p>If possible, go to the management GUI and run the fix procedure for the error alerts listed there. If this is not possible, refer to the service actions for node error 550.</p>

Table 90. Understanding the node canister battery status LEDs

LED state		Description
 Fault (amber)	 Status (green)	
On	On	<p>A battery fault exists that caused the node canister to do an emergency shutdown operation (storing its cache data to the system disk). The node canister is in service state or is going into service state. The service assistant shows a node error.</p> <p>Replace the battery in the node canister, as described in “Replacing the battery in a Storwize V7000 Gen2 node canister” on page 316.</p>
On	Off	<p>A battery fault exists that caused the node canister to shut down to service state. There was insufficient power in the battery or the supply for the node canister to do an emergency shutdown operation. The service assistant shows a node error.</p> <p>Replace the battery in the node canister, as described in “Replacing the battery in a Storwize V7000 Gen2 node canister” on page 316.</p>
Off	Flashing fast	The battery is charging and has insufficient charge to protect against a single ac loss. The node is held in service state with a fatal node error until the battery has charged.
	Flashing	<p>The battery has sufficient charge to complete one emergency shutdown operation.</p> <p>This state does not prevent the canister from participating in the system.</p> <p>The battery continues to charge until it is able to complete two emergency shutdown operations. No action is necessary.</p>
	On	<p>The battery is optimally charged, such that the node canister is able to complete two emergency shutdown operations.</p> <p>This state does not prevent the canister from participating in the system. The canister continues to manage the battery charge.</p> <p>No action is necessary.</p>

- Repeat steps 2 on page 257 and 3 on page 257 for each control enclosure in the system.

Procedure: Understanding the Storwize V7000 Gen1 system status using the LEDs

Use this procedure to determine the Storwize V7000 Gen1 system status using the LED indicators on the system.

About this task

The LEDs provide a general idea of the system status. You can obtain more detail from the management GUI and the service assistant. Examine the LEDs when you are not able to access the management GUI or the service assistant, or when the system is not showing any information about a device.

The procedure shows the status for the enclosure chassis, power supply units and batteries, and canisters. It does not show the status for the drives.

The first step is to determine the state of the control enclosure, which includes its power supply units, batteries, and node canisters. Your control enclosure is operational if you can manage the system using the management GUI. You might also want to view the status of the individual power supply units, batteries, or node canisters.

Find the control enclosure for the system that you are troubleshooting. There is one control enclosure in a system. If you are unsure which one is the control enclosure, go to “Procedure: Identifying which Storwize V7000 Gen1 enclosure or canister to service” on page 254.

Procedure

1. Use the state of the ac power failure, power supply OK, fan failure, and dc power failure LEDs on each power supply unit in the enclosure to determine if there is power to the system, or if there are power problems. Figure 73 on page 263 shows the LEDs on the power supply unit for the 2076-112 or 2076-124. The LEDs on the power supply units for the 2076-312 and 2076-324 are similar, but they are not shown here.

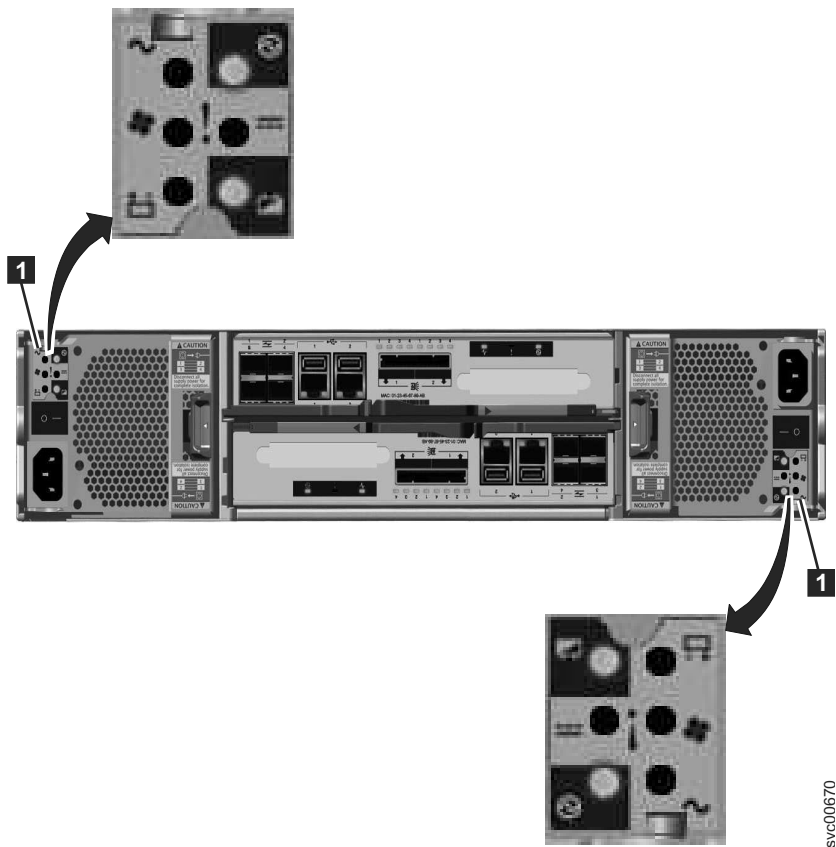


Figure 73. LEDs on the power supply units of the control enclosure

Table 91. Power-supply unit LEDs


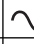






Power supply OK 	AC failure 	Fan failure 	DC failure 	Status	Action
On	On	On	On	Communication failure between the power supply unit and the enclosure chassis	Replace the power supply unit. If failure is still present, replace the enclosure chassis.
Off	Off	Off	Off	No ac power to the enclosure.	Turn on power.
Off	Off	Off	On	The ac power is on but power supply unit is not seated correctly in the enclosure.	Seat the power supply unit correctly in the enclosure.

Table 91. Power-supply unit LEDs (continued)

Power supply OK 	AC failure 	Fan failure 	DC failure 	Status	Action
Off	On	Off	On	No ac power to this power supply	<ol style="list-style-type: none"> 1. Check that the switch on the power supply unit is on. 2. Check that the ac power is on. 3. Reseat and replace the power cable.
On	Off	Off	Off	Power supply is on and operational.	No actions
Off	Off	On	Off	Fan failure	Replace the power supply unit.
Off	On	On	On	Communication failure and power supply problem	Replace the power supply unit. If replacing the power supply unit does not fix the problem, replace the enclosure chassis.
Flashing	X	X	X	No canister is operational.	Both canisters are either off or not seated correctly. Turn off the switch on both power supply units and then turn on both switches. If this action does not resolve the problem, remove both canisters slightly and then push the canisters back in.
Off	Flashing	Flashing	Flashing	Firmware is downloading.	No actions. Do not remove ac power. Note: In this case, if there is a battery in a power supply unit, its LEDs also flash.

2. At least one power supply in the enclosure must indicate Power supply OK or Power supply firmware downloading for the node canisters to operate. For this situation, review the three canister status LEDs on each of the node canisters. Start with the power LED.

Table 92. Power LEDs


Power LED status 	Description
Off	There is no power to the canister. Try reseating the canister. Go to “Procedure: Reseating a node canister” on page 278. If the state persists, follow the hardware replacement procedures for the parts in the following order: node canister, enclosure chassis.

Table 92. Power LEDs (continued)


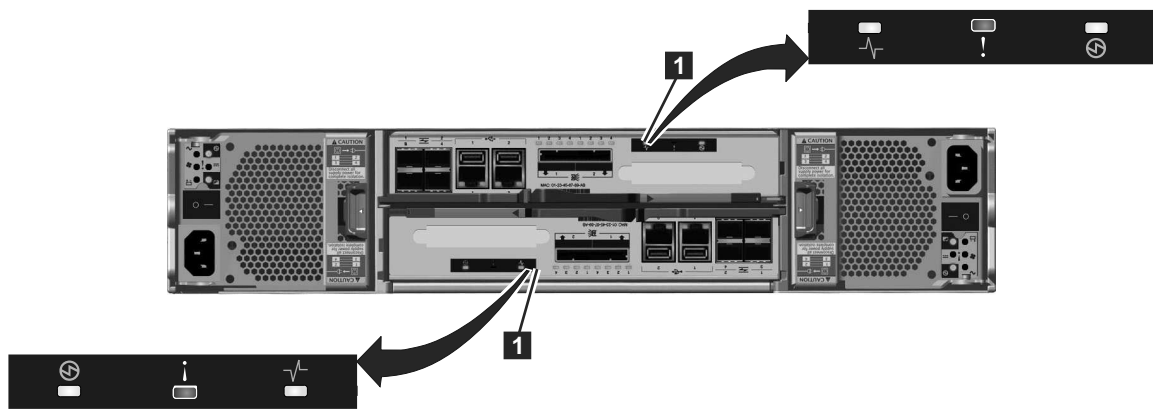
Power LED status 	Description
Slow flashing (1 Hz)	Power is available, but the canister is in standby mode. Try to start the node canister by reseating it. Go to “Procedure: Reseating a node canister” on page 278.
Fast flashing (2 Hz)	The canister is running its power-on self-test (POST). Wait for the test to complete. If the canister remains in this state for more than 10 minutes, try reseating the canister. Go to “Procedure: Reseating a node canister” on page 278. If the state persists, follow the hardware replacement procedure for the node canister.

Figure 74 shows the LEDs on the node canister.



svc00672

Figure 74. LEDs on the node canisters

- If the power LED is on, consider the states of the clustered-system status and fault LEDs.

Table 93. System status and fault LEDs







System status LED 	Fault LED 	Status 	Action
Off	Off	Code is not active.	<ul style="list-style-type: none"> Follow procedures for reviewing power LEDs. If the power LEDs show green, reseat the node canister. See “Procedure: Reseating a node canister” on page 278. If the LED status does not change, see “Replacing a Storwize V7000 Gen1 node canister” on page 297.
Off	On	Code is not active. The BIOS or the service processor has detected a hardware fault.	Follow the hardware replacement procedures for the node canister.
On	Off	Code is active. Node state is active.	No action. The node canister is part of a system and can be managed by the management GUI.

Table 93. System status and fault LEDs (continued)

System status LED 	Fault LED 	Status 	Action
On	On	Code is active and is in starting state. However, it does not have enough resources to form the system.	The node canister cannot become active in a system. There are no detected problems on the node canister itself. However, it cannot connect to enough resources to safely form a system. Follow the procedure to fix the node errors. Go to "Procedure: Fixing node errors" on page 271.
Flashing	Off	Code is active. Node state is candidate.	Create a system on the node canister, or add the node canister to the system. If the other node canister in the enclosure is in active state, it automatically adds this node canister into the system. A node canister in this state can be managed using the service assistant.
Flashing	On	Code is active. Node state is service.	The node canister cannot become active in a system. Several problems can exist: hardware problem, a problem with the environment or its location, or problems with the code or data on the canister. Follow the procedure to fix the node errors. Go to "Procedure: Fixing node errors" on page 271.
Any	Flashing	The node canister is being identified so that you can locate it.	The fix procedures in the management GUI might have identified the component because it requires servicing. Continue to follow the fix procedures. The service assistant has a function to identify node canisters. If the identification LED is on in error, use the service assistant node actions to turn off the LED.

Results

To review the status of the control enclosure batteries, see Table 94.

Table 94. Control enclosure battery LEDs




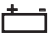
Battery Good 	Battery Fault 	Description	Action
On	Off	Battery is good and fully charged.	None
Flashing	off	Battery is good but not fully charged. The battery is either charging or a maintenance discharge is in process.	None
Off	On	Nonrecoverable battery fault.	Replace the battery. If replacing the battery does not fix the issue, replace the power supply unit.

Table 94. Control enclosure battery LEDs (continued)

Battery Good 	Battery Fault 	Description	Action
Off	Flashing	Recoverable battery fault.	None
Flashing	Flashing	The battery cannot be used because the firmware for the power supply unit is being downloaded.	None

Procedure: Finding the status of Ethernet connections

Use this procedure to find the status of Ethernet connections when you cannot connect. This procedure differs, depending on the generation of your control enclosure model.

About this task

Storwize V7000 Unified Gen1 refers to the enclosure models in the following table:

Table 95. Storwize V7000 Unified Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 Unified control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 Unified control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 Unified control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 Unified control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 Unified expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Storwize V7000 Unified Gen2 refers to the newer generation of enclosures in the following table:

Table 96. Storwize V7000 Unified Gen2 model numbers

Machine type/model	Description
2076-524	Storwize V7000 Unified control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-12F	Storwize V7000 Unified expansion enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Procedure: Finding the status of Storwize V7000 Gen2 Ethernet connections

Use this procedure to find the status of Ethernet connections in the Storwize V7000 Gen2 system when you cannot connect.

About this task

Ethernet port 1 of each node canister must be connected to an active port on your Ethernet network. You can determine the status of the Ethernet ports by using any of the following methods:

- Connect a personal computer directly to the node by following “Procedure: Accessing the service assistant from the technician port” on page 275t. In the service assistant the status, speed and MAC address for each port are shown in the **Ports** tab of the **Node Details** table. Any node errors are shown in the **Node Errors** table.
- Use a USB flash drive to obtain node configuration information (see “Procedure: Getting node canister and system information using a USB flash drive” on page 255). The results file contains the status, speed, and MAC address for each port; whether the node is the configuration node and any node errors being reported.
- Examine the LEDs of the Ethernet ports. The Ethernet ports on the left end of the rear of each node canister are 1 Gbps Ethernet ports. For these ports, the link state LED is ON if the link is connected.

Procedure

1. Verify that each end of the cable is securely connected.
2. Verify that the port on the Ethernet switch or hub is configured correctly.
3. Connect the cable to a different port on your Ethernet network.
4. Replace the Ethernet cable.
5. Review any node errors that are reported in the service assistant or on the USB flash drive. Follow fixing node errors for each node error that is reported.
6. Follow the hardware replacement procedure for a node canister.

Procedure: Finding the status of Storwize V7000 Gen1 Ethernet connections

This procedure explains how to find the status of the Ethernet connections when you cannot connect.

About this task

Ethernet port 1 must be connected to an active port on your Ethernet network. Determine the state of the Ethernet LEDs by using the following methods:

- If the node software is active on the node, use the USB flash drive to obtain the most comprehensive information for the node status. Go to “Procedure: Getting node canister and system information using a USB flash drive” on page 255. The status, speed, and MAC address are returned for each port. Information is returned that identifies whether the node is the configuration node and whether any node errors were reported.
- Examine the LEDs of the Ethernet ports. For the status of the LEDs, go to “Ethernet ports and indicators” on page 28.

Procedure

If your link is not connected, complete the following actions to check the port status each time until it is corrected or connected.

1. Verify that each end of the cable is securely connected.
2. Verify that the port on the Ethernet switch or hub is configured correctly.
3. Connect the cable to a different port on your Ethernet network.

4. If the status is obtained using the USB flash drive, review all the node errors that are reported.
5. Replace the Ethernet cable.

Procedure: Finding the status of Storwize V7000 Gen2 SAS connections

Find the status of the SAS connections between Storwize V7000 Gen2 canisters in different enclosures.

About this task

Ensure that the Storwize V7000 Unified machine code is active on the node before you begin this procedure. To determine if the machine code is active, see “Procedure: Understanding the Storwize V7000 Gen2 system status from the LEDs” on page 256.

Procedure

Determine the state of the SAS ports by using one of the following methods:

- Go to **Monitoring > System** in the management GUI. Use the dynamic image to display the rear of the system. Hover over each of the SAS ports on the canisters to display the status. A port with offline status indicates that its link is not connected.
 - It is normal for port 2 of the canisters in the expansion enclosure at the end of a SAS chain to be offline.
 - If no expansion enclosures are connected to a system, it is normal for port 4 of each canister in the control enclosure to be offline.
- **Attention:** The system can identify some SAS cabling errors and log an event to alert you to the error. Go to the **Monitoring > Events** page of the management GUI and identify alerts concerning hardware errors and SAS cabling errors. Run fix procedures in the recommended order.
- Determine the meaning of the LEDs of the SAS ports, as described in “Expansion canister ports and indicators” on page 32.
 - If the link LED is off, the link is not connected.
 - If the fault LED is on, the link is partially operational, with reduced performance. Consider the state of any other link between the two enclosures before servicing this link.
- To connect a link that is not connected, complete the following actions while checking the link status after each step until the link is connected:
 1. Ensure that both ends of each SAS cable are correctly inserted into their correct ports, as described in “Problem: Storwize V7000 Gen2 SAS cabling not valid” on page 247.
 2. Replace the SAS cable, as described in “Replacing a Storwize V7000 Gen2 expansion enclosure attachment SAS cable” on page 330.
 3. Replace the expansion canister at one end of the connection, as described in “Replacing a Storwize V7000 Gen2 expansion enclosure” on page 345.
 4. Replace the canister at the other end of the connection. If it is a node canister, see “Replacing a Storwize V7000 Gen2 node canister” on page 296.

Procedure: Removing system data from a node canister

You can safely remove system information from a node canister, if you follow the proper guidelines and procedure. The information that is removed includes configuration data, cache data, and location data.

About this task

Attention: Do not remove the system data from a node canister unless instructed to do so by a service procedure. Do not use this procedure to remove the system data from the only online node canister in a system. If the system data is removed or lost from all node canisters in the system, the system is effectively deleted. Attempting a system recovery procedure to restore a deleted system is not guaranteed to recover all of your volumes.

Procedure

1. Log into the service assistant of the node canister.
2. Use the service assistant node action to hold the node canister in service state.
3. Click **Manage System**; then click **Remove system data** to remove the system data from the node canister.

Results

The node canister restarts in service state.

What to do next

When you want the node canister to be active again, use the service assistant to leave service state. The node canister moves to candidate state, and can be added to the system. If the partner node canister is already active, the candidate node canister is added automatically.

Procedure: Deleting a system completely

You might need to completely remove all system information. When the procedure is finished, the system operates like a new installation. No data is retained.

About this task

Attention: This procedure makes all the volume data that you have on your system inaccessible. You cannot recover the data. This procedure affects all volumes that are managed by your system.

Do not continue unless you are certain that you want to remove all the volume data and configuration data from your system. This procedure is not used as part of any recovery action.

There are two stages to this procedure. First, the node canisters are reset. Second, the enclosure data is reset.

Procedure

1. Start the service assistant on one of the node canisters.
2. Use the service assistant node action to hold the node in service state.
3. Use the **Manage System** option to remove the system data from the node.
4. Repeat steps 1 through 3 on the second node canister in the enclosure.

5. On one node, open the service assistant **Configure Enclosure** and select the **Reset System ID** option. This action causes the system to reset.

Procedure: Fixing node errors

To fix node errors that are detected by node canisters in your system, use this procedure.

About this task

Node errors are reported in the service assistant when a node detects erroneous conditions in a node canister.

Procedure

1. Use the service assistant to obtain (and better understand) node canister and system information about the state of each node.
2. If possible, log into the management GUI and use the monitoring page to run the recommended fix procedure.
 - a. Follow the fix procedure instructions to completion.
 - b. Repeat this step for each subsequent recommended fix procedure.
3. If it is not possible to access the management GUI, or no recommended actions are listed, follow the identified user response for each reported node error.

Procedure: Changing the service IP address of a node canister

This procedure identifies many methods that you can use to change the service IP address of a node canister.

About this task

When you change an IPv4 address, you change the IP address, the subnet, mask, and gateway. When you change an IPv6 address, you change the IP address, prefix, and gateway.

Which method to use depends on the status of the system and the other node canisters in the system. Follow the methods in the order shown until you are successful in setting the IP address to the required value.

You can set an IPv4 address, an IPv6 address, or both, as the service address of a node. Enter the required address correctly. If you set the address to 0.0.0.0 or 0000:0000:0000:0000:0000:0000:0000:0000, you disable the access to the port on that protocol.

Procedure

Change the service IP address.

- Use the control enclosure management GUI when the system is operating and the system is able to connect to the node with the service IP address that you want to change.
 1. Select **Settings > Network** from the navigation.
 2. Select **Service IP Addresses**.
 3. Complete the panel. Be sure to select the correct node to configure.

- Use the service assistant when you can connect to the service assistant on either the node canister that you want to configure or on a node canister that can connect to the node canister that you want to configure:
 1. Make the node canister that you want to configure the current node.
 2. Select **Change Service IP** from the menu.
 3. Complete the panel.
- Use one of the following procedures if you cannot connect to the node canister from another node:
 - Use the initialization tool to write the correct command file to the USB flash drive. Go to “Using the initialization tool” on page 207.
 - Use a text editor to create the command file on the USB flash drive. Go to “Using a USB flash drive” on page 207.

Procedure: Initializing a clustered system with a USB flash drive without using the initialization tool

Use this procedure to initialize a clustered system using a USB flash drive when you do not have a workstation to run the initialization tool or you do not have a copy of the tool.

About this task

In these situations, you must manually create an `satask.txt` file on a USB flash drive to initialize your clustered system. Use the USB flash drive that was supplied with your system or any USB flash drive that is formatted with a FAT32 file system on its first partition. (For a complete list of commands you can use in a `satask.txt` file, see “`satask.txt` commands” on page 209.)

Note: The USB flash drive must include only the `satask.txt` file on the formatted first partition of the drive to initialize the system. If other files are included on the formatted first partition of the drive, it might cause the initialization to fail.

Procedure

1. Open a file editor that can create ASCII text files.
2. Create a file called `satask.txt`.
3. Add a single line of command text to the file.

If you are creating a clustered system with an IPv4 address, the command line is like the following string:

```
satask mkcluster -clusterip aaa.aaa.aaa.aaa
-gw ggg.ggg.ggg.ggg -mask mmm.mmm.mmm.mmm
```

where you must replace `aaa.aaa.aaa.aaa` with the management IP address, `ggg.ggg.ggg.ggg` with the network gateway address, and `mmm.mmm.mmm.mmm` with the subnet mask address.

If you are creating a clustered system with an IPv6 address, the command line is like the following string:

```
satask mkcluster -clusterip_6 aaaa:aaaa:aaaa:aaaa:aaaa:aaaa:aaaa:aaaa
-gw_6 gggg:gggg:gggg:gggg:gggg:gggg:gggg:gggg -prefix_6 pp
```

where you must replace `aaaa:aaaa:aaaa:aaaa:aaaa:aaaa:aaaa:aaaa` with the management IPv6 address, `gggg:gggg:gggg:gggg:gggg:gggg:gggg:gggg` with the network gateway IPv6 address, and `pp` with the prefix value.

For other command options, see “Create system command” on page 213.

4. Save the file to a USB flash drive.
5. Plug the USB flash drive into a USB port on a control canister.
6. The system detects the USB flash drive, reads the `satask.txt` file, runs the command, and writes the results to the USB flash drive. The `satask.txt` file is deleted after the command is run.
7. Wait for the fault LED on the node canister to stop flashing before removing the USB flash drive.
8. Remove the USB flash drive and insert it into your workstation to view the results.
9. Use a web browser to view the results file, `satask_result.html`.
Check that there were no errors returned by the command. If there is insufficient battery charge to protect the system, the clustered system creates successfully, but it does not start immediately. In the results, look for the `time_to_charge` field for the battery. The results provide an estimate of the time, in minutes, before the system can start. If the time is not 0, wait for the required time. Check that the node canister that you inserted the USB flash drive into has its clustered-state LED on permanently. For additional information, see “Procedure: Understanding the Storwize V7000 Gen1 system status using the LEDs” on page 262.
10. If the initialization was successful and the batteries are sufficiently charged, point a supported browser to the management IP address that you specified to start the management GUI. You see the management GUI logon panel.
11. Log on as superuser. Use `passwd` for the password.
12. Follow the on-screen instructions.

Results

For more information about using the USB flash drive, see “USB flash drive and Initialization tool interface” on page 206.

Procedure: Initializing a clustered system using the service assistant

Use this procedure to initialize a clustered system using the service assistant. This procedure differs, depending on the generation of your control enclosure model.

About this task

Storwize V7000 Unified *Gen1* refers to the enclosure models in the following table:

Table 97. Storwize V7000 Unified Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 Unified control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 Unified control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 Unified control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 Unified control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 Unified expansion enclosure for 3.5-inch drives

Table 97. Storwize V7000 Unified Gen1 model numbers (continued)

Machine type/model	Description
2076-224	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Storwize V7000 Unified Gen2 refers to the newer generation of enclosures in the following table:

Table 98. Storwize V7000 Unified Gen2 model numbers

Machine type/model	Description
2076-524	Storwize V7000 Unified control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-12F	Storwize V7000 Unified expansion enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Procedure: Initializing the Storwize V7000 Gen2 system using the technician port

To initialize a new Storwize V7000 Gen2 system, you must connect a personal computer to the technician port on the rear of a node canister and run the initialization tool.

Before you begin

You must have the following items:

- A personal computer with an Ethernet port that supports Dynamic Host Configuration Protocol (DHCP)
- A supported browser that is installed on the personal computer
- An Ethernet cable to connect the personal computer to the technician port

Procedure

To initialize the system, complete the following steps.

1. Ensure that the system is powered on.
2. Configure an Ethernet port on the personal computer to enable Dynamic Host Configuration Protocol (DHCP) configuration of its IP address and DNS settings.
3. Locate the Ethernet port that is labeled **T** on the rear of a node canister. This is the Technician port. Connect an Ethernet cable between the port of the personal computer that is configured in step 2 and the technician port. A few moments after the connection is made, the node uses DHCP to configure IP and DNS settings of the personal computer.
4. After the Ethernet port of the personal computer is connected, open a supported browser and browse to address <http://install>. The browser automatically opens the initialization tool.
5. Follow the instructions that are presented by the initialization tool to configure the system with a management IP address.
6. After you complete the initialization process, disconnect the cable between the personal computer and the technician port.

What to do next

The system can now be reached by opening a supported web browser and pointing it to `http://management_IP_address`.

Procedure: Initializing a Storwize V7000 Gen1 system using the service assistant

To initialize a Storwize V7000 Gen1 system using the service assistant rather than the USB flash drive, use this procedure.

About this task

Note: The service assistant gives you the option to create a clustered system only if the node state is candidate.

Procedure

To initialize a clustered system using the service assistant, complete the following steps.

1. Point your web browser to the service assistant address of a node canister. It is best to use the node canister in slot 1; when viewed from the rear of the control enclosure, the left node canister. The default service address for this canister is `192.168.70.121/service`.
2. Log on with the superuser password. The default password is `passwd`. If you cannot connect, see “Problem: Cannot connect to the service assistant” on page 245.
3. Select **Manage System**.
4. Enter the system name and the management IP address.
5. Click **Create System**.
6. Point a supported browser to the management IP address that you specified to start the management GUI. The management GUI logon panel is displayed.
7. Log on as superuser. Use `passwd` for the password.
8. Follow the on-screen instructions.

Results

Attention: Without a USB flash drive to service the system, it is not possible to reset the superuser password or to change the system IP addresses in the event of a fault that prevents access to the management interface. It is essential that you take steps to record this information for use in the event of a failure.

Procedure: Accessing the service assistant from the technician port

If a node canister is inaccessible through your administrative network, use this procedure to connect a personal computer directly to the node canister to access the service assistant.

About this task

This procedure starts the initialization tool if the enclosure is not part of a system because the following conditions are true:

- The node canister is in candidate state.
- No system details are configured.

- The partner node is not in active state.

Otherwise, this procedure starts the service assistant.

Procedure

To connect a personal computer directly to the node canister, complete the following steps:

1. Configure DHCP on the Ethernet port of the personal computer to connect to the node canister.
If the personal computer cannot support DHCP, configure static IPv4 address 192.168.0.2 on the port.
2. Connect an Ethernet cable between the port on the personal computer and the technician port.
The technician port is labeled **T** on the rear of the node canister.
3. Open a supported web browser on the personal computer and browse to this URL:
`http://192.168.0.1`
4. Complete the appropriate procedure.
 - If the initialization tool opens, complete the initialization as described in the installation procedure.
 - If the service assistant dialog opens, use it to service the node canister.
5. Log out of the service assistant.
6. Disconnect the Ethernet cable from the technician port.

Procedure: Accessing a Storwize V7000 Gen1 canister using a directly attached Ethernet cable

If you need to use a direct Ethernet connection to attach a personal computer to a Storwize V7000 Gen1 node canister to run the service assistant or to use the service CLI, use this procedure.

About this task

Follow this procedure if you are not authorized to use a USB flash drive in your data center and when the service address of your nodes cannot be accessed over your Ethernet network. This situation might occur for a new installation where the default service IP addresses cannot be accessed on your network.

The default service addresses are listed in “Problem: Cannot connect to the service assistant” on page 245.

Note: Do not attempt to use a directly attached Ethernet cable to a canister that is active in a clustered system. You might disrupt access from host applications or the management GUI. If the node is active, go to **Settings > Network** in the management GUI to set the service IP address to one that is accessible on the network.

Procedure

Complete the following steps to access a canister using a directly attached Ethernet cable.

1. Connect one end of an Ethernet cable to Ethernet port 1 of the node canister.

Note: A cross-over Ethernet cable is not required.

2. Connect the other end of the Ethernet cable directly to the Ethernet port on a personal computer that has a web browser installed.
3. Get the service IP address of the node canister attached at step 1 on page 276. If the service IP address is unknown, refer to “Problem: node canister service IP address unknown” on page 244.
4. Use the operating system tools on the computer to set the IP address and subnet mask of the Ethernet port that is used in step 2. Set them to the same subnet of the node canister service IP address.
5. Point the web browser to the service IP address for the node canister.
6. Log on with the superuser password. The default password is `passwd`.
7. Set the service address of the node canister to one that can be accessed on the network as soon as possible.
8. Wait for the action to complete.
9. Disconnect your personal computer.
10. Reconnect the node canister to the Ethernet network.

Problem: Reseating a node canister

Use this procedure to reseat a node canister. The procedure differs, depending on the generation of your control enclosure model.

About this task

Storwize V7000 Unified *Gen1* refers to the enclosure models in the following table:

Table 99. Storwize V7000 Unified Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 Unified control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 Unified control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 Unified control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 Unified control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 Unified expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Storwize V7000 Unified *Gen2* refers to the newer generation of enclosures in the following table:

Table 100. Storwize V7000 Unified Gen2 model numbers

Machine type/model	Description
2076-524	Storwize V7000 Unified control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-12F	Storwize V7000 Unified expansion enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Procedure: Reseating a Storwize V7000 Gen2 node canister

Use this procedure to reseat a Storwize V7000 Gen2 node canister that is in service state or because of a service action that requires that the node canister be resealed.

About this task

Verify that you are reseating the correct node canister and that you use the correct canister handle for the node that you are reseating. A handle for each node canister is located above the canister.

Procedure

1. Verify the clustered-system status LED on the node canister. If it is permanently on, the node is active. If the node is active, no reseating is required.
2. Verify that you selected the correct node canister and verify why you are reseating it. Go to "Procedure: Identifying which Storwize V7000 Gen1 enclosure or canister to service" on page 254.
3. Rotate the handle release trigger.
4. Pull out the handle to its full extension.
5. Grasp the canister to pull it out 2 or 3 inches.
6. Push the canister to return it into the slot until the handle starts to move.
7. Finish inserting the canister by closing the handle until the locking catch clicks into place.
8. Verify that the cables were not displaced.
9. Verify that the LEDs are on.

Procedure: Reseating a node canister

Use this procedure to reseat a canister that is in service state or because a service action has directed you.

About this task

Verify that you are reseating the correct node canister and that you use the correct canister handle for the node that you are reseating. Handles for the node canisters are located next to each other. The handle on the right operates the upper canister. The handle on the left operates the lower canister.

Procedure

1. Verify the clustered-system status LED on the node canister. If it is permanently on, the node is active. If the node is active, no reseating is required.
2. Verify that you have selected the correct node canister and verify why you are reseating it. Go to "Procedure: Identifying which Storwize V7000 Gen1 enclosure or canister to service" on page 254.
3. Grasp the handle between the thumb and forefinger.
4. Squeeze them together to release the handle.
5. Pull out the handle to its full extension.
6. Grasp the canister and pull it out 2 or 3 inches.
7. Push the canister back into the slot until the handle starts to move.

8. Finish inserting the canister by closing the handle until the locking catch clicks into place.
9. Verify that the cables were not displaced.
10. Verify that the LEDs are on.

Results

Procedure: Removing a Storwize V7000 Gen2 node canister

Follow this procedure to remove a node canister.

About this task

Attention: Before a node canister can be removed it must be powered off or in service state; otherwise, loss of data or loss of access to data can result.

If a node canister was recently removed from the system and then readded, ensure that the canister is online for at least 25 minutes before you remove its partner canister. This delay allows multipath drivers to fail over to the online canister when the partner canister is removed.

Procedure

1. Read the safety information referred to in “Preparing to remove and replace parts” on page 295.
2. Follow the steps in “Procedure: Powering off a Storwize V7000 Gen2 node canister” on page 283
3. Use the LEDs on the canister to confirm that it is safe to remove the canister from the enclosure, as described in “Procedure: Understanding the Storwize V7000 Gen2 system status from the LEDs” on page 256.
4. Record which data cables are plugged into the specific ports on the rear of the node canister. The cables must be inserted back into the same ports after the replacement is complete; otherwise, the system cannot function properly.
5. Disconnect the data cables that are connected to the node canister.
6. On the canister, unlatch the release lever and pull it open (see Figure 75 on page 280). The canister moves out of the slot approximately 0.6 cm (0.25 inch). Be careful that you do not inadvertently disturb or remove any cables that are connected to other components of the system.

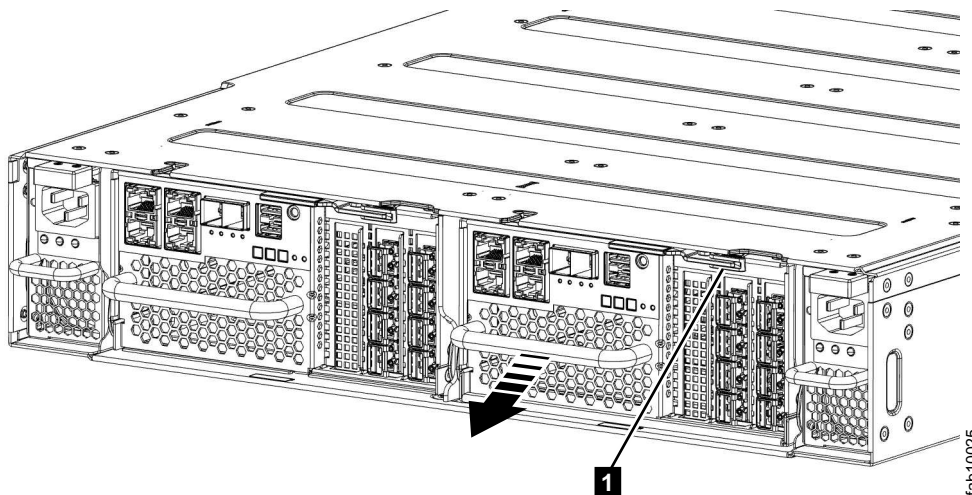


Figure 75. Removing a node canister

Note: The number scale that is etched along the top and sides of the canister indicates how much of the canister is being supported by the enclosure. When you remove the canister, ensure that you are supporting the full weight of the canister before you reach "1" on the scale.

7. As you pay attention to the number scale, slide the canister out of the slot.

Procedure: Powering off your system

You must power off your Storwize V7000 Unified system in order to service it, or to permit other maintenance actions in your data center. To turn off the Storwize V7000 Unified system, see "Turning off the system" in the Storwize V7000 Unified information center.

Procedure: Powering off your Storwize V7000 Gen1 system

You must power off your Storwize V7000 Unified system in order to service it, or to permit other maintenance actions in your data center. To turn off the Storwize V7000 Unified system, see "Turning off the system" in the Storwize V7000 Unified information center.

About this task

Procedure: Powering off your Storwize V7000 Gen2 system

You must power off your Storwize V7000 Gen2 system to service it, or to allow for other maintenance actions in your data center.

Procedure

To power off your Storwize V7000 Unified system, complete the following steps:

1. Stop all host I/O to volumes on the system.
2. Shut down the system by using the management GUI. Click **Monitoring > System**. From the **Actions** menu, select **Power off**.
3. Wait for the power LEDs on all node canisters in all control enclosures to blink at 1 Hz, indicating that the shutdown operation completed, as shown in Figure 76 on page 281.

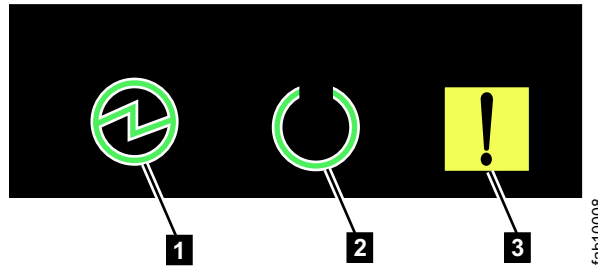


Figure 76. Power LEDs on a node canister

- 1** Power
- 2** Status
- 3** Fault

4. Disconnect the power cords from both power supplies in each control enclosure.
5. Disconnect the power cords from both power supplies in each expansion enclosure.

Procedure: Powering on the Storwize V7000 Gen2 system

After installing all hardware components, you must power on the Storwize V7000 Gen2 system and check its status.

About this task

Attention: Do not power on the system with any open bays or slots.

- Every unused drive bay must be occupied by a filler panel.
- Filler plates must be installed in all empty host interface adapter slots.

Open bays or slots disrupt the internal air flow, causing the drives to receive insufficient cooling.

Procedure

To power on the system, complete the following steps.

1. Power on the control enclosure by connecting both power supply units of the enclosure to their power sources, using the supplied power cables. If the power sources have circuit breakers or switches, ensure that they are turned on. The enclosure does not have power switches.

Note: Each enclosure has two power supply units. To provide power failure redundancy, connect the two power cords to separate power circuits.

2. Check the LEDs on each node canister in the control enclosure, as displayed in Figure 77.

Figure 77. Node canister LEDs

- 1** Power
- 2** Status
- 3** Fault

The canister is ready with no critical errors when **Power** is illuminated, **Status** is blinking, and **Fault** is off. If a canister is *not* ready, refer to the “Procedure:

Understanding the system status using the LEDs” topic in the troubleshooting section of the Storwize V7000 Unified information center.

Procedure: Powering off a Storwize V7000 Gen2 control enclosure

To service a Storwize V7000 Gen2 control enclosure, you must safely power off both of the node canisters in the enclosure.

Before you begin

Host connectivity for hosts that are connected to the control enclosure is lost when the control enclosure is shut down. You must quiesce I/O activity of these hosts before doing this procedure.

About this task

If your system has a single control enclosure, complete the steps in “Procedure: Powering off your system” on page 280 instead of following this procedure. Doing so provides a more coordinated shutdown of the whole system.

Attention:

- If your system is powered on and doing I/O operations, you must power off the control enclosures correctly to ensure that no data is lost. If possible, always use the fix procedures that are presented by the management GUI to manage and maintain your system. The fix procedures ensure that node canisters are powered off safely.
- If the system includes two control enclosures, some volumes might become inaccessible when a control enclosure shuts down. Refer to “Procedure: Understanding Storwize V7000 Gen2 volume dependencies” on page 287 to determine whether it is appropriate to continue this procedure.
- If you need to power off the node canister that is operating as the configuration node, power down the other node canister first and then power down the configuration node second when following this procedure. This prevents two failovers from happening, reducing delays in powering off the control enclosure.

Procedure

To power off a control enclosure, complete the following steps:

1. Use the management GUI to determine which two node canisters are in the control enclosure that is to be powered off. Note whether one of the two nodes is the configuration node so that you can power it off second.
2. Go to the service assistant for the first node to be powered off.
3. On the home page, select the node canister to be powered off.
4. Use the **Power off** action to power off the canister.
5. Wait for the node to appear offline.
6. Repeat steps 2 through 5 on the second node canister in the enclosure that is to be powered off.
7. If another control enclosure in the system is online, the management GUI can be used to confirm that the node status of both nodes is offline. The status LEDs on both canisters indicate whether the node is powered off, as described in “Procedure: Understanding the Storwize V7000 Gen2 system status from the LEDs” on page 256.

8. Turn off the power to the enclosure and disconnect both power cables from the enclosure.

What to do next

After powering off a control enclosure using this procedure, you must reconnect the power cables and turn the power on. The node canisters will start.

Procedure: Powering off a Storwize V7000 Gen2 node canister

You can safely power off a node canister to service the node canister.

About this task

Attention: After powering off a node canister using this procedure, a physical reseal of the canister will be required to power it back on. The reseal procedure requires physical access to the enclosure and is described in “Procedure: Reseating a Storwize V7000 Gen2 node canister” on page 278.

While a node canister is powered off, some volumes can become inaccessible. Refer to “Procedure: Understanding Storwize V7000 Gen2 volume dependencies” on page 287 to determine whether it is appropriate to continue this procedure.

If your system is powered on and doing I/O operations, it is important that the system is powered off correctly to ensure that no data is lost. If possible, always use the fix procedures that are presented by the management GUI to manage and maintain your system. The fix procedures ensure that the canister is powered off safely.

Procedure

To power off a node canister, complete the following steps:

1. Go to the service assistant for the node with the canister to shut down.
2. On the home page, select the node canister to shut down.
3. If you intend to do maintenance of the node canister, click **Identify** to light the Identify LED on the canister. Confirm that you know the location of the node canister.
4. Use the **Power off** action to power off the canister.
5. After the node is powered off, the service assistant shows that the node status is offline. The status LEDs on the canister indicate that the node is powered off.

Procedure: Collecting information for support

IBM support might ask you to collect trace files and dump files from your system to help them resolve a problem. Typically, you perform this task from the Storwize V7000 Unified management GUI. You can also collect information from the Storwize V7000 control enclosure itself.

About this task

The control enclosure management GUI and the service assistant have features to assist you in collecting the required information. The management GUI collects information from all the components in the system. The service assistant collects information from a single node canister. When the information that is collected is packaged together in a single file, the file is called a *snap*.

Special tools that are only available to the support teams are required to interpret the contents of the support package. The files are not designed for customer use.

Procedure

Always follow the instructions that are given by the support team to determine whether to collect the package by using the management GUI or the service assistant. Instruction is also given for which package content option is required.

- If you are collecting the package by using the management GUI, select **Settings > Support > Download Logs**. Click **Download Support Package**. Follow the instructions to download either the full logs or the block-storage logs.
- If you are collecting the package by using the service assistant, ensure that the node that you want to collect logs from is the current node. Select the **Collect Logs** option from the navigation. You can collect a support package or copy an individual file from the node canister. Follow the instructions to collect the information.

Procedure: Rescuing node canister software from another node (node rescue)

Use this procedure to complete a node rescue. This procedure differs, depending on the generation of your control enclosure model.

About this task

Storwize V7000 Unified *Gen1* refers to the enclosure models in the following table:

Table 101. Storwize V7000 Unified Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 Unified control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 Unified control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 Unified control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 Unified control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 Unified expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Storwize V7000 Unified *Gen2* refers to the newer generation of enclosures in the following table:

Table 102. Storwize V7000 Unified Gen2 model numbers

Machine type/model	Description
2076-524	Storwize V7000 Unified control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-12F	Storwize V7000 Unified expansion enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Procedure: Rescuing Storwize V7000 Gen2 node canister software from another node (node rescue)

Use this procedure to rescue a node.

About this task

A failure indicates that the node software is damaged and must be reinstalled.

Procedure

1. Ensure that the node you want to reinstall the code on is the current node. Go to “Accessing the service assistant” on page 204.
2. Select **Reinstall Machine Code** from the navigation.
3. Select **Rescue from another node**.

Procedure: Rescuing node canister software from another node (node rescue)

You can use the service assistant to re-install software from another node to rescue a node canister.

About this task

A failure has indicated that the node software is damaged and must be reinstalled.

Procedure

1. Log on to the service assistant. See “Accessing the service assistant” on page 204 for details.
2. Ensure that the node you want to reinstall the code on is the current node.
3. Select **Rescue from another node**.

Procedure: FCoE host-linking

About this task

If you are having problems attaching to the FCoE hosts, your problem might be related to the network, the Storwize V7000 Unified system, or the host.

Procedure

1. If you are seeing error code 705 on the node, this means that the Fibre Channel I/O port is inactive. Note that FCoE uses Fibre Channel as a protocol and an Ethernet as an interconnect. If you are dealing with an FCoE enabled port, this means that either the Fibre Channel Forwarder (FCF) is not seen or the FCoE feature is not configured on the switch.
 - a. Check that the FCoE feature is enabled on the FCF.
 - b. Check the remote port (switch port) properties on the FCF.
2. If you are connecting the host through a Converged Enhanced Ethernet (CEE) switch, for network problems, you can attempt any of the following actions:
 - a. Test your connectivity between the host and CEE switch.
 - b. Ask the Ethernet network administrator to check the firewall and router settings.
3. Run **svcinfo lsfabric** and check that the host is seen as a remote port in the output. If not, then do the following tasks in order:
 - a. Verify that Storwize V7000 Unified and host get an fcid on FCF. If not, check the VLAN configuration.

- b. Verify that Storwize V7000 Unified and host port are part of a zone and that zone is currently in force.
 - c. Verify the volumes are mapped to the host and that they are online. See **lshostvdiskmap** and **lsvdisk** in the CLI configuration guide for more information.
4. If you still have FCoE problems, you can attempt the following action:
 - a. Verify that the host adapter is in good state. You can unload and load the device driver and see the operating system utilities to verify that the device driver is installed, loaded, and operating correctly.

Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister

To remove or replace the lid of a Storwize V7000 Gen2 node canister, use this procedure.

About this task

It might be necessary to service a node canister or to replace a part that is either a customer replaceable unit (CRU) or a field replaceable unit (FRU) that is contained within the canister. To remove the lid of a node canister, follow these steps.

Attention: The lid of a node canister can be removed only after the canister was removed from its enclosure. Unless you are otherwise instructed, follow the procedure for removing a node canister to remove a node canister from its enclosure.

To remove a canister lid:

1. Place the node canister upside down on a work surface, with the release levers facing toward you.
2. Open the cover of the canister by depressing the recessed, blue touch points on the lid and sliding the lid away from you, as shown in Figure 78 on page 287.

To replace a canister lid, slide the canister lid onto the canister until the catch clicks and the lid edges are flush with the canister.

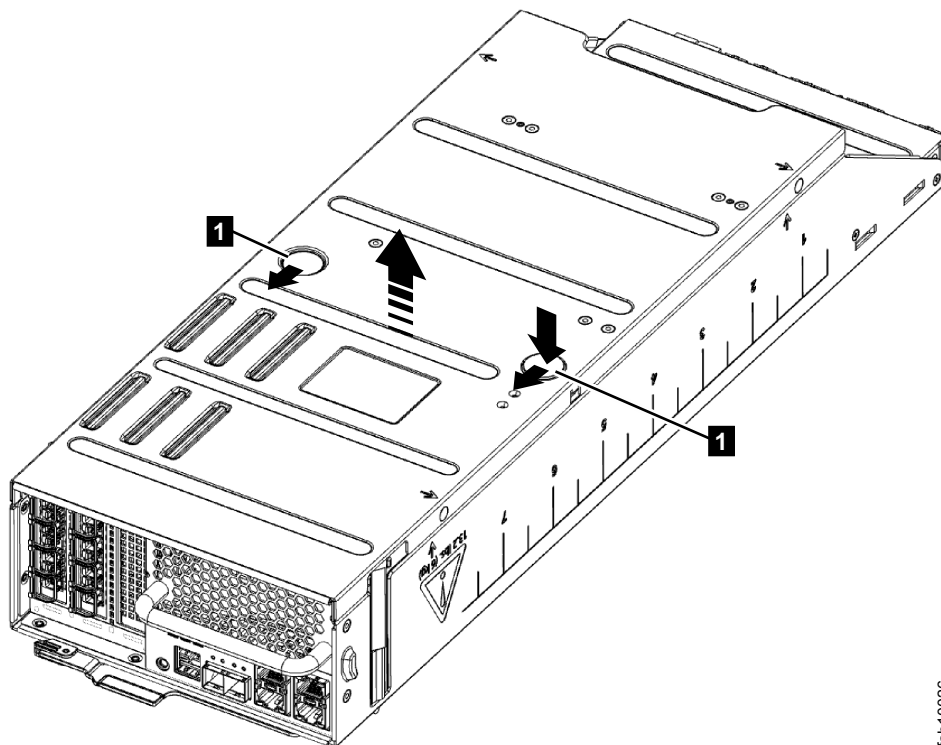


Figure 78. Replacing the canister cover

Procedure: Understanding Storwize V7000 Gen2 volume dependencies

If one component in a redundant pair is offline or powered off, host access to volumes depends on a Storwize V7000 Gen2 enclosure or canister in the system.

- If a control enclosure only has one node canister online, access to a volume depends on the online node canister if the volume is stored partially or wholly on an array that uses drives in the control enclosure or its expansion enclosures.
- If one expansion canister in an expansion enclosure is powered off, any expansion canisters further down that side of the chain become isolated from the control canister on that side of the chain. In this case, host access to volumes depends on the online canister if the volume uses drives in an isolated enclosure or the enclosure with the offline canister.
- If an entire expansion enclosure is powered off, both the left and the right side of the SAS chain are broken. In this case, host access to some volumes can be considered to depend on the entire expansion enclosure.

The impact that a service procedure might have on host access to data can be understood by using the management GUI.

1. Log on to the management GUI. Go to **Monitoring > System**.
2. From the dynamic graphic, right click the canister and select **Show Dependent Volumes** to see which volumes would be inaccessible if the component was taken offline or powered off.

If during a maintenance procedure, the **Show Dependent Volumes** action indicates that there are dependent volumes, you might choose to stop the procedure to

investigate whether it is possible to reinstate the redundancy in the system so that a procedure can be carried out without loss of access to data. An example would be to do procedures to ensure that both canisters in the enclosure are online before doing another procedure that powers off the only online canister in the enclosure.

Storwize V7000 replaceable units

Each Storwize V7000 model consists of several replaceable units. Generic replaceable units are cables, SFP transceivers, canisters, power supply units, battery assemblies, and enclosure chassis. The parts list varies, depending on the generation of your control enclosure model.

Storwize V7000 Unified *Gen1* refers to the enclosure models in the following table:

Table 103. Storwize V7000 Unified Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 Unified control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 Unified control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 Unified control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 Unified control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 Unified expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Storwize V7000 Unified *Gen2* refers to the newer generation of enclosures in the following table:

Table 104. Storwize V7000 Unified Gen2 model numbers

Machine type/model	Description
2076-524	Storwize V7000 Unified control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-12F	Storwize V7000 Unified expansion enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Storwize V7000 2076-524 Gen2 replaceable units

You might have to replace a Storwize V7000 2076-524 part. Some replaceable parts are customer-replaceable units (CRUs) and some are field-replaceable units (FRUs), which are replaced by IBM trained service technicians.

Table 105. Control enclosure replaceable units

Part number	Part name	CRU or FRU	Notes
31P1854	Control enclosure midplane assembly	FRU	Excludes drives, drive blanks, canisters, fan modules, bezel covers, PSUs.
31P1845	Node canister	CRU	Includes 2x 16 GB DIMMs, system drive, CMOS battery. Excludes Node battery, interface adapters, compression adapter/passthrough.
31P1849	Control enclosure power supply unit	CRU	
31P1847	Fan module	CRU	
45W8680	Drive blank, 2.5-inch form factor	CRU	
31P1851	Control enclosure left bezel	CRU	No MTM/Serial number label on the FRU.
00Y2512	2.5-inch enclosure right bezel	CRU	
31P1856	Control enclosure rail kit	CRU	
31P1807	Node battery	CRU	
64P8453	Node canister memory (16 GB DIMM)	CRU	
33F8354	CMOS coin battery	CRU	For real-time clock
31P1861	Compression pass-through adapter	CRU	
31P1863	Compression accelerator	CRU	
64P8473	4-port 8 Gbps Fibre Channel host interface adapter	CRU	No SFPs
00AR316	4-port 10 Gbps Ethernet host interface adapter	CRU	No SFPs
00WY984	4-port 16Gbps Fibre Channel host interface adapter	CRU	No SFPs
00RY007	2-port 16 Gbps Fibre Channel host interface adapter	CRU	No SFPs.
85Y6278	8 Gbps SW SFP	CRU	For 8 Gbps Fibre Channel host interface adapter

Table 105. Control enclosure replaceable units (continued)

Part number	Part name	CRU or FRU	Notes
00AR096	8 Gbps LW SFP	CRU	For 8 Gbps Fibre Channel host interface adapter
00RY190	16 Gbps SW SFP	CRU	For 16 Gbps Fibre Channel host interface adapter
00RY191	16 Gbps LW SFP	CRU	For 16 Gbps Fibre Channel host interface adapter
31P1630	10Gbps SFP	CRU	For 10Gbps Ethernet host interface adapter

Table 106. Expansion enclosure replaceable units

Part number	Part name	CRU or FRU	Notes
45W8680	Drive blank, 2.5-inch form factor	CRU	For models 2076-524, 2076-24F only.
64P8446	Expansion enclosure midplane assembly, 12-slot, 3.5-inch	FRU	For model 2076-12F only. Excludes drives; drive blanks; canisters; bezel covers; PSUs.
64P8447	Expansion enclosure midplane assembly, 24-slot, 2.5-inch	FRU	For model 2076-24F only. Excludes drives; drive blanks; canisters; bezel covers; PSUs.
64P8448	Expansion Canister	CRU	
98Y2218	Expansion enclosure power supply unit	CRU	
45W8680	Drive blank, 2.5-inch form factor	CRU	For models 2076-524, 2076-24F only.
42R7992	Drive blank, 3.5-inch form factor	CRU	For model 2076-12F only.
64P8450	Expansion enclosure left bezel	CRU	No MTM/Serial number label on the FRU.
00Y2512	Enclosure right bezel, 2.5-inch form factor	CRU	
00Y2436	Enclosure right bezel, 3.5-inch form factor	CRU	
64P8449	Expansion enclosure rail kit	CRU	

Table 107. Drive replaceable units

Part number	Part name	CRU or FRU	Notes
2.5-inch form factor			
00AR323	SFF HDD - 600 GB 15 K RPM	CRU	

Table 107. Drive replaceable units (continued)

Part number	Part name	CRU or FRU	Notes
00AR324	SFF HDD - 300 GB 15 K RPM	CRU	
00AR325	SFF HDD - 600 GB 10K RPM	CRU	
00AR326	SFF HDD - 900 GB 10K RPM	CRU	
00AR327	SFF HDD - 1.2 TB 10K RPM	CRU	
00RX908	SFF HDD - 1.8 TB 10K RPM 12 Gbps	CRU	Requires system software version 7.4 or later.
00AR328	SFF HDD - 1 TB NL SAS 7.2 K RPM	CRU	
00AR329	SFF 200 GB SSD	CRU	
00AR330	SFF 400 GB SSD	CRU	
00AR331	SFF 800 GB SSD	CRU	
3.5-inch form factor			
00AR320	LFF HDD - 2 TB NL SAS	CRU	
00AR321	LFF HDD - 3 TB NL SAS	CRU	
00AR322	LFF HDD - 4 TB NL SAS	CRU	
00RX911	LFF HDD - 6 TB NL 12 Gbps SAS	CRU	Requires system software version 7.4 or later.

Table 108. Cable replaceable units

Part number	Part name	CRU or FRU	Notes
Optical			
39M5699	1 m FC cable	CRU	
39M5700	5 m FC cable	CRU	
39M5701	25 m FC cable	CRU	
41V2120	10 m OM3 FC cable	CRU	
SAS			
00AR272	0.6 m 12 Gbps SAS Cable (mini SAS HD to mini SAS HD)	CRU	For connecting expansion enclosures.
00AR311	1.5 m 12 Gbps SAS Cable (mini SAS HD to mini SAS HD)	CRU	For connecting expansion enclosures.
00AR317	3.0 m 12 Gbps SAS Cable (mini SAS HD to mini SAS HD)	CRU	For connecting expansion enclosures.

Table 108. Cable replaceable units (continued)

Part number	Part name	CRU or FRU	Notes
00AR439	6.0 m 12 Gbps SAS Cable (mini SAS HD to mini SAS HD)	CRU	For connecting expansion enclosures.
Power			
39M5068	Argentina 2.8 m	CRU	
39M5199	Japan 2.8 m	CRU	
39M5123	Europe 2.8 m	CRU	
39M5165	Italy 2.8 m	CRU	
39M5102	Aus/NZ 2.8 m	CRU	
39M5130	Denmark 2.8 m	CRU	
39M5144	S. Africa 2.8 m	CRU	
39M5151	UK 2.8 m	CRU	
39M5158	Switzerland 2.8 m	CRU	
39M5172	Israel 2.8 m	CRU	
39M5206	China 2.8 m	CRU	
39M5219	Korea 2.8 m	CRU	
39M5226	India 2.8 m	CRU	
39M5240	Brazil 2.8 m	CRU	
39M5247	Taiwan 2.8 m	CRU	
39M5081	US/Canada 2.8 m	CRU	
39M5377	Power jumper cord - 2.8 m	CRU	

Storwize V7000 2076-1xx and 2076-3xx Gen1 replaceable units

The Storwize V7000 system consists of several replaceable units. Generic replaceable units are cables, SFP transceivers, canisters, power supply units, battery assemblies, and enclosure chassis.

Table 109 provides a brief description of each replaceable unit.

Table 109. Replaceable units

Part	Part number	Applicable models	FRU or customer replaced
2U24 enclosure chassis (empty chassis)	85Y5897	124, 224, 324	FRU
2U12 enclosure chassis (empty chassis)	85Y5896	112, 212, 312	FRU
Type 100 node canister	85Y5899	112, 124	Customer replaced
Type 300 node canister with 10 Gbps Ethernet ports	85Y6116	312, 324	Customer replaced
Expansion canister	85Y5850	212, 224	Customer replaced

Table 109. Replaceable units (continued)

Part	Part number	Applicable models	FRU or customer replaced
764 W power supply unit	85Y5847	112, 124, 312, 324	Customer replaced
580 W power supply unit	85Y5846	212, 224	Customer replaced
Battery backup unit	85Y5898	112, 124, 312, 324	Customer replaced
1 m SAS cable	44V4041	212, 224	Customer replaced
3 m SAS cable	44V4163	212, 224	Customer replaced
6 m SAS cable	44V4164	212, 224	Customer replaced
1 m Fibre Channel cable	39M5699	112, 124, 312, 324	Customer replaced
5 m Fibre Channel cable	39M5700	112, 124, 312, 324	Customer replaced
25 m Fibre Channel cable	39M5701	112, 124, 312, 324	Customer replaced
1.8 m power cord (Chicago)	39M5080	All	Customer replaced
2.8 m power cord (EMEA)	39M5151	All	Customer replaced
2.8 m power cord (Australia)	39M5102	All	Customer replaced
2.8 m power cord (Africa)	39M5123	All	Customer replaced
2.8 m power cord (Denmark)	39M5130	All	Customer replaced
2.8 m power cord (South Africa)	39M5144	All	Customer replaced
2.8 m power cord (Switzerland)	39M5158	All	Customer replaced
2.8 m power cord (Chile)	39M5165	All	Customer replaced
2.8 m power cord (Israel)	39M5172	All	Customer replaced
2.8 m power cord (Group 1 including the United States)	39M5081	All	Customer replaced
2.8 m power cord (Argentina)	39M5068	All	Customer replaced
2.8 m power cord (China)	39M5206	All	Customer replaced
2.8 m power cord (Taiwan)	39M5247	All	Customer replaced
2.8 m power cord (Brazil)	39M5233	All	Customer replaced

Table 109. Replaceable units (continued)

Part	Part number	Applicable models	FRU or customer replaced
2.0 m jumper cable	39M5376	All	Customer replaced
2.8 m power cord (India)	39M5226	All	Customer replaced
4.3 m power cord (Japan)	39M5200	All	Customer replaced
2.8 m power cord (Korea)	39M5219	All	Customer replaced
2.5" flash drive, 300 GB, in carrier assembly	85Y5861	124, 224, 324	Customer replaced
2.5" 10 K, 300 GB, in carrier assembly	85Y5862	124, 224, 324	Customer replaced
2.5" 10 K, 450 GB, in carrier assembly	85Y5863	124, 224, 324	Customer replaced
2.5" 10 K, 600 GB drive, in carrier assembly	85Y5864	124, 224, 324	Customer replaced
2.5" 15 K, 146 GB drive, in carrier assembly	85Y6088	124, 224, 324	Customer replaced
2.5" 15 K, 300 GB drive, in carrier assembly	85Y6185	124, 224, 324	Customer replaced
2.5" 10 K, 900 GB drive, in carrier assembly	00L4680	124, 224, 324	Customer replaced
2.5" 10 K, 1.2 TB SAS drive, in carrier assembly	85Y6156	124, 224, 324	Customer replaced
2.5" 10 K, 300 GB drive, in carrier assembly	85Y6256	124, 224, 324	Customer replaced
2.5" 10 K, 600 GB drive, in carrier assembly	85Y6268	124, 224, 324	Customer replaced
2.5" 10 K, 900 GB drive, in carrier assembly	85Y6274	124, 224, 324	Customer replaced
2.5" 10 K 1.8 TB drive, in carrier assembly	00RX915	124, 224, 324	Customer replaceable. Requires system software version 7.4 or later.
2.5" 7.2 K, Nearline SAS, 1 TB drive, in carrier assembly	85Y6186	124, 224, 324	Customer replaced
2.5" flash drive, 200 GB drive, in carrier assembly	85Y6188	124, 224, 324	Customer replaced
2.5" flash drive, 400 GB drive, in carrier assembly	85Y6189	124, 224, 324	Customer replaced
2.5" flash drive, 800 GB drive, in carrier assembly	00AR252	124, 224, 324	Customer replaced
3.5" 7.2 K Nearline SAS - 2 TB in carrier assembly	85Y5869	112, 212, 312	Customer replaced
3.5" 7.2 K Nearline SAS - 3 TB in carrier assembly	85Y6187	112, 212, 312	Customer replaced

Table 109. Replaceable units (continued)

Part	Part number	Applicable models	FRU or customer replaced
3.5" 7.2 K Nearline SAS - 6 TB in carrier assembly	00RX918	112, 212, 312	Customer replaceable. Requires system software version 7.4 or later.
Blank 2.5" carrier	85Y5893	124, 224, 324	Customer replaced
Blank 3.5" carrier	85Y5894	112, 212, 312	Customer replaced
Fibre Channel shortwave small form-factor pluggable (SFP)	85Y5958	112, 124, 312, 324	Customer replaced
Fibre Channel longwave small form-factor pluggable (SFP)	85Y5957	112, 124, 312, 324	Customer replaced
Ethernet small form-factor pluggable (SFP)	31P1549	312, 324	Customer replaced
Rail kit	85Y5852	All	Customer replaced
Left enclosure cap including RID tag but no black MTM label	85Y5901	All	Customer replaced
Right enclosure cap (2U12)	85Y5903	112, 212, 312	Customer replaced
Right enclosure cap (2U24)	85Y5904	124, 224, 324	Customer replaced

Replacing parts

You can remove and replace customer-replaceable units (CRUs) in control enclosures or expansion enclosures.

Attention: Even though many of these components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.

Each replaceable unit has its own removal procedure. Sometimes you can find that a step within a procedure might refer you to a different remove and replace procedure. You might want to complete the new procedure before you continue with the first procedure that you started.

Remove or replace parts only when you are directed to do so.

Preparing to remove and replace parts

Before you remove and replace parts, you must be aware of all safety issues.

Before you begin

First, read the safety precautions in the *IBM Systems Safety Notices*. These guidelines help you safely work with the Storwize V7000 Unified.

Replacing a node canister

Remove and replace a node canister.

Replacing a Storwize V7000 Gen2 node canister

To replace a faulty node canister with a new one received from CRU or FRU stock, use this procedure. When replacing a node canister, aim to maximize drive and system availability by maintaining one online node in the control enclosure with the faulty node canister. If you cannot maintain at least one node canister online in the system, then you might need to follow the system recovery procedure after replacing the faulty node canister.

Procedure

1. Follow “Procedure: Removing a Storwize V7000 Gen2 node canister” on page 279 to remove the faulty node canister.
2. Remove the lid of the faulty canister, as described in “Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister” on page 286. Do the same to the replacement canister.
3. Some components inside the faulty node canister must be transferred to the replacement canister. Transfer each of the following components as necessary:
 - The battery, as described in “Replacing the battery in a Storwize V7000 Gen2 node canister” on page 316.
 - The host interface adapters in one control enclosure, as described in “Replacing a Storwize V7000 Gen2 host interface adapter” on page 372.
 - Memory modules, as described in “Replacing a Storwize V7000 Gen2 node canister memory module (16 GB DIMM)” on page 371.
 - The compression pass-through adapter or compression accelerator, as described in the installation description of upgrading the hardware to install the compression accelerator.
4. Replace the lid of the faulty canister and the lid of the replacement canister, as described in “Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister” on page 286.
5. Open the release lever of the replacement canister.
6. Push the replacement canister into the slot of the enclosure that the faulty canister was removed from, until it stops.
7. Finish inserting the replacement canister by closing its release lever so that the orange latch engages the enclosure.
8. If the enclosure is powered and the canister is correctly installed, the canister starts automatically. Repeat from step 5, if the canister is not correctly installed.
9. Reinsert the data cables into the ports that they were originally connected.
10. If no node canisters are online, the system is not online. To recover the system in the case when no node canisters are online, see “Recover system procedure” on page 381.
11. If only the replacement node is in service state with node error 503, apply “Procedure: Rescuing Storwize V7000 Gen2 node canister software from another node (node rescue)” on page 285 to rescue the replacement node canister.
12. When the node canister is powered up, it is automatically added to the system and the system automatically ensures that the machine code version on the new canister matches that of the other node canister in the control enclosure. This is reflected in the system event log.

13. When the canister is back online, check the event log for any new events that might indicate a problem with the reassembly.

Replacing a Storwize V7000 Gen1 node canister

This topic describes how to replace a node canister.

About this task

Attention: Even though many of these components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.

Attention: Do not replace one type of node canister with another type. For example, do not replace a model 2076-112 node canister with a model 2076-312 node canister.

Be aware of the following canister LED states:

- If both the power LED and system status LED are on, do not remove a node canister unless directed to do so by a service procedure.
- If the system status is off, it is acceptable to remove a node canister. However, do not remove a node canister unless directed to do so by a service procedure.
- If the power LED is flashing or off, it is safe to remove a node canister. However, do not remove a node canister unless directed to do so by a service procedure.

Attention: Even if a node canister is powered off, it is still possible to lose data. Do not remove a node canister unless directed to do so by a service procedure.

To replace the node canister, perform the following steps:

Procedure

1. Read the safety information to which “Preparing to remove and replace parts” on page 295 refers.
2. Confirm that you know which canister to replace. Go to “Procedure: Identifying which Storwize V7000 Gen1 enclosure or canister to service” on page 254.
3. Record which data cables are plugged into the specific ports of the node canister. The cables must be inserted back into the same ports after the replacement is complete; otherwise, the system cannot function properly.
4. Disconnect the data cables for each canister.
5. Grasp the handle between the thumb and forefinger.

Note: Ensure that you are opening the correct handle. The handle locations for the node canisters and expansion canisters are slightly different.

Handles for the node canisters are located in close proximity to each other. The handle with the finger grip on the right removes the upper canister (**1**). The handle with the finger grip on the left removes the lower canister (**2**).

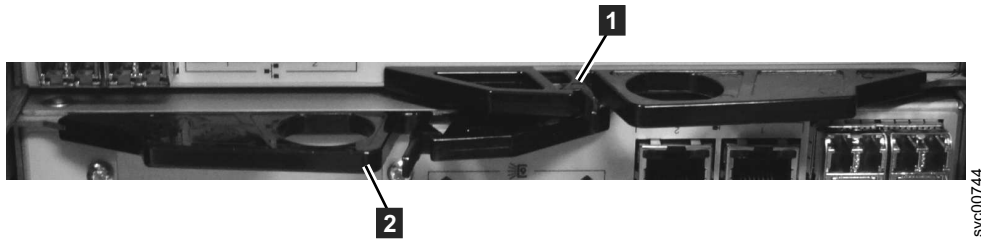


Figure 79. Rear of node canisters that shows the handles.

6. Squeeze them together to release the handle.

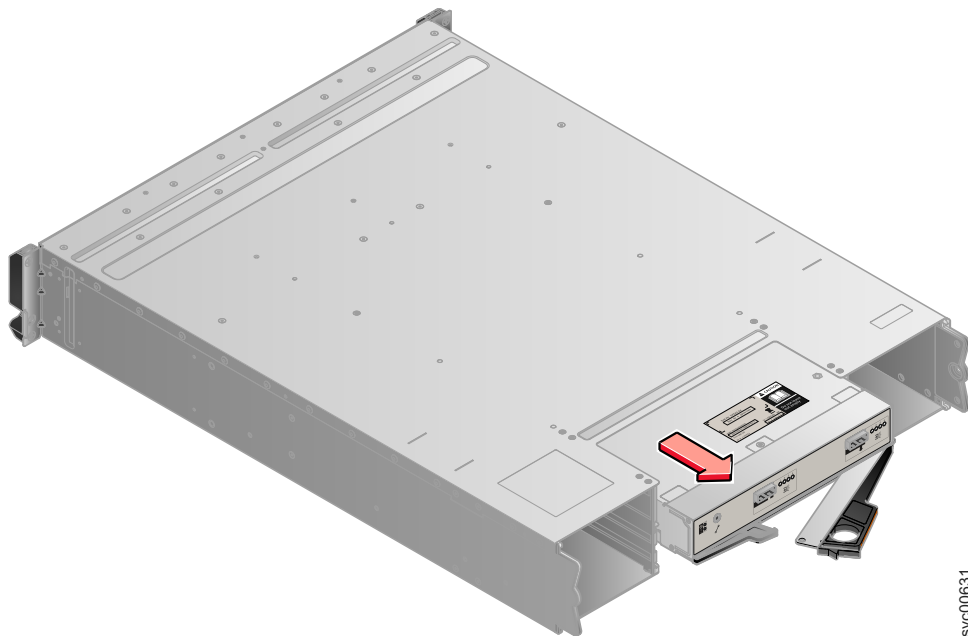


Figure 80. Removing the canister from the enclosure

7. Pull out the handle to its full extension.
8. Grasp canister and pull it out.
9. Insert the new canister into the slot with the handle pointing towards the center of the slot. Insert the unit in the same orientation as the one that you removed.
10. Push the canister back into the slot until the handle starts to move.
11. Finish inserting the canister by closing the handle until the locking catch clicks into place. Make sure that the canister is firmly and correctly seated, otherwise this can lead to problems.
If the enclosure is powered on, the canister starts automatically.
12. Reattach the data cables.

Replacing a fan module

Remove and replace a fan module.

Replacing a Storwize V7000 Gen2 fan module

Use this procedure to replace a faulty fan module with a new one received from CRU or FRU stock.

About this task

A fan module is located behind each node canister, and is accessed using the node canister slot after the node canister is removed.

Do not remove the node canister and faulty fan module before the replacement fan is on hand. The replacement procedure described must be completed within 5 minutes of the faulty fan module being removed to ensure that components do not shut down due to excessive temperatures.

When removing a node canister, aim to maximize drive and system availability by maintaining one online node in the control enclosure. If you cannot maintain at least one node canister online in the system, then you might need to follow the system recovery procedure after the node canister is replaced into the enclosure.

Procedure

1. Remove the replacement fan module from its packaging. Familiarize yourself with the part by reading through this procedure.
2. Remove the node canister that is on the same side of the enclosure as the faulty fan module. See "Procedure: Removing a Storwize V7000 Gen2 node canister" on page 279.
3. Locate the two orange locking rings of the fan module inside the left and right edges of the node canister slot. Note their position relative to the inside of the canister slot.
4. Simultaneously rotate both rings upwards through 90 degrees, releasing the fan module from the slot. Pull the locking rings to slide the faulty fan module out from the canister slot.
5. Ensure that the orange locking rings on the replacement fan module are rotated open so that they extend out from the fan module.
6. Slide the replacement fan module into the canister slot until it stops.
7. Simultaneously rotate both locking rings downwards through 90 degrees while applying gentle pressure to push the fan module into the slot. The fan module is installed correctly when the back edges of the locking rings are flush with the relief detail inside the canister slot.
8. Replace the node canister into the canister slot until it stops.
9. Finish inserting the node canister by closing its release lever so that the orange catch engages the enclosure.
10. If the enclosure is powered and the canister is correctly installed, the canister starts automatically. Remove the canister and repeat the procedure from step 5, if the canister is not correctly installed.
11. Reinsert the data cables into the ports that they were originally connected.
12. When the canister is back online, check the event log for any new events that might indicate a problem with the reassembly.

Replacing an expansion canister

Remove and replace an expansion canister.

Replacing a Storwize V7000 Gen2 expansion canister

To replace a faulty expansion canister with a new one received from CRU / FRU stock, use this procedure.

About this task

Attention: Even though many of these components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.

Do not remove an expansion canister unless directed to do so by a service procedure.

To replace an expansion canister, do the following steps:

Procedure

1. Read the safety information in “Preparing to remove and replace parts” on page 295.
2. Refer to “Procedure: Understanding Storwize V7000 Gen2 volume dependencies” on page 287 to determine whether to do this procedure.
3. Carefully identify the expansion canister that you are replacing. If possible, go to **Monitoring > System** in the management GUI. Select the expansion enclosure that you are replacing and select **Actions > Identify** to set the canister fault LED blinking.
4. Record which SAS cables are plugged into the specific ports of the expansion canister. The cables must be inserted back into the same ports after the replacement is complete; otherwise, the system cannot function properly.
5. Disconnect the SAS cables from the canister.
6. Open the two release levers as shown in Figure 81 on page 301. The canister moves out of the slot approximately 0.6 cm (0.25 inch).
7. Slide the canister out of the slot.
8. Open the release levers of the replacement canister.
9. Push the replacement canister into the slot until it stops.
10. Finish inserting the canister by closing both release levers so that both orange latches click into place.
11. The canister is correctly installed when the rear face of the canister is flush with the rear edge of the enclosure.
If the enclosure is powered on and the canister is correctly installed, the canister starts automatically.
12. Reattach each SAS cable into the port from which it was removed in step 5.
 - a. Ensuring the SAS cable connectors are inserted with the pull tab to the bottom of the connector, gently push the connector in until a slight click is felt or heard.
 - b. Verify that the connector is fully inserted by gently pulling on it (not on the tab).

You should not be able to remove it.

If the enclosure is powered on and the SAS connector is correctly inserted into the port, the green SAS link LED above the port lights up.

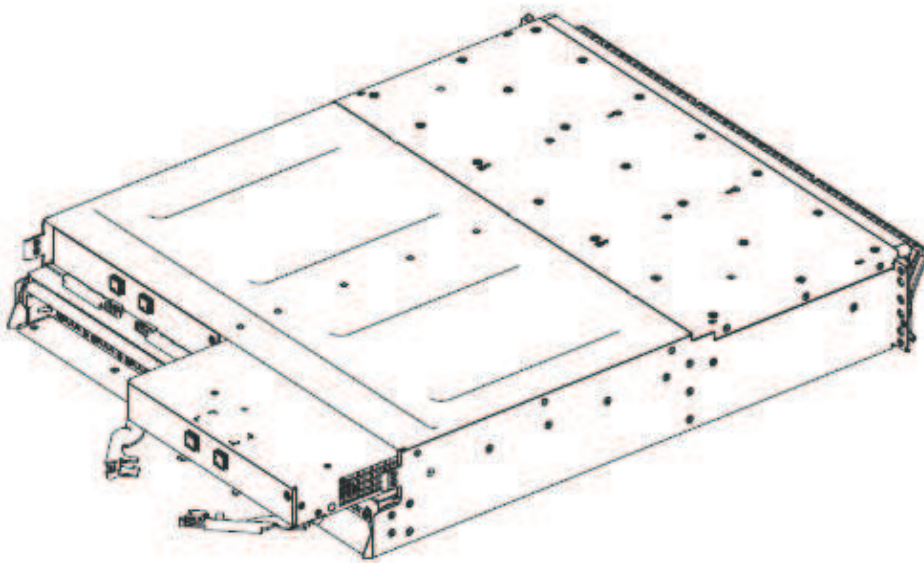


Figure 81. Removing and replacing the Storwize V7000 Gen2 expansion canister

Replacing an expansion canister

This topic describes how to replace an expansion canister.

About this task

Attention: Even though many of these components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.

Be aware of the following canister LED states:

- If the power LED is on, do not remove an expansion canister unless directed to do so by a service procedure.
- If the power LED is flashing or off, it is safe to remove an expansion canister. However, do not remove an expansion canister unless directed to do so by a service procedure.

Attention: Even if an expansion canister is powered off, it is still possible to lose data. Do not remove an expansion canister unless directed to do so by a service procedure.

To replace an expansion canister, perform the following steps:

Procedure

1. Read the safety information to which “Preparing to remove and replace parts” on page 295 refers.
2. Record which SAS cables are plugged into the specific ports of the expansion canister. The cables must be inserted back into the same ports after the replacement is complete; otherwise, the system cannot function properly.
3. Disconnect the SAS cables for each canister.
4. Grasp the handle between the thumb and forefinger.

Note: Ensure that you are opening the correct handle. The handle locations for the node canisters and expansion canisters are slightly different.

Handles for the upper and lower expansion canisters overlap each other. The handle with the finger grip on the left removes the upper canister (**1**). The handle with the finger grip on the right removes the lower canister (**2**).

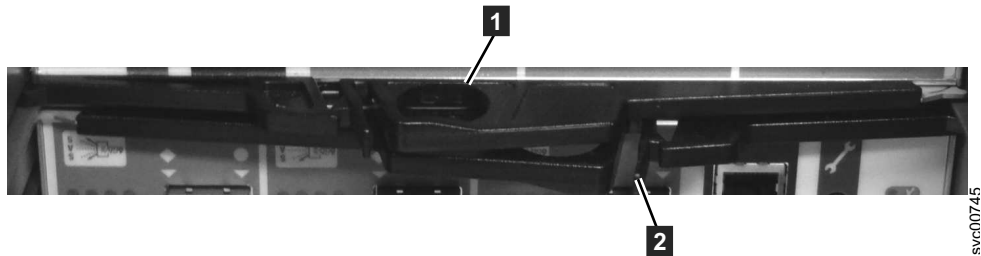


Figure 82. Rear of expansion canisters that shows the handles.

5. Squeeze them together to release the handle.

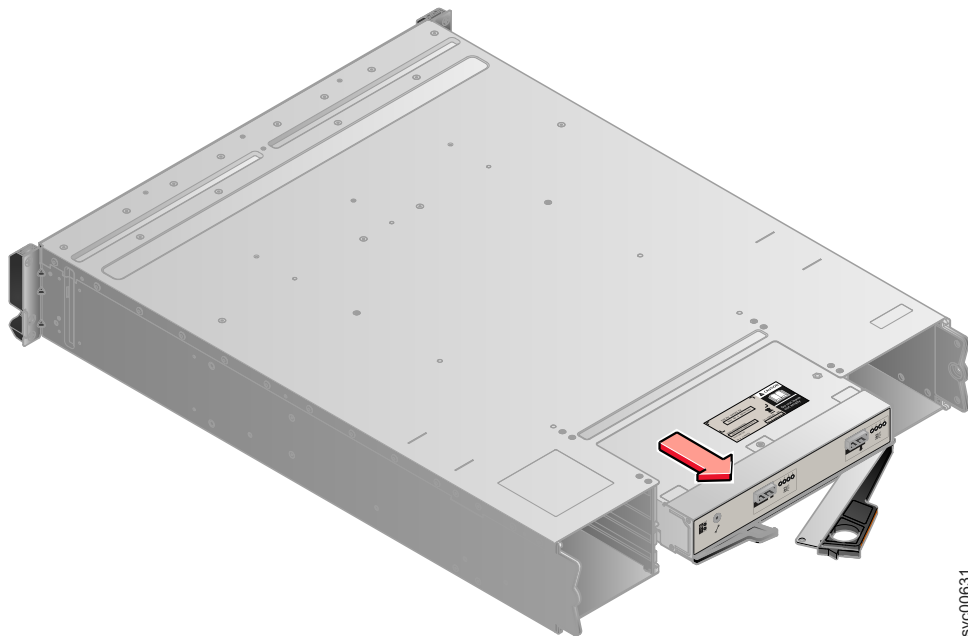


Figure 83. Removing the canister from the enclosure

6. Pull out the handle to its full extension.
7. Grasp canister and pull it out.

8. Insert the new canister into the slot with the handle pointing towards the center of the slot. Insert the unit in the same orientation as the one that you removed.
9. Push the canister back into the slot until the handle starts to move.
10. Finish inserting the canister by closing the handle until the locking catch clicks into place. Make sure that the canister is firmly and correctly seated, otherwise this can lead to problems.
11. Reattach the SAS cables.

Replacing an SFP transceiver

Remove and replace an SFP transceiver.

Replacing an SFP transceiver in a Storwize V7000 2076-524 control enclosure

When a failure occurs on an optical link, the SFP transceiver in the port providing the link might need to be replaced. To replace a faulty SFP transceiver with a new one received from CRU or FRU stock, use this procedure.

Before you begin

Even though many of these components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.

CAUTION:

Some laser products contain an embedded Class 3A or Class 3B laser diode.

Note the following information: laser radiation when open. Do not stare into the beam, do not view directly with optical instruments, and avoid direct exposure to the beam. (C030)

Attention: When replacing this part, you must follow recommended procedures for handling electrostatic discharge (ESD) sensitive devices.

Important: For correct operation, use the correct SFP transceivers with each adapter card. The topic “Storwize V7000 2076-524 Gen2 replaceable units” identifies the suitable IBM parts.

- Use only 8G bps SFP transceivers in the 8 Gbps Fibre Channel adapter cards.
- Use only 16 Gbps SFP transceivers in the 16 Gbps Fibre Channel adapter cards.
- Use only 10 Gbps SFP transceivers in the 10 Gbps Ethernet (FCoE/iSCSI) adapter card.

Procedure

Complete the following steps to remove and then replace an SFP transceiver.

1. Carefully determine the failing physical port connection.

Important: Removing the wrong SFP transceiver might result in loss of data access.

2. Remove the cable from the SFP.

Figure 84 illustrates an SFP transceiver.



Figure 84. SFP transceiver

3. Remove the faulty SFP transceiver from its aperture.
 - a. Unclip the handle of the SFP transceiver.
 - b. Pull on the handle of the SFP transceiver.
 - c. The SFP transceiver slides out of its slot.
4. Install the replacement SFP transceiver into the aperture that is vacated in step 3.
 - a. Open the lock on the replacement SFP transceiver.
 - b. Push the new SFP transceiver into the aperture until it stops.
 - c. Close the release handle.
 - d. Gently pull the SFP transceiver. If it is installed correctly, it does not move from its aperture.
5. Reconnect the optical cable.
6. Confirm that the error is now fixed. Either mark the error as fixed or restart the node, depending on the failure indication originally noted.

Replacing an SFP transceiver in a control enclosure

When a failure occurs on a single link, the SFP transceiver might need to be replaced.

Before you begin

Even though many of these components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.

CAUTION:

Some laser products contain an embedded Class 3A or Class 3B laser diode. Note the following information: laser radiation when open. Do not stare into the beam, do not view directly with optical instruments, and avoid direct exposure to the beam. (C030)

About this task

Perform the following steps to remove and then replace an SFP transceiver:

Procedure

1. Carefully determine the failing physical port connection.

Important: The Fibre Channel links in the enclosures are supported with both longwave SFP transceivers and shortwave SFP transceivers. A longwave SFP transceiver has some blue components that are visible even when the SFP transceiver is plugged in. You must replace an SFP transceiver with the same type of SFP transceiver that you are replacing. If the SFP transceiver to replace is a longwave SFP transceiver, for example, you must replace with another longwave SFP transceiver. Removing the wrong SFP transceiver might result in loss of data access.

2. Remove the optical cable by pressing the release tab and pulling the cable out. Be careful to exert pressure only on the connector and do not pull on the optical cables.
3. Remove the SFP transceiver. There are a number of different handling or locking mechanisms that are used on the SFP transceivers. Some SFP transceivers might have a plastic tag. If so, pull the tag to remove the SFP transceiver.

Important: Always check that the SFP transceiver that you replace matches the SFP transceiver that you remove.

4. Push the new SFP transceiver into the aperture and ensure that it is securely pushed home. The SFP transceiver usually locks into place without having to swing the release handle until it locks flush with the SFP transceiver. Figure 85 illustrates an SFP transceiver and its release handle.



svc00418

Figure 85. SFP transceiver

5. Reconnect the optical cable.
6. Confirm that the error is now fixed. Either mark the error as fixed or restart the node depending on the failure indication that you originally noted.

Replacing a power supply unit for a control enclosure

Remove and replace the power supply units in the control enclosure.

Replacing a Storwize V7000 Gen2 power supply unit for a control enclosure

You can replace either of the two hot-swap redundant power supplies in an enclosure. These redundant power supplies operate in parallel, one continuing to power the enclosure if the other fails.

Before you begin

Attention:

- Even though many of these components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.
- Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.
- Ensure that you are aware of the procedures for handling static-sensitive devices before you replace the power supply.

About this task

To replace the power supply, do the following steps:

Procedure

1. Read the safety information in “Preparing to remove and replace parts” on page 295.
2. Confirm that you know which power supply must be replaced. Go to “Procedure: Identifying which Storwize V7000 Gen2 enclosure or canister to service” on page 252.
3. Disconnect the power cord from the electrical outlet. Release the cable retention clip and disconnect the power cord from the power supply that you are replacing.
4. Locate the orange release tab at the top edge of the power supply unit. Press the release tab gently until it stops.
5. Using the handle, firmly pull the power supply out of the enclosure shown in Figure 86 on page 307.

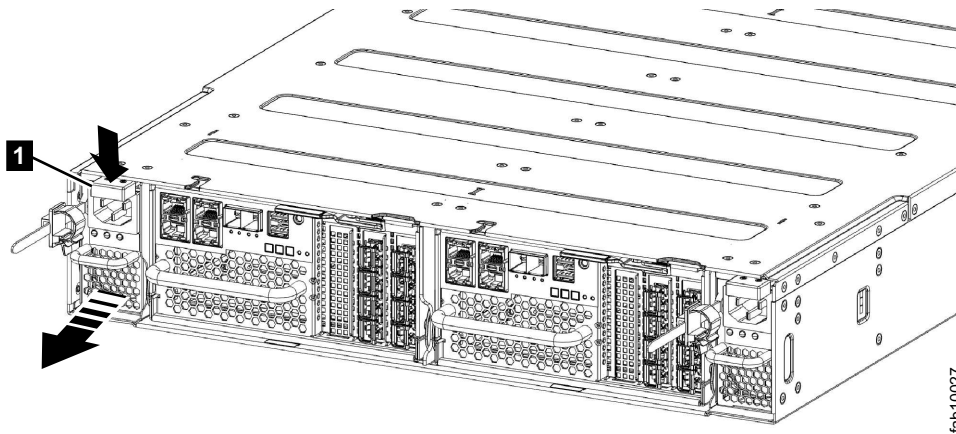


Figure 86. Removing the power supply unit (left side of enclosure)

6. Hold the new power supply so that the handle is fully extended.
7. Slide the power supply into the enclosure until it stops. Push it firmly into position until it clicks.
8. Connect the power cord to the power supply and to a properly grounded electrical outlet. Secure the cable with the cable retention clip on the rear of the power supply unit.

Note: After the power cord is connected to the electrical outlet, make sure that the ac and dc power (green) LEDs are lit and the fault (amber) LED is off.

Replacing a Storwize V7000 Gen1 power supply unit for a control enclosure

You can replace either of the two 764 watt hot-swap redundant power supplies in the control enclosure. These redundant power supplies operate in parallel, one continuing to power the canister if the other fails.

Before you begin

DANGER

When working on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard:

- If IBM supplied a power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.
- Do not open or service any power supply assembly.
- Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.
- Connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate.
- Connect any equipment that will be attached to this product to properly wired outlets.
- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.

To disconnect:

1. Turn off everything (unless instructed otherwise).
2. Remove the power cords from the outlets.
3. Remove the signal cables from the connectors.
4. Remove all cables from the devices.

To connect:

1. Turn off everything (unless instructed otherwise).
 2. Attach all cables to the devices.
 3. Attach the signal cables to the connectors.
 4. Attach the power cords to the outlets.
 5. Turn on the devices.
- Sharp edges, corners and joints might be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching.
- (D005)

Attention: Even though many of these components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.

Attention: A powered-on enclosure must not have a power supply removed for more than five minutes because the cooling does not function correctly with an empty slot. Ensure that you have read and understood all these instructions and have the replacement available, and unpacked, before you remove the existing power supply.

Even though many of these components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.

Attention: In some instances, it might not be advisable to remove a power supply unit when a system is performing I/O. For example, the charge in the backup battery might not be sufficient enough within the partner power-supply unit to continue operations without causing a loss of access to the data. Wait until the partner battery is 100% charged before replacing the power supply unit.

Ensure that you are aware of the procedures for handling static-sensitive devices before you replace the power supply.

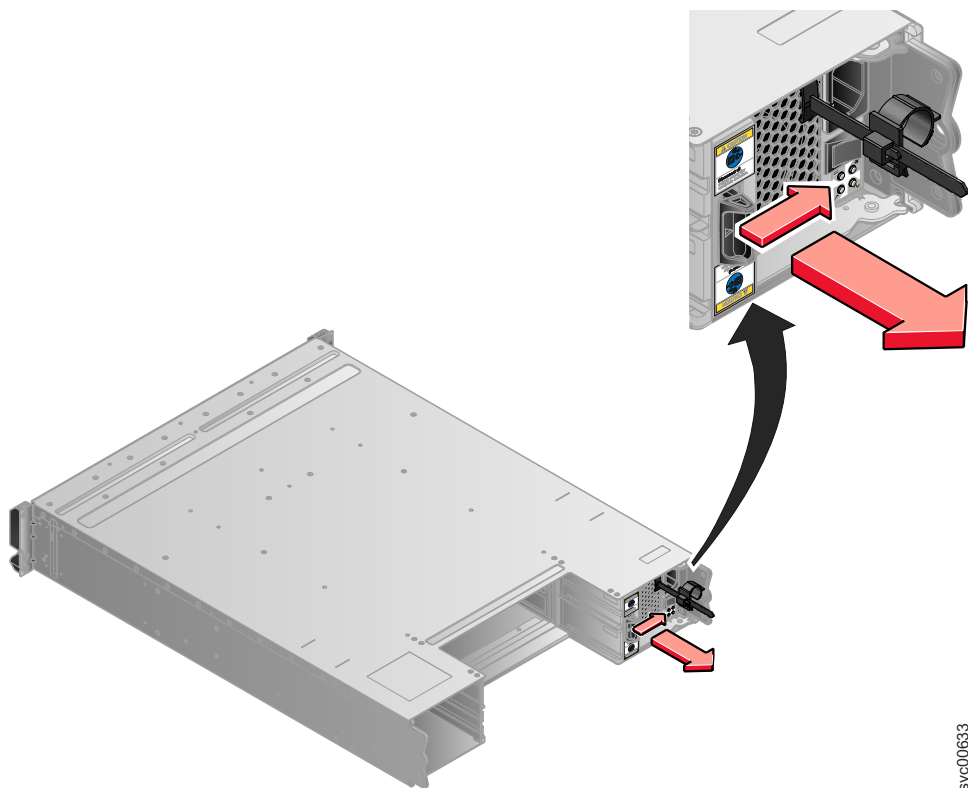
About this task

A replacement power supply unit is not shipped with a battery; therefore, transfer the battery from the existing power supply unit to the replacement unit. To transfer a battery, go to “Replacing a battery in a Storwize V7000 Gen1 power supply unit” on page 318.

To replace the power supply, perform the following steps:

Procedure

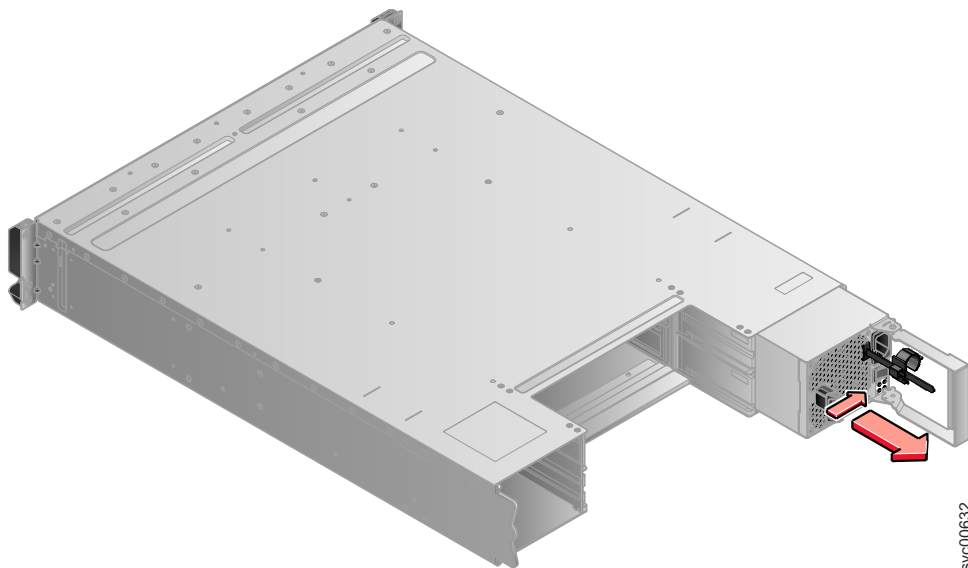
1. Read the safety information to which “Preparing to remove and replace parts” on page 295 refers.
2. Examine the Identify LED that is lit on the front of the enclosure to identify the correct enclosure.
3. Turn off the power to the power supply unit using the switch at the back.
4. Disconnect the cable retention bracket and the power cord from the power supply that you are replacing.
5. Remove the power supply unit. Record the orientation of the power supply unit. Power supply unit 1 is top side up, and power supply unit 2 is inverted.
 - a. Depress the black locking catch from the side with the colored sticker as shown in Figure 87 on page 310.



svc00633

Figure 87. Directions for lifting the handle on the power supply unit

- b. Grip the handle to pull the power supply out of the enclosure as shown in Figure 88.



svc00632

Figure 88. Using the handle to remove a power supply unit

6. Insert the replacement power supply unit into the enclosure with the handle pointing towards the center of the enclosure. Insert the unit in the same orientation as the one that you removed.

7. Push the power supply unit back into the enclosure until the handle starts to move.
8. Finish inserting the power supply unit into the enclosure by closing the handle until the locking catch clicks into place.
9. Reattach the power cable and cable retention bracket.
10. Turn on the power switch to the power supply unit.

What to do next

If required, return the power supply. Follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Replacing a power supply unit for an expansion enclosure

Remove and replace the hot-swap redundant power supplies in the expansion enclosure.

Replacing a power supply unit for a Storwize V7000 Gen2 expansion enclosure

You can replace either of the two hot-swap redundant power supplies in an enclosure. These redundant power supplies operate in parallel, one continuing to power the canister if the other fails.

Before you begin

Attention:

- Even though many of these components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.
- Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.
- Ensure that you are aware of the procedures for handling static-sensitive devices before you replace the power supply.

About this task

To replace the power supply, do the following steps:

Procedure

1. Read the safety information in “Preparing to remove and replace parts” on page 295.
2. Confirm that you know which power supply must be replaced. Go to “Procedure: Identifying which Storwize V7000 Gen2 enclosure or canister to service” on page 252.
3. Disconnect the power cord from the electrical outlet. Release the cable retention clip and disconnect the power cord from the power supply that you are replacing.
4. On the left side of the power supply, press the orange release tab to the right just enough to release the handle (no more than 6 mm [0.25 in.]) as you rotate the handle downward.

5. Using the handle, gently slide the power supply out of the enclosure, as shown in Figure 89.

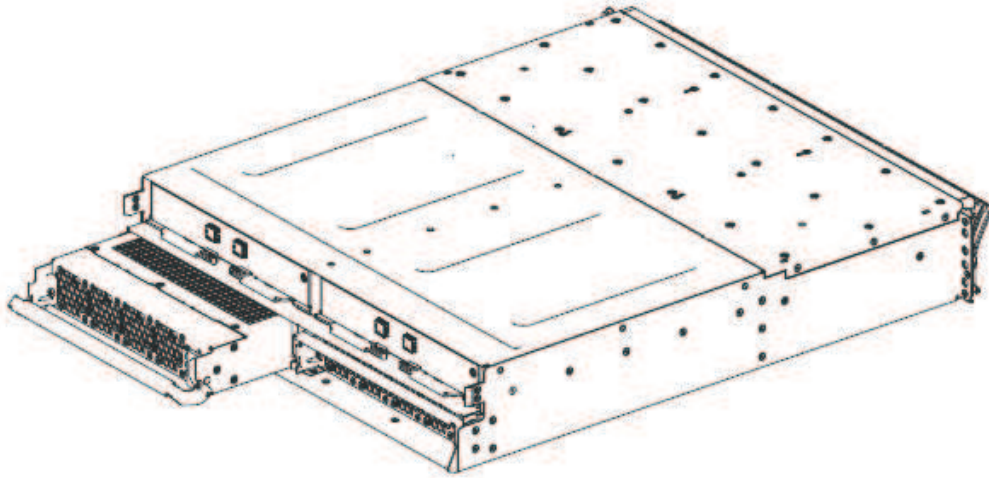


Figure 89. Removing the power supply unit from the left side of the expansion enclosure

6. Hold the new power supply so that the handle is fully extended.
7. Slide the power supply into the enclosure until it stops. Rotate the handle upward into the closed position until it clicks.
8. Hold the new power supply so that the handle is fully extended.
9. Connect the power cord to the power supply and to a properly grounded electrical outlet.

Note: After the power cord is connected to the electrical outlet, make sure that the ac and dc power (green) LEDs are lit and the fault (amber) LED is off.

Replacing a power supply unit for an expansion enclosure

You can replace either of the two 580 watt hot-swap redundant power supplies in the expansion enclosure. These redundant power supplies operate in parallel, one continuing to power the canister if the other fails.

Before you begin

DANGER

When working on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard:

- If IBM supplied a power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.
- Do not open or service any power supply assembly.
- Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.
- Connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate.
- Connect any equipment that will be attached to this product to properly wired outlets.
- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.

To disconnect:

1. Turn off everything (unless instructed otherwise).
2. Remove the power cords from the outlets.
3. Remove the signal cables from the connectors.
4. Remove all cables from the devices.

To connect:

1. Turn off everything (unless instructed otherwise).
 2. Attach all cables to the devices.
 3. Attach the signal cables to the connectors.
 4. Attach the power cords to the outlets.
 5. Turn on the devices.
- Sharp edges, corners and joints might be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching.
- (D005)

Attention: Even though many of these components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.

Attention: A powered-on enclosure must not have a power supply removed for more than five minutes because the cooling does not function correctly with an empty slot. Ensure that you have read and understood all these instructions and have the replacement available, and unpacked, before you remove the existing power supply.

Even though many of these components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.

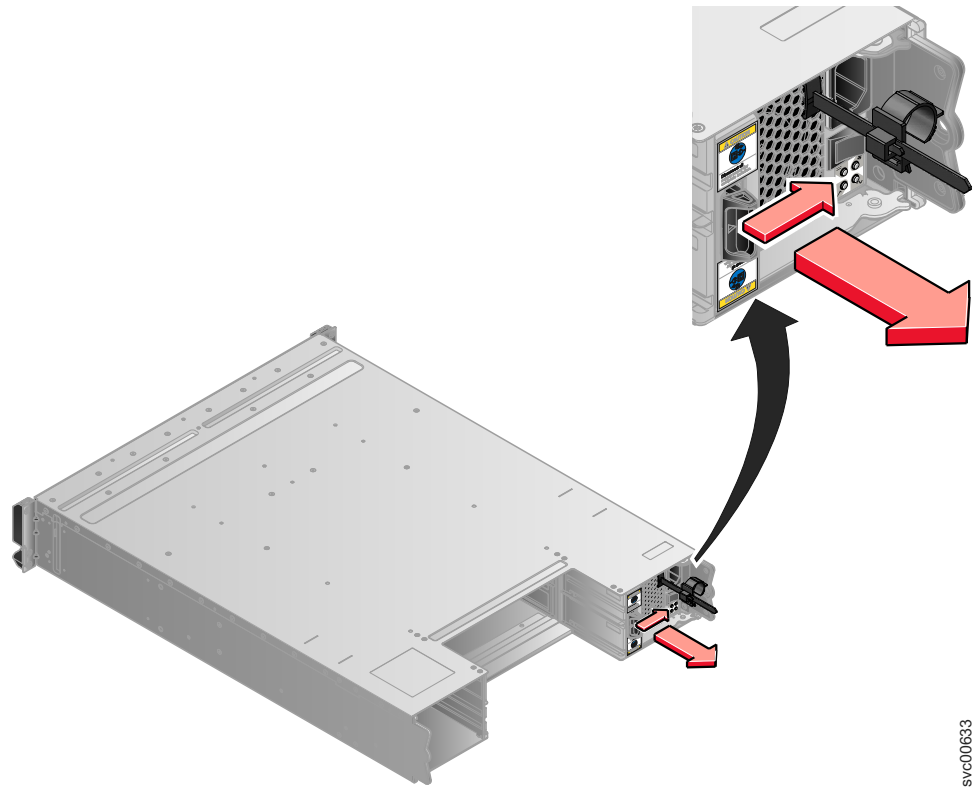
Ensure that you are aware of the procedures for handling static-sensitive devices before you replace the power supply.

About this task

To replace the power supply unit in an expansion enclosure, perform the following steps:

Procedure

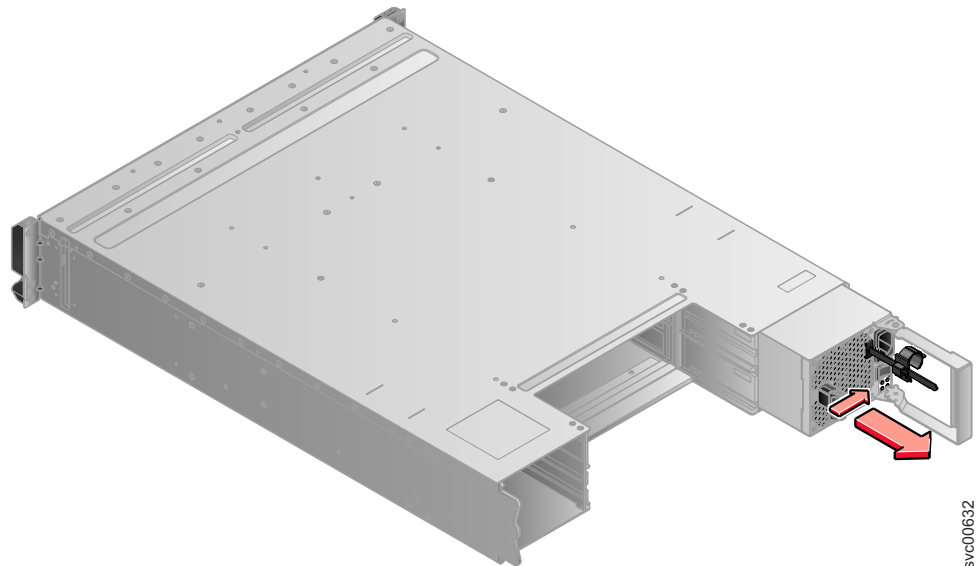
1. Read the safety information to which “Preparing to remove and replace parts” on page 295 refers.
2. Examine the Identify LED that is lit on the front of the enclosure to identify the correct enclosure.
3. Turn off the power to the power supply unit using the switch at the back of the unit.
4. Disconnect the cable retention bracket and the power cord from the power supply that you are replacing.
5. Remove the power supply unit. Record the orientation of the power supply unit. Power supply unit 1 is top side up, and power supply unit 2 is inverted.
 - a. Depress the black locking catch from the side with the colored sticker as shown in Figure 90 on page 315.



svc00633

Figure 90. Directions for lifting the handle on the power supply unit

- b. Grip the handle to pull the power supply out of the enclosure as shown in Figure 91.



svc00632

Figure 91. Using the handle to remove a power supply unit

6. Insert the replacement power supply unit into the enclosure with the handle pointing towards the center of the enclosure. Insert the unit in the same orientation as the one that you removed.

7. Push the power supply unit back into the enclosure until the handle starts to move.
8. Finish inserting the power supply unit in the enclosure by closing the handle until the locking catch clicks into place.
9. Reattach the power cable and cable retention bracket.
10. Turn on the power switch to the power supply unit.

What to do next

If required, return the power supply. Follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Replacing the battery in a node canister

Remove and replace the battery in a node canister.

Replacing the battery in a Storwize V7000 Gen2 node canister

To replace a faulty battery with a new one received from customer replaceable unit (CRU) or field replaceable unit (FRU) stock, use this procedure.

About this task

CAUTION:

The battery is a lithium ion battery. To avoid possible explosions, do not burn. Exchange only with the approved part. Recycle or discard the battery as instructed by local regulations. (C007a)

To replace a battery:

Procedure

1. Identify the node canister with the faulty battery by following the procedure “Procedure: Understanding the Storwize V7000 Gen2 system status from the LEDs” on page 256.
2. Follow the procedure “Procedure: Removing a Storwize V7000 Gen2 node canister” on page 279 to remove the node canister with the faulty battery.
3. Open the lid of the canister, as described in “Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister” on page 286.
4. Spread the two blue battery latches outward as shown in Figure 92 on page 317. Raise open both latching arms of the battery simultaneously to disconnect the battery.

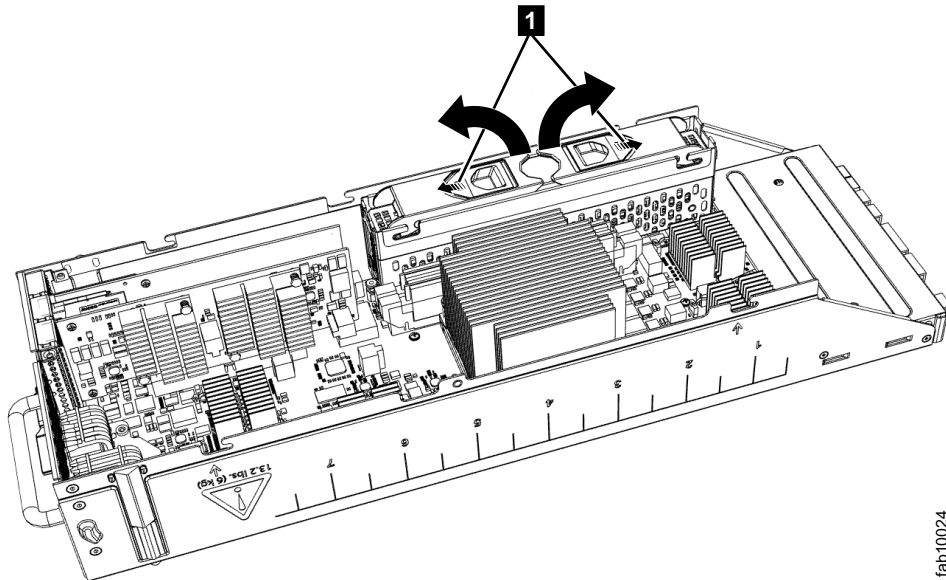


Figure 92. Opening latching arms to disconnect a Storwize V7000 Gen2 node canister battery

5. Holding the battery by its latching arms, lift it from its cradle. Place the battery in a safe place.
6. Remove the replacement battery from its package.
7. Open the latching arms of the replacement battery, then place the replacement battery into the battery cradle of the node canister so that the connectors align.
8. Apply gentle downward pressure to both battery latches so that the battery is drawn into the battery cradle. Ensure that both latches are fully engaged by spreading the two blue latches outward while you apply gentle downwards pressure.
9. Replace the canister lid, as described in “Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister” on page 286.
10. Reinstall the canister into the enclosure from which it was removed in step 2 on page 316.

Notes:

- If the replacement battery is sufficiently charged, the node canister starts.
 - If the replacement battery is not sufficiently charged, the node canister does not come online. The battery continues to charge.
11. Refer to “Procedure: Understanding the Storwize V7000 Gen2 system status from the LEDs” on page 256 to understand the charge level of the replacement battery. If the canister did not restart, use the management GUI to monitor the canister and battery status.
 12. Reconnect the cables to the canister, ensuring that each cable goes into the same port from which it was removed in step 2 on page 316.
 13. When the canister is back online, check the event log for any new events that might indicate a problem with the reassembly.

Replacing a battery in a power supply unit

Remove and replace the battery in a control enclosure power-supply unit.

Replacing a battery in a Storwize V7000 Gen1 power supply unit

This topic describes how to replace the battery in the control enclosure power-supply unit.

Before you begin

DANGER

When working on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard:

- If IBM supplied a power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.
- Do not open or service any power supply assembly.
- Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.
- Connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate.
- Connect any equipment that will be attached to this product to properly wired outlets.
- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.

To disconnect:

1. Turn off everything (unless instructed otherwise).
2. Remove the power cords from the outlets.
3. Remove the signal cables from the connectors.
4. Remove all cables from the devices.

To connect:

1. Turn off everything (unless instructed otherwise).
 2. Attach all cables to the devices.
 3. Attach the signal cables to the connectors.
 4. Attach the power cords to the outlets.
 5. Turn on the devices.
- Sharp edges, corners and joints might be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching. (D005)

CAUTION:

The battery is a lithium ion battery. To avoid possible explosion, do not burn. (C007)

Attention: Even though many of these components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.

Even though many of these components are hot-swappable, they are intended to be used only when your system is not active (no I/O operations). If your system is powered on and processing I/O operations, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures can result in loss of data or loss of access to data.

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.

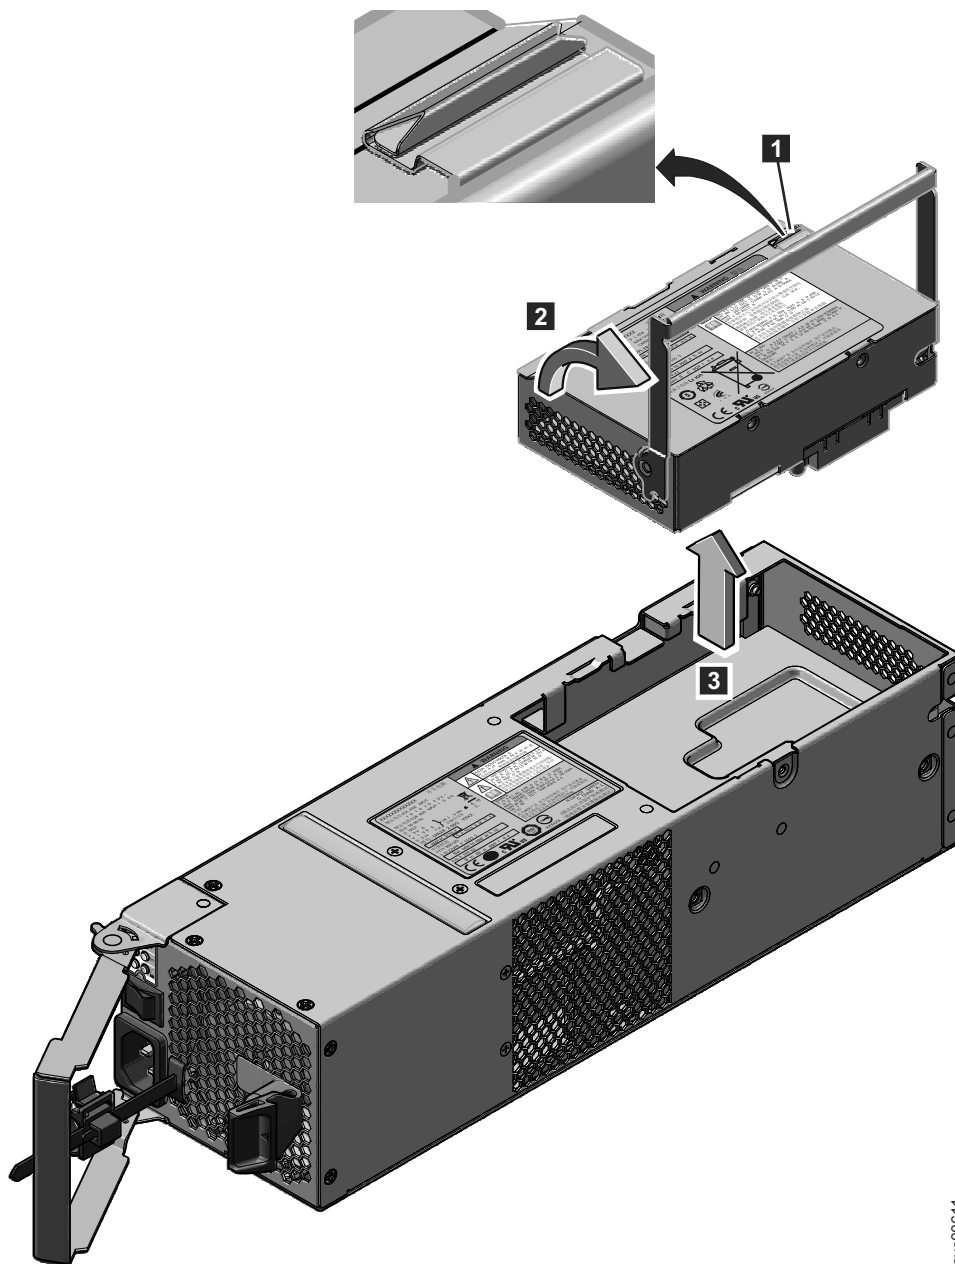
About this task

Each power supply unit in a control enclosure contains an integrated battery that is used during temporary short-term power outages. You must replace the battery with the exact same model.

To replace the battery in the power supply unit of the control enclosure, perform the following steps:

Procedure

1. Read the safety information to which “Preparing to remove and replace parts” on page 295 refers.
2. Follow the removing steps of the replacing a power-supply unit procedure. Go to “Replacing a Storwize V7000 Gen1 power supply unit for a control enclosure” on page 307.
3. Remove the battery, as shown in Figure 93 on page 320.



svc00611

Figure 93. Removing the battery from the control enclosure power-supply unit

- a. Press the catch to release the handle **1**.
- b. Lift the handle on the battery **2**.
- c. Lift the battery out of the power supply unit **3**.
4. Install the replacement battery.

Attention: The replacement battery has protective end caps that must be removed prior to use.

 - a. Remove the battery from the packaging.
 - b. Remove the end caps.
 - c. Attach the end caps to both ends of the battery that you removed and place the battery in the original packaging.

- d. Place the replacement battery in the opening on top of the power supply in its proper orientation.
 - e. Press the battery to seat the connector.
 - f. Place the handle in its downward location
5. Push the power supply unit back into the enclosure until the handle starts to move.
6. Finish inserting the power supply unit into the enclosure by closing the handle until the locking catch clicks into place.
7. Reattach the power cable and cable retention bracket.
8. Turn on the power switch to the power supply unit.

What to do next

If required, return the battery. Follow all packaging instructions, and use any packaging materials for shipping that are supplied to you.

Releasing the cable retention bracket

Release the cable retention bracket when removing the power cords from the power supply unit.

Releasing the cable retention bracket

This topic provides instructions for releasing the cable retention bracket when removing the power cords from the power supply unit.

About this task

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.

Each cable retention bracket comes attached to the back of the power supply unit by the power cord plug-in.

To release a cable retention bracket, perform these steps:

Procedure

1. Unlock the cable retention bracket that is around the end of the power cord.
2. Pull the lever next to the black plastic loop slightly towards the center of the canister.
3. Continue to pull the lever towards you as you slide the cable retention bracket away from the end of the cable.

Replacing a 3.5 inch drive assembly or blank carrier

Remove and replace a 3.5 inch drive assembly or a blank carrier.

Replacing a Storwize V7000 Gen2 3.5-inch drive assembly

You can replace a faulty 3.5-inch drive assembly with a new one received from CRU / FRU stock.

About this task

The status of the drive must be such that it is not a spare or a member. The status is shown in **Pools > Internal Storage** in the management GUI.

Attention:

- Do not replace a drive unless the drive fault LED is on or you are instructed to do so by a fix procedure.
- If the drive is a member of an array, go to the management GUI and follow the fix procedures. The fix procedures mitigate loss of data and loss of access to data and manage the system's use of the drive.
- Do not leave a drive slot empty for extended periods. Do not remove a drive assembly or a blank filler without having a replacement drive or a blank filler with which to replace it.

Procedure

To prepare to replace a drive assembly, complete the following steps.

1. Read the safety information in "Preparing to remove and replace parts" on page 295.
2. Locate the slot that contains the drive assembly that you want to replace.
 - a. Refer to "Procedure: Identifying which Storwize V7000 Gen2 enclosure or canister to service" on page 252 to ensure correct identification of the correct system or enclosure.
 - b. The drive slots on the front are numbered 1 - 12. For example, the numbering is from left to right and top to bottom:

1 2 3 4
5 6 7 8
9 10 11 12
 - c. If the drive in the slot is faulty, the lit, amber fault LED on the drive helps identify it.
3. To further help identify the drive assembly, go to the management GUI to **Pools > Internal Storage**, select the drive to replace, and click **Actions > Identify**. Verify that the correct drive fault LED begins to flash.

Attention: Never hot-swap a hard disk drive when its green activity LED is flashing. Hot-swap a drive only when its amber fault LED is lit (not flashing) or when the drive activity LED is off.

To remove a drive assembly, complete the following steps.

4. Press the latch on the right end of the tray handle to release it.
5. Pull out the tray handle to the open position (see Figure 94 on page 323).
6. Grasp the handle and pull the drive partially out of the bay.
7. Wait at least 20 seconds before you remove the drive assembly from the enclosure to enable the drive to spin down. This avoids possible damage to the drive.
8. Make sure that there is proper identification (such as a label) on the hard disk drive.
9. Gently slide it completely out of the enclosure.
10. If the drive failed, record that information on its label.

To install a drive assembly, complete the following steps.

11. Touch the static-protective package that contains the drive assembly to any unpainted surface on the outside of the enclosure.
12. Remove the drive assembly from its package.
13. Make sure that its drive-tray handle is in the open (unlocked) position.
14. Align the drive assembly with the guide rails in the bay (see Figure 95 on page 323).

15. Gently push the drive assembly into the bay until the drive stops.
16. Rotate its handle to the closed (locked) position.

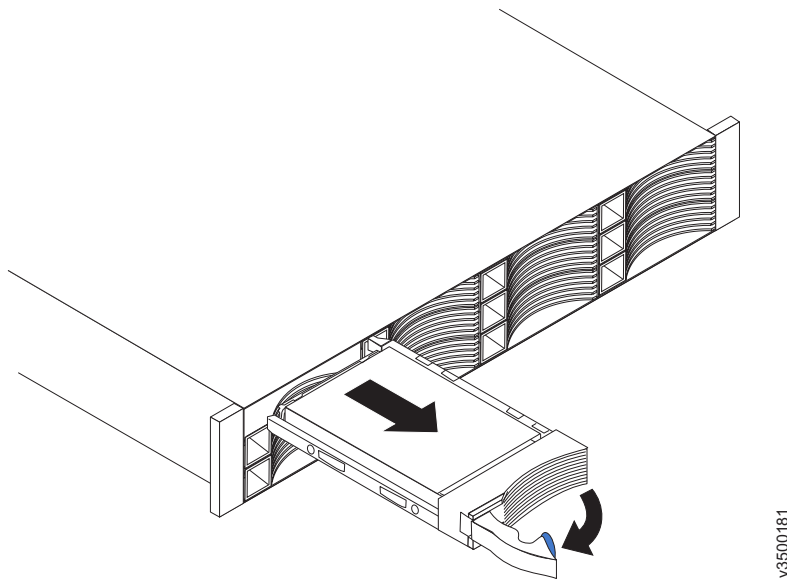


Figure 94. Unlocking and removing a 3.5-inch drive from its slot

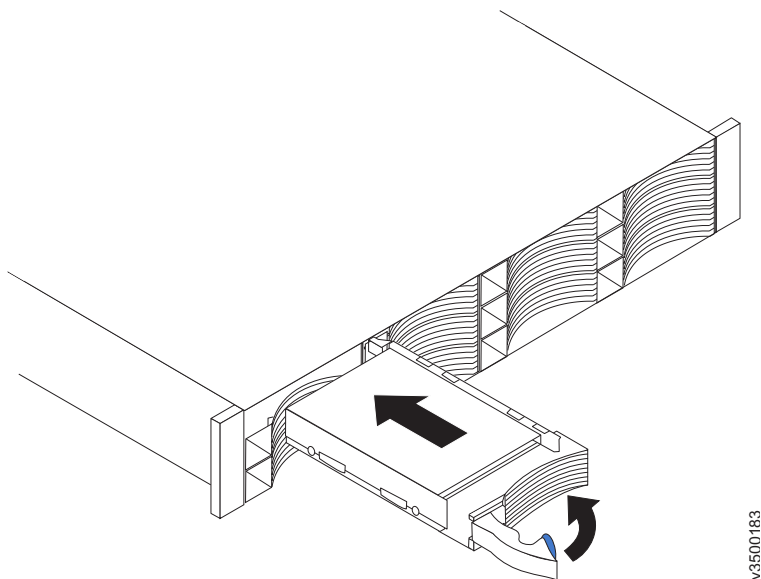


Figure 95. Installing and locking a 3.5-inch drive into its slot

Results

If the replaced drive was a failed drive, the system automatically reconfigures the replacement drive as a spare and the replaced drive is removed from the configuration. The process can take a few minutes.

Replacing a 3.5-inch drive assembly or blank carrier

This topic describes how to replace a 3.5-inch drive assembly or blank carrier.

About this task

Attention: If your drive is configured for use, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures results in loss of data or loss of access to data.

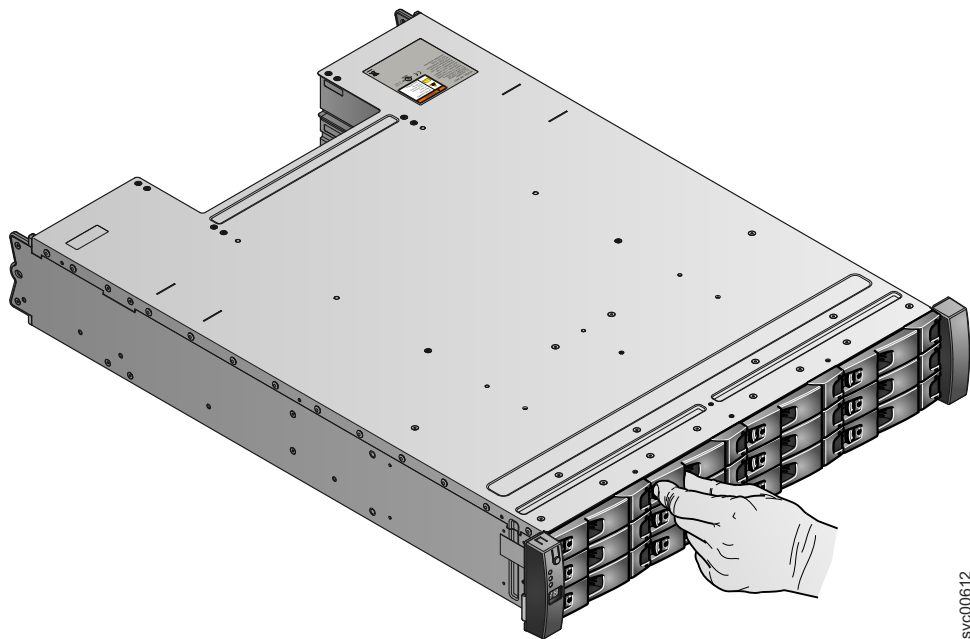
Attention: Do not leave a drive slot empty. Do not remove a drive or drive assembly before you have a replacement available.

The drives can be distinguished from the blank carriers by the color-coded striping on the drive. The drives are marked with an orange striping. The blank carriers are marked with a blue striping.

To replace the drive assembly or blank carrier, perform the following steps:

Procedure

1. Read the safety information to which “Preparing to remove and replace parts” on page 295 refers.
2. Unlock the assembly by squeezing together the tabs on the side.



svc00612

Figure 96. Unlocking the 3.5 inch drive

3. Open the handle to the full extension.

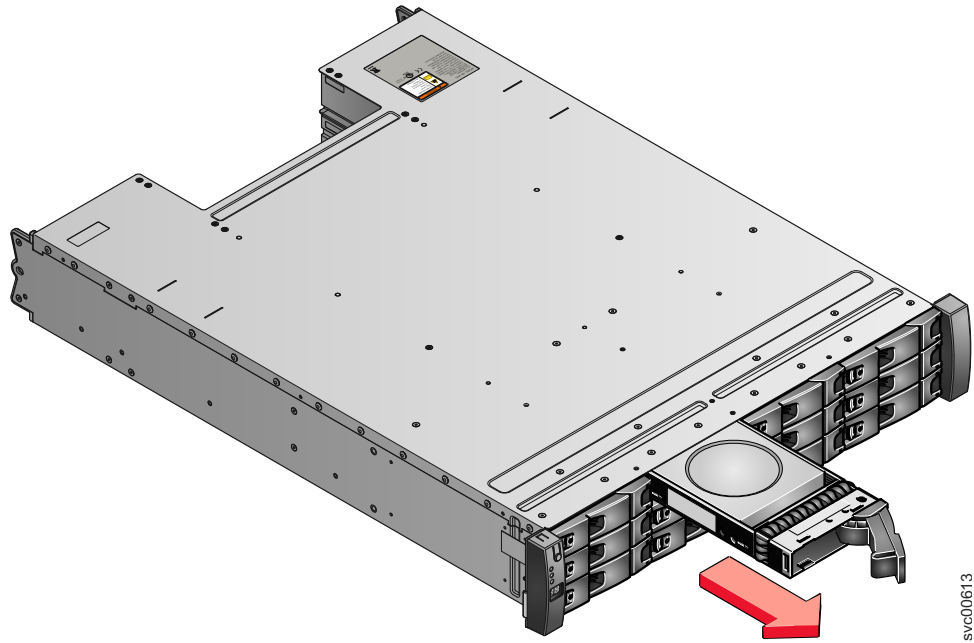


Figure 97. Removing the 3.5 inch drive

4. Pull out the drive.
5. Push the new drive back into the slot until the handle starts to move.
6. Finish inserting the drive by closing the handle until the locking catch clicks into place.

Replacing a 2.5 inch drive assembly or blank carrier

Remove and replace a 2.5 inch drive assembly or a blank carrier.

Replacing a Storwize V7000 Gen2 2.5-inch drive assembly

You can replace a faulty 2.5-inch drive assembly with a new one received from CRU / FRU stock.

About this task

The status of the drive must be such that it is not a spare or a member. The status is shown in **Pools > Internal Storage** in the management GUI.

Attention:

- Do not replace a drive unless the drive fault LED is on or you are instructed to do so by a fix procedure.
- If the drive is a member of an array, go to the management GUI and follow the fix procedures. The fix procedures mitigate loss of data and loss of access to data and manage use of the drive by the system.
- Do not leave a drive slot empty for extended periods. Do not remove a drive assembly or a blank filler without having a replacement drive or a blank filler with which to replace it.

Procedure

To prepare to replace a drive assembly, complete the following steps.

1. Read the safety information in “Preparing to remove and replace parts” on page 295.
2. Locate the slot that contains the drive assembly that you want to replace.
 - a. Refer to “Procedure: Identifying which Storwize V7000 Gen2 enclosure or canister to service” on page 252 to ensure correct identification of the correct system or enclosure.
 - b. The drive slots on the front are numbered 1 - 24, starting from the far left slot of the enclosure.
 - c. If the drive in the slot is faulty, the lit, amber fault LED on the drive helps to identify it.
3. To further help identify the drive assembly, go to the management GUI to **Pools > Internal Storage**, select the drive to replace, and click **Actions > Identify**. Verify that the correct drive fault LED begins to flash.

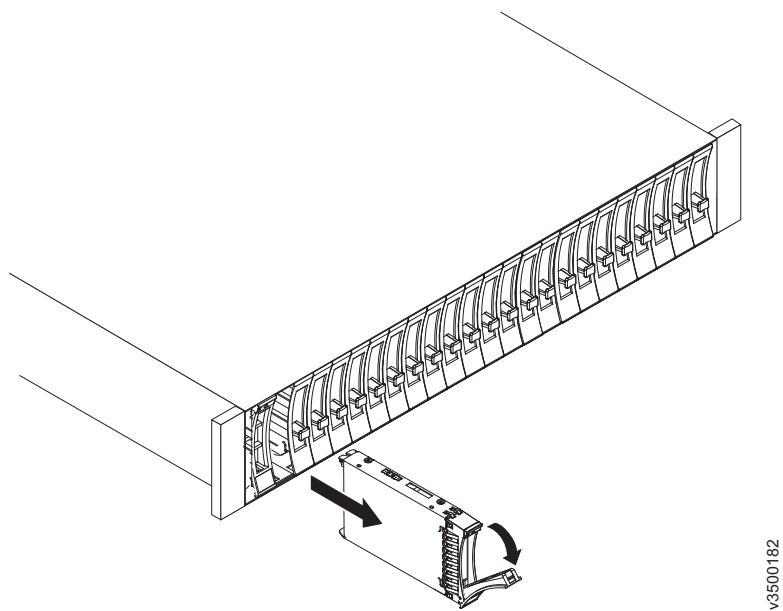
Attention: Never hot-swap a hard disk drive when its green activity LED is flashing. Hot-swap a drive only when its amber fault LED is lit (not flashing) or when the drive activity LED is off.

To remove a drive assembly, complete the following steps.

4. Gently slide the orange release latch up to unlock the handle.
5. Pull out the tray handle to the open position (see Figure 98 on page 327).
6. Grasp its handle and pull the drive partially out of the bay.
7. Wait at least 20 seconds before you remove the drive assembly from the enclosure to enable the drive to spin down. This avoids possible damage to the drive.
8. Gently slide it completely out of the enclosure.
9. Make sure that there is proper identification (such as a label) on the drive assembly. If the drive failed, record that information on the label.

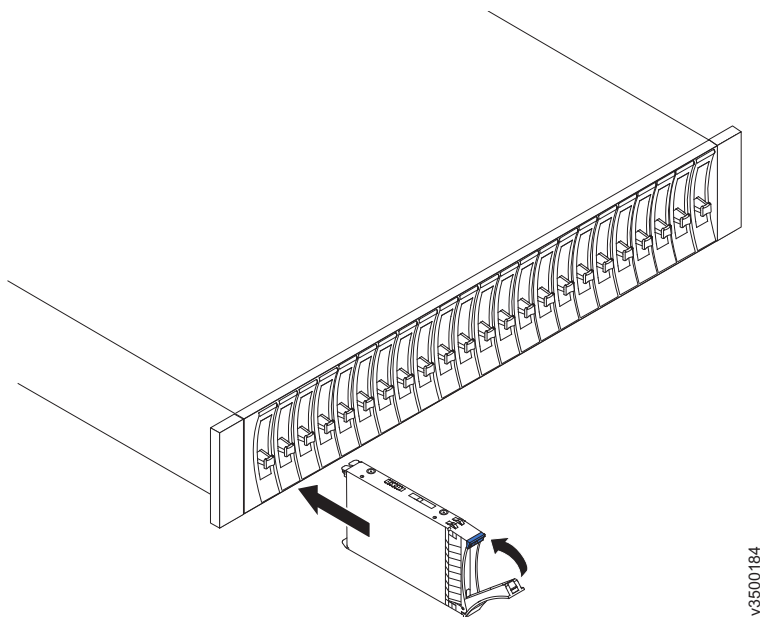
To install a drive assembly, complete the following steps.

10. Touch the static-protective package that contains the drive assembly to any unpainted surface on the outside of the enclosure.
11. Remove the drive assembly from its package.
12. Make sure that its drive-tray handle is in the open (unlocked) position.
13. Align the drive assembly with the guide rails in the bay (see Figure 99 on page 327).
14. Gently push the drive assembly into the bay until the drive stops.
15. Rotate the drive handle to the closed (locked) position.



v3500182

Figure 98. Unlocking and removing a 2.5-inch drive from its slot



v3500184

Figure 99. Installing and locking a 2.5-inch drive into its slot

Results

If the replaced drive was a failed drive, the system automatically reconfigures the replacement drive as a spare and the replaced drive is removed from the configuration. The process can take a few minutes.

Replacing a 2.5-inch drive assembly or blank carrier

This topic describes how to remove a 2.5-inch drive assembly or blank carrier.

About this task

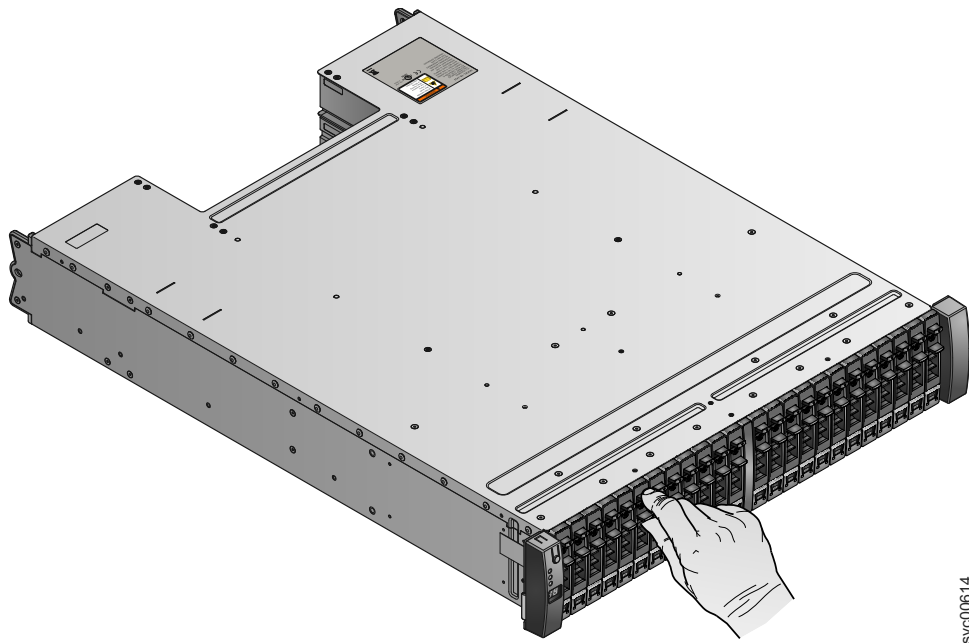
Attention: If your drive is configured for use, go to the management GUI and follow the fix procedures. Initiating the replacement actions without the assistance of the fix procedures results in loss of data or loss of access to data.

Attention: Do not leave a drive slot empty. Do not remove a drive or drive assembly before you have a replacement available.

To replace the drive assembly or blank carrier, perform the following steps:

Procedure

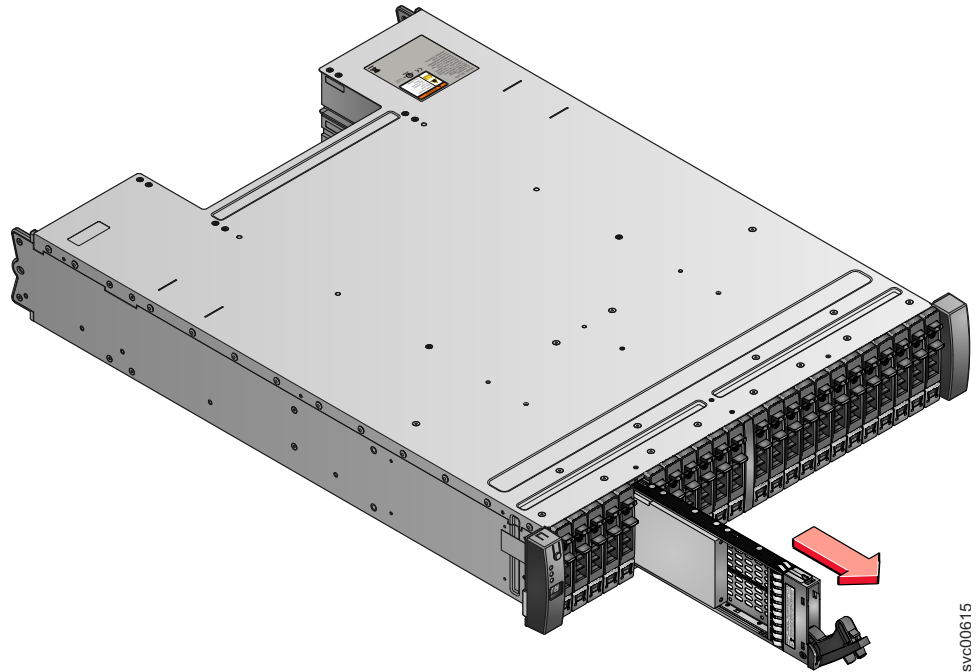
1. Read the safety information to which “Preparing to remove and replace parts” on page 295 refers.
2. Unlock the module by squeezing together the tabs at the top.



svc00614

Figure 100. Unlocking the 2.5 inch drive

3. Open the handle to the full extension.



svc00615

Figure 101. Removing the 2.5 inch drive

4. Pull out the drive.
5. Push the new drive back into the slot until the handle starts to move.
6. Finish inserting the drive by closing the handle until the locking catch clicks into place.

Replacing enclosure end caps

Remove and replace enclosure end caps.

Replacing Storwize V7000 Gen2 enclosure end caps

You can remove and replace enclosure end caps.

About this task

Attention: The left end cap is printed with information that helps identify the enclosure.

- Machine type and model
- Enclosure serial number

The information on the end cap should always match the information that is printed on the rear of the enclosure, and it should also match the information that is stored on the enclosure midplane.

Procedure

To remove and replace either the left or right end cap, complete the following steps.

1. If the enclosure is on a table or other flat surface, elevate the enclosure front slightly or carefully extend the front over the table edge.
2. Grasp the end cap by the blue touch point and pull it until the bottom edge of the end cap is clear of the bottom tab on the chassis flange.

3. Lift the end cap off the chassis flange.
4. Fit the slot on the top of the new end cap over the tab on the top of the chassis flange.
5. Rotate the end cap down until it snaps into place. Ensure that the inside surface of the end cap is flush with the chassis.

Replacing enclosure end caps

You can remove and replace enclosure end caps.

About this task

Attention: The left end cap is printed with information that helps identify the enclosure.

- Machine type and model
- Enclosure serial number
- Machine part number

The information on the end cap should always match the information printed on the rear of the enclosure, and it should also match the information that is stored on the enclosure midplane.

Procedure

To remove and replace either the left or right end cap, complete the following steps.

1. If the enclosure is on a table or other flat surface, elevate the enclosure front slightly or carefully extend the front over the table edge.
2. Grasp the end cap by the blue touch point and pull it until the bottom edge of the end cap is clear of the bottom tab on the chassis flange.
3. Lift the end cap off the chassis flange.
4. Fit the slot on the top of the new end cap over the tab on the top of the chassis flange.
5. Rotate the end cap down until it snaps into place. Ensure that the inside surface of the end cap is flush with the chassis.

Replacing a SAS cable to an expansion enclosure

Remove and replace a SAS cable to an expansion enclosure.

Replacing a Storwize V7000 Gen2 expansion enclosure attachment SAS cable

To replace a faulty Storwize V7000 Gen2 expansion enclosure attachment SAS cable with a new one received from CRU / FRU stock, use this procedure.

About this task

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.

Attention:

If you need to replace more than one cable, record which two ports, canisters, and enclosures each cable connects, so you can match the connections with the replacement cables. The system cannot operate if the expansion enclosure attachment SAS cabling is incorrect.

Expansion enclosure attachment SAS cables are connected only between SAS port 3 or 4 of a node canister and SAS port 1 of an expansion canister, or between SAS ports 1 and 2 of different expansion canisters.

More information about correct expansion enclosure attachment SAS cabling can be found in the troubleshooting description of a problem with Storwize V7000 Gen2 SAS cabling.

Procedure

To replace a SAS cable, complete the following steps.

1. Locate the connector at one end of the SAS cable that is to be removed.
2. Grasp the connector by its blue tag. Pull the tag.
3. The connector is released and slides out of the port.
4. Repeat steps 2 and 3 on the other end of the SAS cable.
5. To connect the replacement expansion enclosure attachment SAS cable, connect each end to the vacated ports.

Attention: When inserting a SAS connector into a SAS port, ensure that the orientation of the connector matches the orientation of the port before pushing the connector into the port.

- The cable connector and socket are keyed and it is important that you have proper alignment of the keys when the cable is inserted.
- Before inserting the connector into the port, ensure that the connector is rotated such that the blue tag is the lowest part.
- Figure 102 shows the correct orientation. The blue tab is always below the port for expansion enclosure attachment SAS cables.

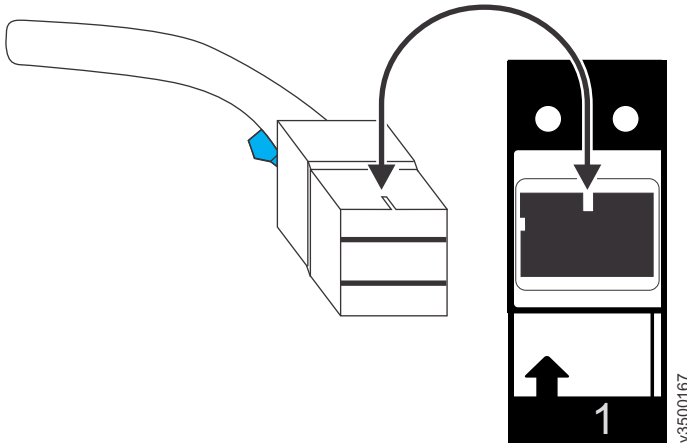


Figure 102. Proper orientation for SAS cable connector

- A click is heard or felt when the cable is successfully inserted and you should not be able to disconnect the cable without pulling on the blue tag.
- When both ends of a SAS cable are correctly connected, the green link LED next to the connected SAS ports are lit.

See the troubleshooting procedure for finding the status of SAS connections for more information.

Replacing a SAS cable

This topic describes how to replace a SAS cable.

About this task

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.

To replace a SAS cable, perform the following steps:

Procedure

1. Record which SAS cable is plugged into the specific port of the expansion canister. The cable must be inserted back into the same port after the replacement is complete; otherwise, the system cannot function properly.

Note: If you are replacing a single cable, this step is not necessary.

2. Pull the tab with the arrow away from the connector.



Figure 103. SAS cable

3. Plug the replacement cable into the specific port.
4. Ensure that the SAS cable is fully inserted. A click is heard when the cable is successfully inserted.

Replacing a control enclosure chassis

Remove and replace a control enclosure chassis. This procedure only applies to Storwize V7000 Gen1 control enclosure models.

About this task

Storwize V7000 Unified *Gen1* refers to the enclosure models in the following table:

Table 110. Storwize V7000 Unified Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 Unified control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 Unified control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 Unified control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 Unified control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 Unified expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Storwize V7000 Unified *Gen2* refers to the newer generation of enclosures in the following table:

Table 111. Storwize V7000 Unified Gen2 model numbers

Machine type/model	Description
2076-524	Storwize V7000 Unified control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-12F	Storwize V7000 Unified expansion enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Replacing a Storwize V7000 Gen2 control enclosure

You can replace a control enclosure.

Before you begin

Note: Ensure that you know the type of enclosure that you are replacing. The procedures for replacing a control enclosure are different from those procedures for replacing an expansion enclosure chassis. For information about replacing an expansion enclosure, see “Replacing an expansion enclosure chassis” on page 344.

DANGER

When working on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard:

- If IBM supplied a power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.
- Do not open or service any power supply assembly.
- Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.
- Connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate.
- Connect any equipment that will be attached to this product to properly wired outlets.
- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.

To disconnect:

1. Turn off everything (unless instructed otherwise).
2. Remove the power cords from the outlets.
3. Remove the signal cables from the connectors.
4. Remove all cables from the devices.

To connect:

1. Turn off everything (unless instructed otherwise).
 2. Attach all cables to the devices.
 3. Attach the signal cables to the connectors.
 4. Attach the power cords to the outlets.
 5. Turn on the devices.
- Sharp edges, corners and joints might be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching.
- (D005)

Attention: Perform this procedure only if instructed to do so by a service action or the IBM support center. If you have a single control enclosure, this procedure requires that you shut down your system to replace the control enclosure. If you have more than one control enclosure, you can keep part of the system running, but you lose access to the volumes that are on the affected I/O group and any volumes that are in other I/O groups that depend on the drives that are in the affected I/O group. If the system is still doing I/O requests in all the I/O groups, schedule the replacement during a maintenance period or other time when the I/O can be stopped.

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.

Ensure that you are aware of the procedures for handling static-sensitive devices before you remove the enclosure.

Procedure

To replace a control enclosure, complete the following steps:

1. If you are able to access either of the node canisters with the service assistant, record the machine type and model of the enclosure, the serial number of the enclosure, and the two world-wide node names (WWNNs) for the enclosure.
 - From the service assistant home page, open the location data for the node. Record the machine type and model (MTM), the serial number, WWNN 1 and WWNN 2 from the enclosure column.
 - If you are replacing the enclosure because neither node canister can start, retrieve this information after you finished the replacement.
 - a. Start the service assistant on one of the canisters.
 - b. Go to the node location data on the home page.
 - c. Record the machine type and model, the serial number, WWNN 1 and WWNN 2 from the node copy column.

The machine type and model and the serial number are also shown on the labels at the front and back of the enclosure.

2. If the enclosure is still active, shut down the block host I/O and the Metro Mirror and Global Mirror activity to all the volumes that depend on the affected enclosure.

This statement applies to all volumes in the I/O group that are managed by this enclosure plus any volumes in other I/O groups that depend on the drives in the affected I/O group.

3. If your system contains a single I/O group and if the clustered system is still online, shut the system down by following the procedure in “Turning off the system”.
4. If your system contains more than one I/O group, you might not need to power down the file modules if the enclosure is not connected to the file modules and if the drives that are managed by this enclosure are not used in any file volumes. Use the following procedure to see if any file volumes are affected:
 - a. Use the output from the **lsenclosure** CLI command to determine the enclosure_id for the control enclosure that is to be replaced.
 - b. Use the following CLI command to find the volumes that depend on this enclosure:

```
lsdependentvdisks -enclosure <enclosure_id>
```

Dependent volume names that start with IFS are file volumes that are used by the file modules to provide file systems. Turn off these file modules. See the procedure "Turning off the system".

5. If the I/O group is still online, shut down the I/O group by using the control enclosure CLI.
 - a. Identify the two node canisters in the I/O group that are provided by the control enclosure to be replaced.
 - b. To shut down each node, issue the following CLI command once for each of the two node canisters:
`stopssystem -force -node <node ID>`
 - c. Wait for the shutdown to complete.

6. Verify that it is safe to remove the power from the enclosure.

For each of the canisters, verify the status of the system status LED. If the LED is lit on either of the canisters, do not continue because the system is still online. Determine why the node canisters did not shut down in step 3 on page 336 or step 4 on page 336.

Note: If you continue while the system is still active, you risk losing the clustered system configuration and volume cache data that is stored in the canister.

7. Turn off the power to the enclosure using the switches.
8. Record which data cables are plugged into the specific ports. The cables must be inserted back into the same ports after the replacement is complete; otherwise, the system cannot function properly.
9. Disconnect the cable retention brackets and the power cords from the power supply units.
10. Disconnect the data cables for each canister.
11. Remove the power supply units from the enclosure.
12. Remove the canisters from the enclosure. Record the location of each canister. They must be inserted back into the same location in the new enclosure.
13. Remove the fan modules from the enclosure. Record the location of each fan module. They must be inserted back into the same location in the new enclosure.
14. Remove all the drives and blank drive assemblies from the enclosure. Record the location for each drive. They must be inserted back into the same location in the new enclosure.
15. Remove both enclosure end caps from the enclosure. Keep the left end cap because it is used again.
16. Remove the clamping screws that attached the enclosure to the rack cabinet.
17. Remove the enclosure chassis from the front of the rack cabinet and take the chassis to a work area.
18. Install the new enclosure chassis in the rack cabinet.
19. Remove the end caps from the new enclosure and install the clamping screws that attach the enclosure to the rack cabinet.
20. Replace the end caps. Use the new right end cap and use the left end cap that you removed in step 15.
Using the left end cap that you removed preserves the model and serial number identification.
21. Reinstall the drives in the new enclosure. The drives must be inserted back into the same location from which they were removed on the old enclosure.

22. Reinstall the canisters in the enclosure. The canisters must be inserted back into the same location from which they were removed on the old enclosure.
23. Reinstall the fan modules in the enclosure. The fan modules must be inserted back into the same location from which they were removed on the old enclosure.
24. Install the power supply units.
25. Reattach the data cables to each canister using the information that you recorded previously.

Note: The cables must be inserted back into the same ports from which they were removed on the old enclosure; otherwise, the system cannot function properly.

26. Attach the power cords and the cable retention brackets to the power supply units.
27. Write the old enclosure machine type and model (MTM) and serial number on the repair identification (RID) tag that is supplied. Attach the tag to the left flange at the back of the enclosure.
28. Turn on the power to the enclosure using the switches.

The node canisters boot up. The fault LEDs are on because the new enclosure was not set with the identity of the old enclosure. The node canisters report that they are in the wrong location.

- a. Connect to the service assistant on one of the node canisters to configure the machine type and model, serial number, and WWNNs that are stored in the enclosure. If you replaced a node canister, connect to the canister that was not replaced.

You can connect using the previous service address. However, it is not always possible to maintain this address. If you cannot connect through the original service address, attempt to connect using the default service address. If you still cannot access the system, see the troubleshooting description of a problem when connecting to the service assistant.

“Problem: Cannot connect to the service assistant” on page 245.

- b. Use the **Configure enclosure** panel.
- c. Select the options to **Update WWNN 1**, **Update WWNN 2**, **Update the machine type and model**, and **Update the serial number**. Do not update the system ID. Use the node copy data for each of the values. Check that these values match the values that you recorded in step 1 on page 336.

If you were not able to record the values, use the node copy values only if none of them have all zeros as their value. If any of the node copy values are all zeros, connect the service assistant to the other node canister and configure the enclosure there. If you still do not have a full set of values, contact IBM support.

After you modify the configuration, the node attempts to restart.

Note: There are situations where the canisters restart and report critical node error 508. If the node canisters fail to become active after they restart when the enclosure is updated, check their status by using the service assistant. If both node canisters show critical node error 508, use the service assistant to restart the nodes. For any other node error, see “Procedure: Fixing node errors” on page 271.

To restart a node from the service assistant, do the following steps:

- 1) Log on to the service assistant.

- 2) From the home page, select the node that you want to restart from the **Changed Node List**.
- 3) Select **Actions > Restart**.
- d. The system starts and can handle I/O requests from the host systems.

Note: The configuration changes that are described in the following steps must be done to ensure that the system is operating correctly. If you do not do these steps, the system is unable to report certain errors.
- e. Power up the file modules. See “Turning on the system”.
29. Start the management GUI and select **Monitoring > System Details**. You see an extra enclosure in the system list because the system detected the replacement control enclosure. The original control enclosure is still listed in its configuration. The original enclosure is listed with its original enclosure ID. It is offline and managed. The new enclosure has a new enclosure ID. It is online and unmanaged.
30. Select the original enclosure in the tree view.

Verify that it is offline and managed and that the serial number is correct.
31. From the **Actions** menu, select **Remove enclosure** and confirm the action. The physical hardware was already removed. You can ignore the messages about removing the hardware. Verify that the original enclosure is no longer listed in the tree view.
32. Add the new enclosure to the system.
 - a. Select the enclosure from the tree view.
 - b. From the **Actions** menu, select **Add Control and Expansion Enclosures**.
 - c. Because you already added the hardware, select **Next** on the first panel that asks you to install the hardware. The next panel shows the unmanaged new enclosure.
 - d. Follow the steps in the wizard. The wizard changes the control enclosure to Managed.
 - e. Select the enclosure and add it to the system.
33. Select the new enclosure in the tree view and verify that it is now online and managed.
34. Change the enclosure ID of the replaced enclosure to that of the original enclosure. From the **Enclosure ID** field, select the ID value of the original enclosure.
35. Check the status of all volumes and physical storage to ensure that everything is online.
36. Restart the host application and any FlashCopy activities, Global Mirror activities, or Metro Mirror activities that were stopped.

Replacing a Storwize V7000 Gen1 control enclosure chassis

You can replace a control enclosure chassis.

Before you begin

Note: Ensure that you know the type of enclosure chassis that you are replacing. The procedures for replacing a control enclosure chassis are different from those procedures for replacing an expansion enclosure chassis. To replace an expansion enclosure chassis, see “Replacing an expansion enclosure chassis” on page 348.

DANGER

When working on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard:

- If IBM supplied a power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.
- Do not open or service any power supply assembly.
- Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.
- Connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate.
- Connect any equipment that will be attached to this product to properly wired outlets.
- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.

To disconnect:

1. Turn off everything (unless instructed otherwise).
2. Remove the power cords from the outlets.
3. Remove the signal cables from the connectors.
4. Remove all cables from the devices.

To connect:

1. Turn off everything (unless instructed otherwise).
 2. Attach all cables to the devices.
 3. Attach the signal cables to the connectors.
 4. Attach the power cords to the outlets.
 5. Turn on the devices.
- Sharp edges, corners and joints might be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching.
- (D005)

Attention: Perform this procedure only if instructed to do so by a service action or the IBM support center. If you have a single control enclosure, this procedure requires that you shut down your system to replace the control enclosure. If you have more than one control enclosure, you can keep part of the system running, but you lose access to the volumes that are on the affected I/O group and any volumes that are in other I/O groups that depend on the drives that are in the affected I/O group. If the system is still performing I/O requests in all the I/O groups, schedule the replacement during a maintenance period or other time when the I/O can be stopped.

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.

Ensure that you are aware of the procedures for handling static-sensitive devices before you remove the enclosure.

Procedure

To replace a control enclosure chassis, complete the following steps.

1. If you are able to access either of the node canisters with the service assistant, record the machine type and model of the enclosure, the serial number of the enclosure, and the two WWNNs for the enclosure.
 - a. From the service assistant home page, open the location data for the node. Record the machine type and model (MTM), the serial number, WWNN 1 and WWNN 2 from the enclosure column.

Note: If you are replacing the enclosure because neither node canister can start, retrieve this information after you have completed the replacement.
 - b. Start the service assistant on one of the canisters.
 - c. Go to the node location data on the home page.
 - d. Record the machine type and model, the serial number, WWNN 1 and WWNN 2 from the node copy column. The machine type and model and the serial number are also shown on the labels at the front and back of the enclosure.
2. If the enclosure is still active, shut down the block host I/O and the Metro Mirror and Global Mirror activity to all the volumes that depend on the affected enclosure. This statement applies to all volumes in the I/O group that are managed by this enclosure plus any volumes in other I/O groups that depend on the drives in the affected I/O group.
3. If your system contains a single I/O group and if the clustered system is still online, shut the system down by following the procedure in “Turning off the system”.
4. If your system contains more than one I/O group, you might not need to power down the file modules if the enclosure is not connected to the file modules and if the drives that are managed by this enclosure are not used in any file volumes. Use the following procedure to see if any file volumes are affected:
 - a. Use the output from the **lsenclosure** CLI command to determine the enclosure_id for the control enclosure that is to be replaced.
 - b. Use the following CLI command to find the volumes that depend on this enclosure:

```
lsdependentvdisks -enclosure <enclosure_id>
```

Dependent volume names that start with IFS are file volumes that are used by the file modules to provide file systems. Turn off these file modules. See the procedure "Turning off the system".

5. If the I/O group is still online, shut down the I/O group by using the control enclosure CLI.
 - a. Identify the two node canisters in the I/O group that are provided by the control enclosure to be replaced.
 - b. To shut down each node, issue the following CLI command once for each of the two node canisters:
`stopssystem -force -node <node ID>`
 - c. Wait for the shutdown to complete.
6. Verify that it is safe to remove the power from the enclosure. For each of the canisters, verify the status of the system status LED. If the LED is lit on either of the canisters, do not continue because the system is still online. Determine why the node canisters did not shut down in step 3 on page 341 or step 4 on page 341.

Note: If you continue while the system is still active, you risk losing the clustered system configuration and volume cache data that is stored in the canister.

7. Turn off the power to the enclosure using the switches on the power supply units.
8. Record which data cables are plugged into the specific ports. The cables must be inserted back into the same ports after the replacement is complete; otherwise, the system cannot function properly.
9. Disconnect the cable retention brackets and the power cords from the power supply units.
10. Disconnect the data cables for each canister.
11. Remove the power supply units from the enclosure.
12. Remove the canisters from the enclosure. Record the location of each canister. They must be inserted back into the same location in the new enclosure.
13. Remove all the drives and blank drive assemblies from the enclosure. Record the location for each drive. They must be inserted back into the same location in the new enclosure.
14. Remove both enclosure end caps from the enclosure.

Important: Keep the left end cap because it is used again.

15. Remove the clamping screws that attached the enclosure to the rack cabinet.
16. Remove the enclosure chassis from the front of the rack cabinet and take the chassis to a work area.
17. Install the new enclosure chassis in the rack cabinet.
18. Remove the end caps from the new enclosure and install the clamping screws that attach the enclosure to the rack cabinet.
19. Replace the end caps. Use the new right end cap and use the left end cap that you removed in step 14. Using the left end cap that you removed preserves the model and serial number identification.
20. Reinstall the drives in the new enclosure. The drives must be inserted back into the same location from which they were removed on the old enclosure.
21. Reinstall the canisters in the enclosure. The canisters must be inserted back into the same location from which they were removed on the old enclosure.
22. Install the power supply units.

23. Reattach the data cables to each canister using the information that you recorded previously.

Note: The cables must be inserted back into the same ports from which they were removed on the old enclosure; otherwise, the system cannot function properly.

24. Attach the power cords and the cable retention brackets to the power supply units.
25. Write the old enclosure machine type and model (MTM) and serial number on the repair identification (RID) tag that is supplied. Attach the tag to the left flange at the back of the enclosure.
26. Turn on the power to the enclosure using the switches on the power supply units. The node canisters boot up. The fault LEDs are on because the new enclosure has not been set with the identity of the old enclosure. The node canisters log node error 504, reporting that they are in the wrong location. In the system event log these appear with an error code of 1192.
27. Connect to the service assistant on one of the node canisters to configure the machine type and model, serial number, and WWNNs that are stored in the enclosure. If you have replaced a node canister, connect to the canister that has not been replaced. You can connect using the previous service address. However, it is not always possible to maintain this address. If you cannot connect through the original service address, attempt to connect using the default service address. If you still cannot access the system, see “Problem: Cannot connect to the service assistant” on page 245.
28. Use the **Configure enclosure** panel.
29. Select the options to **Update WWNN 1**, **Update WWNN 2**, **Update the machine type and model**, and **Update the serial number**. Do not update the system ID. Use the node copy data for each of the values. Check that these values match the values that you recorded in step 1 on page 341. If you were not able to record the values, use the node copy values only if none of them have all zeroes as their value. If any of the node copy values are all zeroes, connect the service assistant to the other node canister and configure the enclosure there. If you still do not have a full set of values, contact IBM support.

Important: Step 30 writes the enclosure identity into the replacement midplane. The replacement midplane cannot be used as a replacement part for a different enclosure after step 30 is completed.

30. Click the **Modify** button. The node writes the data and attempts to restart.

Note: There are situations where the canisters restart and report critical node error 508. If the node canisters fail to become active after they restart when the enclosure is updated, check their status by using the service assistant. If both node canisters show critical node error 508, use the service assistant to restart the nodes. For any other node error, see “Procedure: Fixing node errors” on page 271. To restart a node from the service assistant, perform the following steps:

- a. Log on to the service assistant.
- b. From the home page, select the node that you want to restart from the **Changed Node List**.
- c. Select **Actions > Restart**.

The system starts and can handle I/O requests from the host systems.

Note: The configuration changes that are described in the following steps must be performed to ensure that the system is operating correctly. If you do not perform these steps, the system is unable to report certain errors.

31. Power up the file modules. See “Turning on the system”.
32. Start the management GUI and select **Monitoring > System**. You see an additional enclosure in the system list because the system has detected the replacement control enclosure. The original control enclosure is still listed in its configuration. The original enclosure is listed with its original enclosure ID. It is offline and managed. The new enclosure has a new enclosure ID. It is online and unmanaged.
33. Select the original enclosure. Verify that it is offline and managed and that the serial number is correct.
34. Right-click the enclosure and select **Remove**. The physical hardware has already been removed. You can ignore the messages about removing the hardware. Verify that the original enclosure is no longer listed in the tree view.
35. Add the new enclosure to the system.
 - a. From the **Actions** menu, select **Add Enclosures**.
 - b. Because you have already added the hardware, select **Next** on the first panel that asks you to install the hardware. The next panel shows the unmanaged new enclosure.
 - c. Follow the steps in the wizard. The wizard changes the control enclosure to Managed.
 - d. Select the enclosure and add it to the system.
36. Select the new enclosure and verify that it is now online and managed.
37. Change the enclosure ID of the replaced enclosure to that of the original enclosure. Right-click the control enclosure and select **Modify ID** and change the ID to that of the original enclosure.
38. Check the status of all volumes and physical storage to ensure everything is online.
39. Restart the host application and any FlashCopy activities, Global Mirror activities, or Metro Mirror activities that were stopped.

Results

Replacing an expansion enclosure chassis

Remove and replace an expansion enclosure chassis. This procedure only applies to Storwize V7000 Gen1 enclosure models.

About this task

Storwize V7000 Unified *Gen1* refers to the enclosure models in the following table:

Table 112. Storwize V7000 Unified Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 Unified control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 Unified control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 Unified control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)

Table 112. Storwize V7000 Unified Gen1 model numbers (continued)

Machine type/model	Description
2076-324	Storwize V7000 Unified control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 Unified expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Storwize V7000 Unified Gen2 refers to the newer generation of enclosures in the following table:

Table 113. Storwize V7000 Unified Gen2 model numbers

Machine type/model	Description
2076-524	Storwize V7000 Unified control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-12F	Storwize V7000 Unified expansion enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Replacing a Storwize V7000 Gen2 expansion enclosure

You can replace an expansion enclosure.

Before you begin

Note: Ensure that you know the type of enclosure chassis that you are replacing. The procedures for replacing an expansion enclosure chassis are different from those procedures for replacing a control enclosure chassis. To replace a control enclosure chassis, see “Replacing a control enclosure chassis” on page 333.

DANGER

When working on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard:

- If IBM supplied a power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.
- Do not open or service any power supply assembly.
- Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.
- Connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate.
- Connect any equipment that will be attached to this product to properly wired outlets.
- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.

To disconnect:

1. Turn off everything (unless instructed otherwise).
2. Remove the power cords from the outlets.
3. Remove the signal cables from the connectors.
4. Remove all cables from the devices.

To connect:

1. Turn off everything (unless instructed otherwise).
 2. Attach all cables to the devices.
 3. Attach the signal cables to the connectors.
 4. Attach the power cords to the outlets.
 5. Turn on the devices.
- Sharp edges, corners and joints might be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching.
- (D005)

Attention: If your system is powered on and performing I/O operations, go to the management GUI and follow the fix procedures. Performing the replacement actions without the assistance of the fix procedures can result in loss of data or access to data.

Even though many parts are hot-swappable, these procedures are intended to be used only when your system is not up and running and performing I/O operations. Unless your system is offline, go to the management GUI and follow the fix procedures.

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.

Ensure that you are aware of the procedures for handling static-sensitive devices before you remove the enclosure.

About this task

Note: If your system is online, replacing an expansion enclosure can cause one or more of your volumes to go offline or your quorum disks to be inaccessible. Before you proceed with these procedures, verify which volumes might go offline. From the management GUI, go to **Home > Manage Devices**. Select the enclosure that you want to replace. Then, select **Show Dependent Volumes** in the **Actions** menu.

Procedure

To replace an expansion enclosure chassis, perform the following steps:

1. Shut down the I/O activity to the enclosure, which includes host access to GPFS file systems, FlashCopy, Metro Mirror, and Global Mirror access.
2. Turn off the power to the enclosure by disconnecting the power cable.
3. Record which data cable is plugged into each specific port. The cables must be inserted back into the same ports after the replacement is complete; otherwise, the system cannot function properly.
4. Disconnect the data cables for each canister.
5. Remove the power supply units from the enclosure.
6. Remove the canisters from the enclosure.
7. Remove all the drives and blank drive assemblies from the enclosure. Record the location for each drive. They must be inserted back into the same location in the new enclosure.
8. Remove both enclosure end caps from the enclosure. Keep the left end cap because it is used again.
9. Remove the clamping screws that attached the enclosure to the rack cabinet.
10. Remove the enclosure chassis from the front of the rack cabinet and take the chassis to a work area.
11. Install the new enclosure chassis in the rack cabinet.
12. Remove the end caps from the new enclosure and install the clamping screws that attach the enclosure to the rack cabinet.
13. Replace the end caps. Use the new right end cap and use the left end cap that you removed in step 8.

Using the left end cap that you removed preserves the model and serial number identification.

14. Reinstall drives in the new enclosure. You must insert the drives back into the same location from which they were removed on the old enclosure.
15. Reinstall the canisters (and drives) in the enclosure.
16. Install the power supply units.
17. Use the information that you recorded previously to reattach the data cables to each canister.

Note: The cables must be inserted back into the same ports from which they were removed on the old enclosure; otherwise, the system cannot function properly.

18. Attach the power cords and the cable retention brackets to the power supply units.
19. Write the old enclosure machine type and model (MTM) and serial number on the repair identification (RID) tag that is supplied. Attach the tag to the left flange at the back of the enclosure.

Results

The system records an error that indicates that an enclosure FRU replacement was detected. Go to the management GUI to use the fix procedure to change the machine type and model and serial number in the expansion enclosure.

Replacing an expansion enclosure chassis

You can replace an expansion enclosure chassis.

Before you begin

DANGER

When working on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard:

- If IBM supplied a power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.
- Do not open or service any power supply assembly.
- Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.
- Connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate.
- Connect any equipment that will be attached to this product to properly wired outlets.
- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.

To disconnect:

1. Turn off everything (unless instructed otherwise).
2. Remove the power cords from the outlets.
3. Remove the signal cables from the connectors.
4. Remove all cables from the devices.

To connect:

1. Turn off everything (unless instructed otherwise).
 2. Attach all cables to the devices.
 3. Attach the signal cables to the connectors.
 4. Attach the power cords to the outlets.
 5. Turn on the devices.
- Sharp edges, corners and joints might be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching.
- (D005)

Attention: If your system is powered on and performing I/O operations, go to the management GUI and follow the fix procedures. Performing the replacement actions without the assistance of the fix procedures can result in loss of data or access to data.

Even though many of the parts are hot-swappable, these procedures are intended to be used only when your system is not up and running and performing I/O operations. Unless your system is offline, go to the management GUI and follow the fix procedures.

Be careful when you are replacing the hardware components that are located in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.

Ensure that you are aware of the procedures for handling static-sensitive devices before you remove the enclosure.

About this task

Note: If your system is online, replacing an expansion enclosure can cause one or more of your volumes to go offline or your quorum disks to be inaccessible. Before you proceed with these procedures, verify which volumes might go offline. From the management GUI, go to **Home > Manage Devices**. Select the enclosure that you want to replace. Then select **Show Dependent Volumes** in the **Actions** menu.

To replace an expansion enclosure chassis, perform the following steps:

Procedure

1. Shut down the I/O activity to the enclosure, which includes host access to GPFS file systems, FlashCopy, Metro Mirror and Global Mirror access.
2. Turn off the power to the enclosure by using the switches on the power supply units.
3. Record which data cables are plugged into the specific ports. The cables must be inserted back into the same ports after the replacement is complete; otherwise, the system cannot function properly.
4. Disconnect the cable retention brackets and the power cords from the power supply units.
5. Disconnect the data cables for each canister.
6. Remove the power supply units from the enclosure.
7. Remove the canisters from the enclosure.
8. Remove all the drives and blank drive assemblies from the enclosure. Record the location for each drive. They must be inserted back into the same location in the new enclosure.
9. Remove both enclosure end caps from the enclosure. Keep the left end cap because it is used again.
10. Remove the clamping screws that attached the enclosure to the rack cabinet.
11. Remove the enclosure chassis from the front of the rack cabinet and take the chassis to a work area.
12. Install the new enclosure chassis in the rack cabinet.
13. Remove the end caps from the new enclosure and install the clamping screws that attach the enclosure to the rack cabinet.

14. Replace the end caps. Use the new right end cap and use the left end cap that you removed in step 9 on page 350. Using the left end cap that you removed preserves the model and serial number identification.
15. Reinstall the drives in the new enclosure. The drives must be inserted back into the same location from which they were removed on the old enclosure.
16. Reinstall the canisters in the enclosure.
17. Install the power supply units.
18. Reattach the data cables to each canister by using the information that you recorded previously.

Note: The cables must be inserted back into the same ports from which they were removed on the old enclosure; otherwise, the system cannot function properly.

19. Attach the power cords and the cable retention brackets to the power supply units.
20. Write the old enclosure machine type and model (MTM) and serial number on the repair identification (RID) tag that is supplied. Attach the tag to the left flange at the back of the enclosure.
21. Turn on the power to the enclosure by using the switches on the power supply units. The system records an error that indicates that an enclosure FRU replacement was detected.

Important: Step 22 writes the enclosure identity into the replacement midplane. The replacement midplane cannot be used as a replacement part for a different enclosure after step 22 is completed.

22. Go to the management GUI to use the fix procedure to change the machine type and model and serial number in the expansion enclosure.

Replacing a Storwize V7000 Gen2 enclosure midplane

A trained service provider must replace the midplane assembly of a Storwize V7000 Gen2 enclosure.

About this task

DANGER

When working on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard:

- If IBM supplied a power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.
- Do not open or service any power supply assembly.
- Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.
- Connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate.
- Connect any equipment that will be attached to this product to properly wired outlets.
- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.

To disconnect:

1. Turn off everything (unless instructed otherwise).
2. Remove the power cords from the outlets.
3. Remove the signal cables from the connectors.
4. Remove all cables from the devices.

To connect:

1. Turn off everything (unless instructed otherwise).
 2. Attach all cables to the devices.
 3. Attach the signal cables to the connectors.
 4. Attach the power cords to the outlets.
 5. Turn on the devices.
- Sharp edges, corners and joints might be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching.
- (D005)

Attention:

- The enclosure midplane must be replaced only by a trained service provider. Perform this procedure only if instructed to do so by a service action or the IBM support center.
- Be careful when you are replacing the hardware components that are in the back of the system that you do not inadvertently disturb or remove any cables that you are not instructed to remove.
- Ensure that you are aware of the procedures for handling static-sensitive devices before you remove the enclosure.

Replacing a Storwize V7000 Gen2 control enclosure midplane assembly

A trained service provider can use this procedure to replace a faulty Storwize V7000 Gen2 control enclosure midplane with a new one received from CRU / FRU stock. Ensure that your control enclosure midplane assembly is replaced only by a trained service provider.

Before you begin

Three persons are required at step 14 on page 356.

About this task

Follow all safety precautions when completing this procedure.

DANGER

When working on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard:

- If IBM supplied a power cord(s), connect power to this unit only with the IBM provided power cord. Do not use the IBM provided power cord for any other product.
- Do not open or service any power supply assembly.
- Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.
- The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.
- Connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate.
- Connect any equipment that will be attached to this product to properly wired outlets.
- When possible, use one hand only to connect or disconnect signal cables.
- Never turn on any equipment when there is evidence of fire, water, or structural damage.
- Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.
- Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.

To disconnect:

1. Turn off everything (unless instructed otherwise).
2. Remove the power cords from the outlets.
3. Remove the signal cables from the connectors.
4. Remove all cables from the devices.

To connect:

1. Turn off everything (unless instructed otherwise).
 2. Attach all cables to the devices.
 3. Attach the signal cables to the connectors.
 4. Attach the power cords to the outlets.
 5. Turn on the devices.
- Sharp edges, corners and joints might be present in and around the system. Use care when handling equipment to avoid cuts, scrapes and pinching. (D005)

Attention:

The control enclosure must be replaced only by a trained service provider. Complete this procedure only if instructed to do so by a service action or the IBM support center.

If you have a single control enclosure, this procedure requires that you shut down your system to replace the control enclosure midplane assembly. If you have more than one control enclosure, you can keep part of the system running, but you lose access to the volumes that are on the affected I/O group and any volumes that are in other I/O groups that depend on the drives that are in the affected I/O group. If the system is still doing I/O requests in all the I/O groups, schedule the replacement during a maintenance period or other time when the I/O can be stopped.

When replacing hardware components in the back of the enclosure, ensure that you do not inadvertently disturb or remove cables that you are not instructed to remove.

Ensure that you are aware of procedures for handling static-sensitive devices before you remove the enclosure.

Procedure

To replace the control enclosure midplane, complete the following steps:

1. Log in to the service assistant on one of the node canisters in the control enclosure.
2. Navigate to the **Enclosure Information** panel.

Important: Do NOT select the **Reset the system ID** check box.

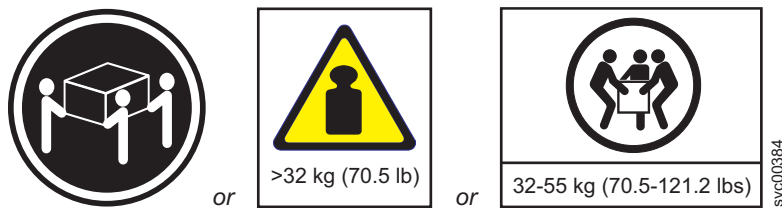
Record the following information for use in subsequent steps:

- WWNN 1
 - WWNN 2
 - Machine type and model
 - Serial number
3. Read the safety information in “Preparing to remove and replace parts” on page 295.
 4. If the control enclosure is still active, stop host I/O and Metro Mirror and Global Mirror activity on all the volumes that depend on the enclosure. This step applies to all I/O group volumes that are managed by this enclosure plus any volumes in other I/O groups that depend on the drives in the affected I/O group.
 5. Complete “Procedure: Powering off a Storwize V7000 Gen2 control enclosure” on page 282 for the control enclosure that requires the midplane assembly replacement.
 6. Disconnect both power cables from the rear of the enclosure.
 7. Write down which port connects to which cable before disconnecting all cables from the rear of the enclosure.
 8. Carefully remove each drive and label it with the drive slot from which it was removed.

You can use the drive-slot information to insert the drives into the correct drive slots at the end of this procedure.

9. Remove the two power supplies from the enclosure. Refer to "Replacing a Storwize V7000 Gen2 power supply unit for a control enclosure" on page 306 for guidance.
10. Remove the node canisters from the enclosure. Label them to indicate what canister came from each canister slot.
11. Remove the fan modules from the enclosure, as described in "Replacing a Storwize V7000 Gen2 fan module" on page 298.
12. Remove the end caps from the enclosure, as described in "Replacing Storwize V7000 Gen2 enclosure end caps" on page 329.
13. Remove the two M5 screws from the front of the enclosure to free the enclosure from the rack.
14. Slide the enclosure from the rack, and then place the enclosure on a work surface, so that the underside of the enclosure faces upward, and the enclosure front is facing toward you.

CAUTION:



The weight of this part or unit is between 32 and 55 kg (70.5 and 121.2 lb). It takes three persons to safely lift this part or unit. (C010)

15. Remove the four screws from the bottom of the enclosure. Three screws are near the front and one is near the middle. Label these screws to indicate the location from which they are removed and place them aside. Figure 104 illustrates the location of the screws on the bottom of the enclosure.



Figure 104. Bottom enclosure screws

Note: A PH1 screw driver is used for the screws in this step and the following steps. A pair of pliers is needed for the screw-pins in the following steps.

16. Turn the enclosure over again so that the top of the enclosure is facing upward and the front is facing towards you.
17. Remove the three screws and one screw-pin on the right side that secure the midplane assembly to the enclosure. Label each screw to indicate the removal location and place the screws aside. Figure 105 illustrates the location of the screws and screw-pin on the right-side of the enclosure.



Figure 105. Right-side enclosure screws

18. Remove the three screws and one screw-pin on the left side that secure the midplane assembly to the enclosure. Label each screw to indicate the removal location and place the screws aside. Figure 106 illustrates the location of the screws and screw-pin on the left-side of the enclosure.



Figure 106. Left-side enclosure screws

19. Remove the midplane assembly from the chassis by rotating up the midplane assembly to about 45°, then withdraw the midplane assembly from the front of the enclosure. Figure 107 on page 358 shows the midplane assembly at a 45 degree angle.



Figure 107. Angled midplane assembly

20. Unpack the replacement midplane assembly. Grasp the midplane assembly with two hands to hold the assembly at a 45° angle.
21. Insert the tabs on the midplane assembly into the tab holes in the enclosure and rotate down the front of the assembly.
22. Secure the midplane assembly to the enclosure chassis on both the right and left sides using six screws and two screw-pins that you removed in steps 16 on page 357 and 17 on page 357.
23. Turn over the bottom of the enclosure to face upward, then insert the four screws on the bottom of the enclosure that were removed in step 15 on page 356.
24. Reinstall the enclosure in the rack cabinet, securing it with two screws that were removed at step 13 on page 356.
25. Reinstall the end caps at the front of the enclosure, as described in “Replacing Storwize V7000 Gen2 enclosure end caps” on page 329.
26. Reinstall the hard disk drives at the front of the enclosure, making sure that each drive is inserted into the same slot from which it was removed.
27. Replace the fan modules, as described in “Replacing a Storwize V7000 Gen2 fan module” on page 298.
28. Reinstall the canisters into the same canister slots from which you removed them.
29. Reinstall the two power supplies.
30. Reconnect the data cables at the rear of the enclosure into the same connectors from which you removed them.
31. Reconnect power to the control enclosure. The node canisters restart. The fault LEDs are on because the new enclosure has not been set with the identity of the old enclosure. The node canisters log node error 504, reporting that they are in the wrong location. In the system event log, the error code is 1192.

32. Connect to the service assistant on one of the node canisters to configure the machine type and model, serial number, and WWNNs that are stored in the enclosure. If you have replaced a node canister, connect to the canister that has not been replaced. The service assistant retains a copy of the same information that was on the faulty enclosure midplane assembly. You can connect using the previous service address. However, it is not always possible to maintain this address. If you cannot connect through the original service address, attempt to connect using the default service address. If you still cannot access the system, see “Problem: Cannot connect to the service assistant” on page 245.

33. Use the **Configure enclosure** panel.

34. Use the node copy data that you recorded in step 2 on page 355 to update each of these values: **Update WWNN 1**, **Update WWNN 2**, **Update the machine type and model**, and **Update the serial number**.

Attention: Do **not** update the system ID.

If you were not able to record the values, use the node copy values only if none of them have all zeros as their value. If any of the node copy values are all zeros, connect the service assistant to the other node canister and configure the enclosure there. If you still do not have a full set of values, contact IBM support.

Important: Step 35 writes the enclosure identity into the replacement midplane. The replacement midplane cannot be used as a replacement part for a different enclosure after step 35 is completed.

35. In the **Enclosure Information** panel, click **Modify**. The node canisters restart. When the restart finishes, the system comes online with both node canisters online.

Note: In some situations, the canisters restart and report critical node error 508. If the node canisters fail to become active after they restart when the enclosure is updated, check their status by using the service assistant. If both node canisters show critical node error 508, use the service assistant to restart the nodes. For any other node error, see “Procedure: Fixing node errors” on page 271. To restart a node from the service assistant, complete the following steps:

- a. Log on to the service assistant.
- b. From the home page, select the node that you want to restart from the **Changed Node List**.
- c. Select **Actions > Restart**.

The system starts and can handle I/O requests from any host systems.

36. Power up the file modules. See “Turning on the system”.

37. Use the management GUI to check the status of all volumes and physical storage to ensure that everything is online.

38. Go to **Monitoring > Events** to check the event log for other events or errors.

39. Restart the host application and any FlashCopy activities, Global Mirror activities, or Metro Mirror activities that were stopped.

Replacing a Storwize V7000 Gen2 expansion enclosure midplane assembly

A trained service provider can use this procedure to replace a faulty Storwize V7000 Gen2 expansion enclosure midplane assembly with a new one received from CRU / FRU stock.

Before you begin

Three persons are required at step 11 on page 361.

About this task

Attention: To prevent data loss, you must shut down the system before you begin the procedure to replace an expansion enclosure midplane assembly.

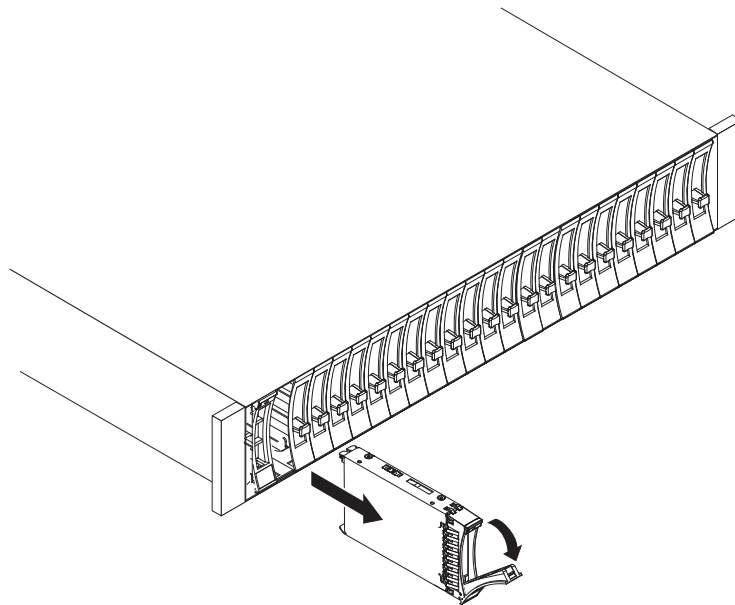
The expansion enclosure midplane assembly must be replaced only by a trained service provider.

There are two models of expansion enclosure. Before proceeding to replace an expansion enclosure midplane assembly, ensure the FRU part number of the replacement part matches that of the enclosure being repaired.

Procedure

To replace the expansion enclosure midplane, complete the following steps.

1. Read the safety information in “Preparing to remove and replace parts” on page 295.
2. Read “Procedure: Understanding Storwize V7000 Gen2 volume dependencies” on page 287 to determine whether to continue this procedure.
3. Disconnect each power supply unit in the expansion enclosure from its power outlet, so that the expansion enclosure is powered off.
4. Confirm that all the LEDs on the rear of the enclosure are off.
5. Disconnect all cables, labeling each cable to record exactly which port it was attached to (so that the cables can be inserted back into the same ports).
6. Carefully remove each hard disk drive and label it with the drive slot from which it was removed (so that the drives can be inserted back into the same slots). Refer to Figure 108 or Figure 109 on page 361.



v3500182

Figure 108. Removing a vertical style hard disk drive

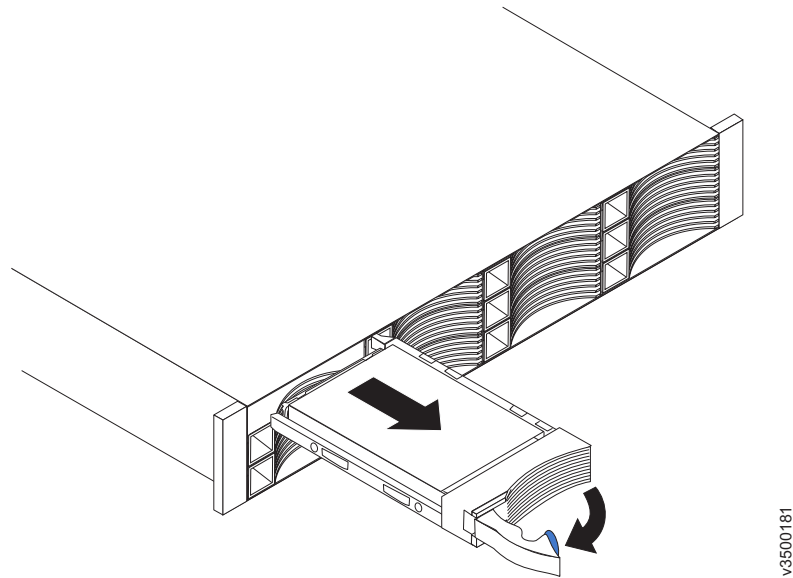
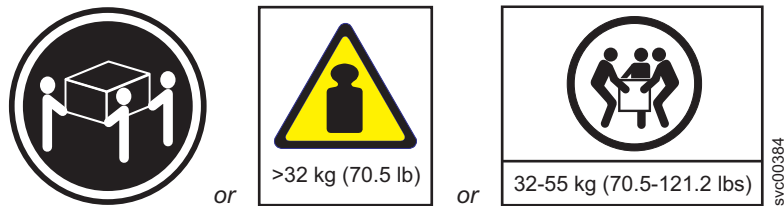


Figure 109. Removing a horizontal style hard disk drive

7. Remove the two power supplies from the enclosure. Refer to “Replacing a power supply unit for a Storwize V7000 Gen2 expansion enclosure” on page 311 for guidance.
8. Remove the expansion canisters from the enclosure. Label them to indicate which canister came from which slot.
9. Remove the end caps from the enclosure, as described in “Replacing Storwize V7000 Gen2 enclosure end caps” on page 329.
10. Remove the two screws securing the front of the enclosure into the rack. Label these screws to indicate the location from which they are removed and place them aside.
11. Slide the enclosure from the rack cabinet, turn it onto its back so that the bottom is facing upwards, and place the enclosure on a flat surface.

CAUTION:



The weight of this part or unit is between 32 and 55 kg (70.5 and 121.2 lb). It takes three persons to safely lift this part or unit. (C010)

12. Remove the four screws from the bottom of the enclosure (see Figure 110 on page 362). Remove the three screws that are near the front and the screw that is near the middle. Label these screws to indicate the location from which they are removed and place them aside.

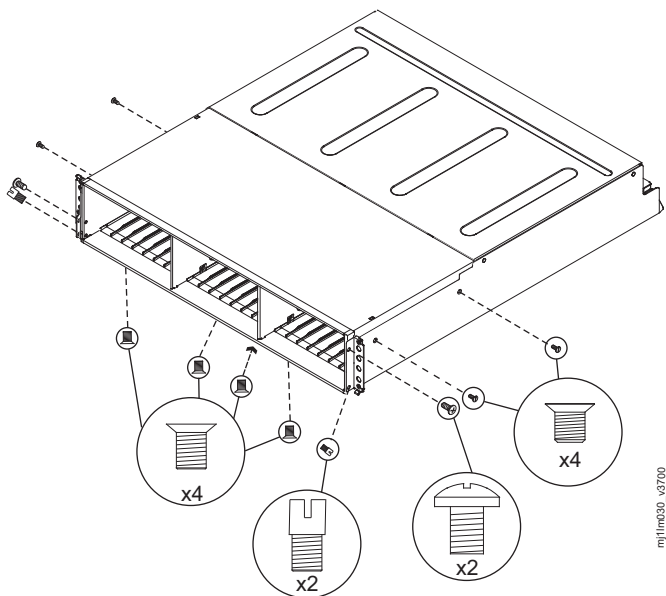


Figure 110. Removing the screws of an expansion enclosure assembly

13. Turn the enclosure top side up and place it on a flat surface.
14. Remove the three screws and one screw-pin on the right side that secure the midplane assembly to the enclosure (see Figure 110). Label the screws to indicate the location from which they are removed and place them aside.
15. Remove the three screws and one screw-pin on the left side that secure the midplane assembly to the enclosure (see Figure 110). Label the screws to indicate the location from which they are removed and place them aside. See Figure 4.
16. Remove the midplane assembly from the chassis by rotating the midplane assembly up about 45° and then lifting it out. Set the midplane assembly on a flat surface.
17. Unpack the replacement midplane assembly. Grasp the midplane assembly with two hands and hold it at a 45° angle.
18. Insert the tabs on the midplane assembly into the tab holes in the enclosure and rotate the front of the assembly down.
19. Secure the midplane assembly to the chassis on both the right and left sides of the enclosure by using the six screws and two screw-pins that you removed in steps 14 and 15.
20. Turn the enclosure over so the bottom faces upwards and insert the four screws on the bottom of the enclosure that you removed in step 12 on page 361.
21. Reinstall the enclosure in the rack cabinet, securing it with the two screws that are removed at step 10 on page 361.
22. Reinstall the end caps at the front of the enclosure, as described in “Replacing Storwize V7000 Gen2 enclosure end caps” on page 329.
23. Reinstall the hard disk drives at the front of the enclosure. Ensure that each drive is inserted back in the same slot from which it was removed.
24. Reinstall the canisters into the same slots they were removed from.
25. Reinstall the two power supplies.
26. Reconnect the data cables at the rear of the enclosure.

27. Reconnect the power to the expansion enclosure. The expansion canisters restart and the system logs an error in the event log alerting you to the unrecognized enclosure.

Important: Step 28 writes the enclosure identity into the replacement midplane. The replacement midplane cannot be used as a replacement part for a different enclosure after step 28 is completed.

28. Go to **Monitoring > Events** in the management GUI. Find the error relating to the enclosure ID of the replaced enclosure and run the fix procedure for the error.

Replacing the support rails

Remove and replace the support rails. The procedure differs, depending on the generation of your control enclosure model.

About this task

Storwize V7000 Unified *Gen1* refers to the enclosure models in the following table:

Table 114. Storwize V7000 Unified Gen1 model numbers

Machine type/model	Description
2076-112	Storwize V7000 Unified control enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-124	Storwize V7000 Unified control enclosure for up to 24 2.5-inch (6.35 cm) drives
2076-312	Storwize V7000 Unified control enclosure for 3.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-324	Storwize V7000 Unified control enclosure for 2.5-inch drives (with two 10 Gbps iSCSI/FCoE Ethernet ports)
2076-212	Storwize V7000 Unified expansion enclosure for 3.5-inch drives
2076-224	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Storwize V7000 Unified *Gen2* refers to the newer generation of enclosures in the following table:

Table 115. Storwize V7000 Unified Gen2 model numbers

Machine type/model	Description
2076-524	Storwize V7000 Unified control enclosure, with up to 24 2.5-inch (6.35 cm) drives
2076-12F	Storwize V7000 Unified expansion enclosure for up to 12 3.5-inch (8.89 cm) drives
2076-24F	Storwize V7000 Unified expansion enclosure for 2.5-inch drives

Replacing the Storwize V7000 Gen2 control enclosure support rails

You can replace faulty support rails with new ones received from CRU / FRU stock.

Before you begin

Three persons are required at step 7

About this task

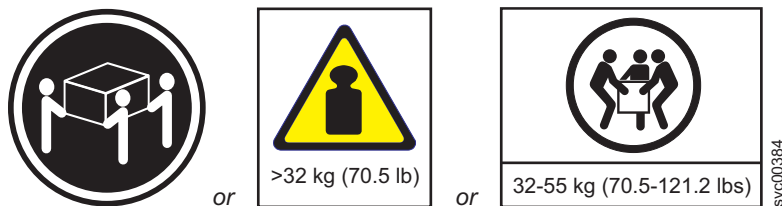
Follow all safety precautions when completing this procedure.

Procedure

To replace the support rails, complete the following steps.

1. Identify the enclosure mounted on the rails being replaced.
Follow the steps in “Procedure: Identifying which Storwize V7000 Gen2 enclosure or canister to service” on page 252 to ensure that you identify the correct enclosure.
2. Shut down the system by following the steps in “Procedure: Powering off your Storwize V7000 Gen2 system” on page 280.
3. Remove power from the enclosure by unplugging both power cables from the electrical outlets.
4. Ensuring you identify which port each cable connects to, remove all cables from the back of the enclosure that has faulty support rails.
5. Remove the end caps from the front flanges of the enclosure by following the removal instructions in topic “Replacing Storwize V7000 Gen2 enclosure end caps” on page 329.
6. Unscrew the M5 screw from the left flange.
Repeat with the M5 screw in the right flange.
7. Slide the enclosure from the rack.

CAUTION:



The weight of this part or unit is between 32 and 55 kg (70.5 and 121.2 lb). It takes three persons to safely lift this part or unit. (C010)

8. Locate the left support rail.
Record the shelf number of the support rail so that the replacement rails can be installed into the same position.
9. At the rear of the rack, remove the securing M5 screw from the bottom hole of the rear bracket of the rail, then open the rear hinge bracket (Figure 111 on page 365).

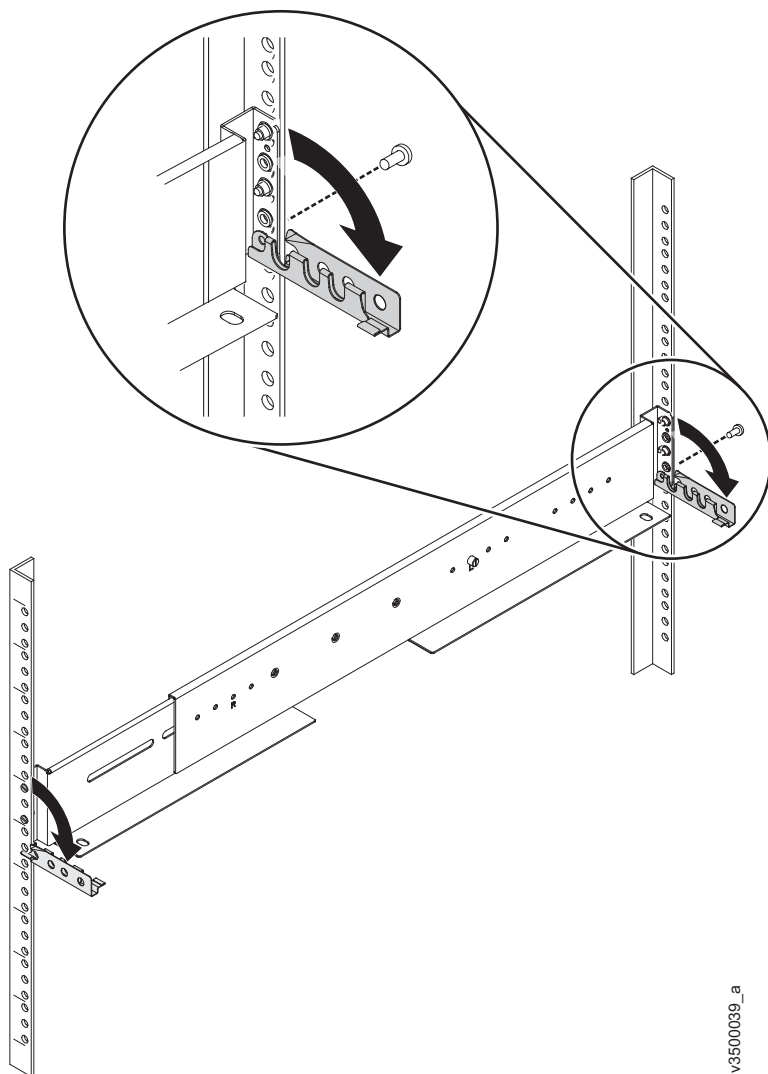


Figure 111. Opening rear hinge bracket of mounting rail

10. At the front of the rack, hold onto the rail and open the front hinge bracket.
11. Compress the rail against its spring to shorten it, then remove it from inside the rack (Figure 112 on page 366).

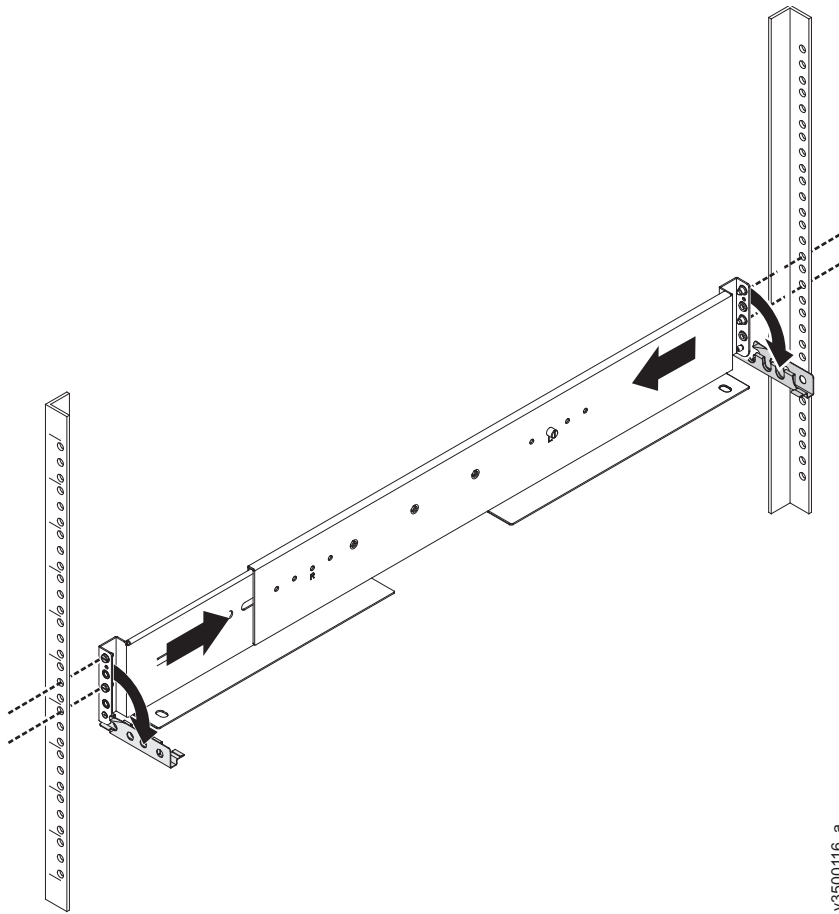


Figure 112. Compressing rail for removal from rack

12. Repeat steps 9 on page 364 to 11 on page 365 on the right support rail.
13. Install the new support rails at the rack position that is recorded at step 8 on page 364 by following the instructions in Step 6. Installing the support rails for the enclosures.
14. Reinstall the enclosure (removed at step 7 on page 364) and the end caps (removed at step 5 on page 364) by following the instructions in Step 7. Installing the enclosures.
15. If components were removed from the enclosure at step 7 on page 364, return each canister, drive assembly, and power supply unit to its labeled slot.
16. Reconnect the cables, ensuring that they are connected to their original ports.
17. Reconnect the power supply cables to their original power supply and electrical outlet.
The system starts.
18. After the system is online, use the management GUI to verify that the system is correct.

Replacing the Storwize V7000 Gen2 expansion enclosure support rails

You can replace faulty support rails with new ones received from CRU / FRU stock.

Before you begin

Two persons are required at step 7

Procedure

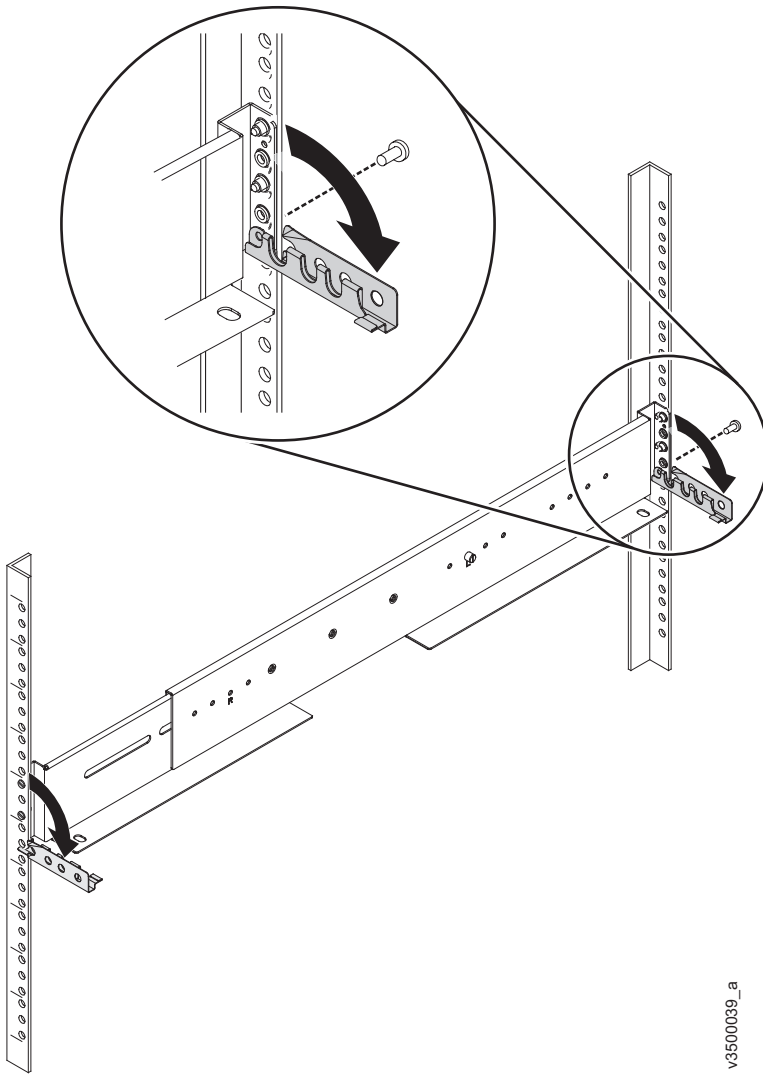
To replace the support rails, complete the following steps.

1. Identify the enclosure mounted on the rails being replaced.
Follow the steps in “Procedure: Identifying which Storwize V7000 Gen2 enclosure or canister to service” on page 252 to ensure that you identify the correct enclosure.
2. Shut down the system by following the steps in “Procedure: Powering off your Storwize V7000 Gen2 system” on page 280.
3. Remove power from the enclosure by unplugging both power cables from the electrical outlets.
4. Ensuring you identify which port each cable connects to, remove all cables from the back of the enclosure that has faulty support rails.
5. Remove the end caps from the front flanges of the enclosure by following the removal instructions in topic “Replacing Storwize V7000 Gen2 enclosure end caps” on page 329.
6. Unscrew the M5 screw from the left flange.
Repeat with the M5 screw in the right flange.
7. Slide the enclosure from the rack.

CAUTION:

The weight of this part or unit is between 18 and 32 kg (39.7 and 70.5 lb). It takes two persons to safely lift this part or unit. (C009)

8. Locate the left support rail.
Record the shelf number of the support rail so that the replacement rails can be installed into the same position.
9. At the rear of the rack, remove the securing M5 screw from the bottom hole of the rear bracket of the rail, then open the rear hinge bracket (Figure 113 on page 368).



v3500039_a

Figure 113. Opening rear hinge bracket of mounting rail

10. At the front of the rack, hold onto the rail and open the front hinge bracket.
11. Compress the rail against its spring to shorten it, then remove it from inside the rack (Figure 114 on page 369).

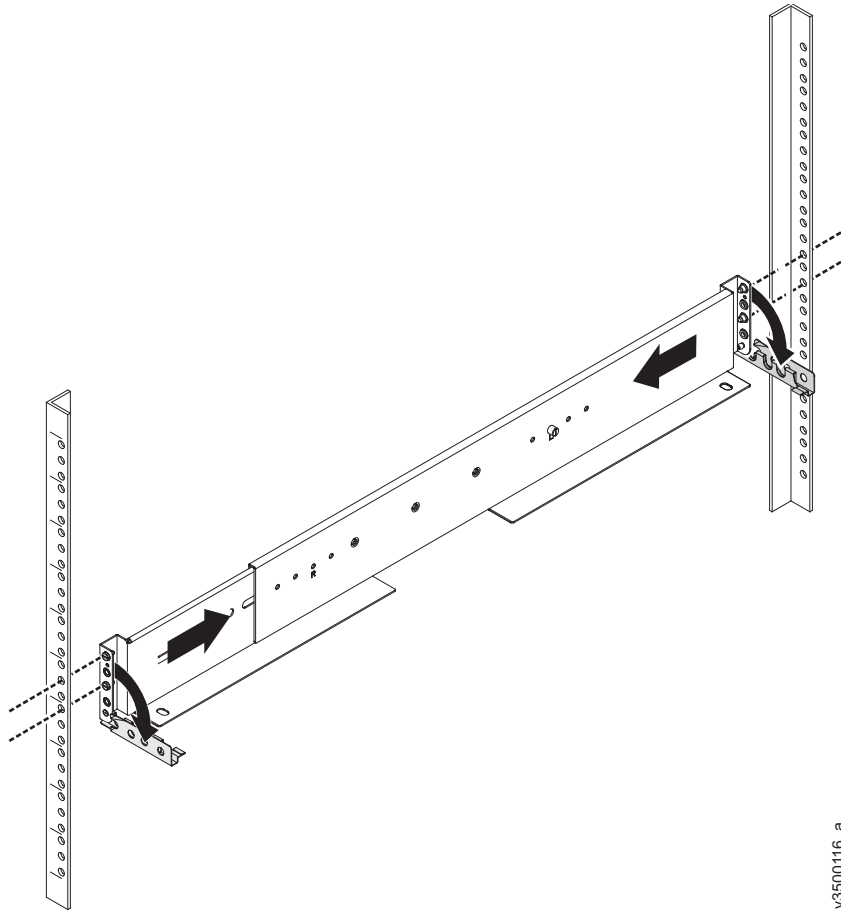


Figure 114. Compressing rail for removal from rack

12. Repeat steps 9 on page 367 to Figure 114 on the right support rail.
13. Install the new support rails at the rack position that is recorded at step 8 on page 367 by following the instructions in Step 6. Installing the support rails for the enclosures.
14. Reinstall the enclosure (removed at step 7 on page 367) and the end caps (removed at step 5 on page 367) by following the instructions in Step 7. Installing the enclosures.
15. If components were removed from the enclosure at step 7 on page 367, return each canister, drive assembly, and power supply unit to its labeled slot.
16. Reconnect the cables, ensuring that they are connected to their original ports.
17. Reconnect the power supply cables to their original power supply and electrical outlet.
The system starts.
18. After the system is online, use the management GUI to verify that the system is correct.

Replacing the Storwize V7000 Gen1 support rails

You can replace the support rails.

Procedure

To replace the support rails, complete the following steps:

1. Remove the enclosure.
2. Record the location of the rail assembly in the rack cabinet.
3. Working from the back of the rack cabinet, remove the clamping screw **1** from the rail assembly on both sides of the rack cabinet.

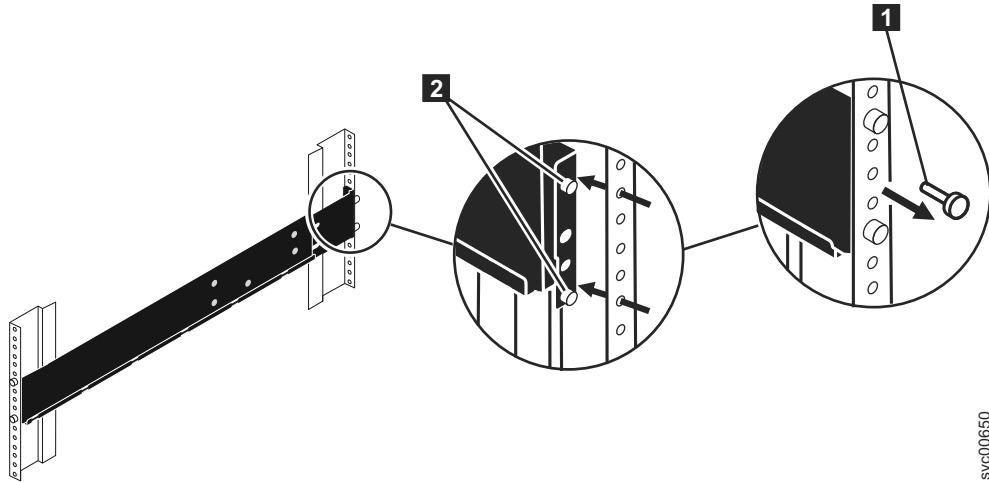


Figure 115. Removing a rail assembly from a rack cabinet

4. Working from the front of the rack cabinet, remove the clamping screw from the rail assembly on both sides of the rack cabinet.
5. From one side of the rack cabinet, grip the rail and slide the rail pieces together to shorten the rail.
6. Disengage the rail location pins **2**.
7. From the other side the rack cabinet, grip the rail and slide the rail pieces together to shorten the rail.
8. Disengage the rail location pins **2**.
9. Starting from the location of the previous rail assembly, align the bottom of the rail with the bottom of the two rack units. Insert the rail location pins through the holes in the rack cabinet.
10. Insert a clamping screw into the upper mounting hole between the rail location pins.
11. Tighten the screw to secure the rail to the rack.
12. Working from the rear of the rack cabinet, extend the rail that you secured to the front to align the bottom of the rail with the bottom of the two rack units.

Note: Ensure that the rail is level between the front and the back.

13. Insert the rail location pins through the holes in the rack cabinet.
14. Insert a clamping screw into the upper mounting hole between the rail location pins.
15. Tighten the screw to secure the rail to the rack from the back side.
16. Repeat the steps to secure the opposite rail to the rack cabinet.

Replacing node canister memory modules

Remove and replace node canister memory modules.

Replacing a Storwize V7000 Gen2 node canister memory module (16 GB DIMM)

You can replace a faulty node canister memory module (16 GB DIMM) with a new one received from CRU / FRU stock.

Procedure

1. Follow “Procedure: Removing a Storwize V7000 Gen2 node canister” on page 279 to disconnect and remove the node canister with the faulty memory.
2. Remove the lid of the canister, as described in “Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister” on page 286.
3. Locate the DIMM slot with the faulty DIMM. Slot 1 is next to the battery area. Slot 2 is next to the processor. The slots are marked 1, 2, 4, 3 as shown in Figure 116.
4. Remove the faulty DIMM by applying gentle, outwards pressure simultaneously to the retaining clips at each end of the DIMM slot until the DIMM is levered out of the slot.
5. Touch the replacement DIMM packaging onto a metal area of the case, then remove the replacement DIMM from its package.
6. Ensure that the retaining clips of the DIMM slot are open.
7. Gently place the DIMM in the slot, ensuring that the notches in the DIMM align with the shape of the slot, as shown in Figure 116.

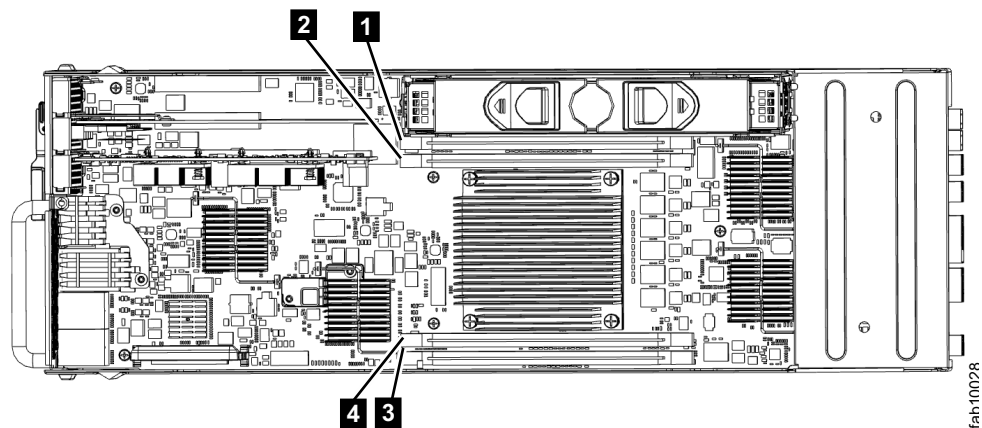


Figure 116. Installing a Storwize V7000 2076-524 node canister memory module

8. Apply even, firm, downwards pressure on the DIMM in its slot until the retaining clips move inwards and engage the edges of the DIMM.
9. Ensure that the retaining clips are fully engaged with the edges of the DIMM. Gently pull the DIMM upwards and ensure that it does not become dislodged.
10. Replace the canister lid, as described in “Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister” on page 286.
11. Reinstall the canister, as described in “Replacing a Storwize V7000 Gen2 node canister” on page 296, into the enclosure from which it was removed in step 1. The node canister starts.
12. Reconnect the cables to the canister, ensuring cables go into the same ports from which they were removed in step 1.
13. When the canister is back online, check the event log for new events, particularly events that relate to hardware changes.

Replacing a host interface adapter

Remove and replace a host interface adapter.

Replacing a Storwize V7000 Gen2 host interface adapter

To replace a faulty host interface adapter in a Storwize V7000 2076-524 with a new one received from customer replaceable unit (CRU) or field replaceable unit (FRU) stock, use this procedure.

About this task

For lists of supported host interface adapters, refer to “Storwize V7000 2076-524 Gen2 replaceable units” on page 288.

Important: For correct operation, use the correct SFP transceivers with each adapter card. The topic “Storwize V7000 2076-524 Gen2 replaceable units” identifies the suitable IBM parts.

- Use only 8G bps SFP transceivers in the 8 Gbps Fibre Channel adapter cards.
- Use only 16 Gbps SFP transceivers in the 16 Gbps Fibre Channel adapter cards.
- Use only 10 Gbps SFP transceivers in the 10 Gbps Ethernet (FCoE/iSCSI) adapter card.

Procedure

Complete the following steps to replace a host interface adapter.

1. Complete “Procedure: Removing a Storwize V7000 Gen2 node canister” on page 279 to remove the Storwize V7000 2076-524 node canister with the faulty host interface adapter.
2. Identify which host interface adapter is to be removed. The interface adapters are in slots numbered 2 and 3
3. Remove any small form-factor pluggable SFP transceiver from each rear-facing port of the host interface adapter and put safely to one side.
4. Complete “Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister” on page 286 to remove and replace the lid of a Storwize V7000 2076-524 node canister.
5. Gently pull the host interface adapter upward to disconnect it **2**, and then carefully remove it from the canister **1**. Figure 117 on page 373 displays removing the host interface adapter.

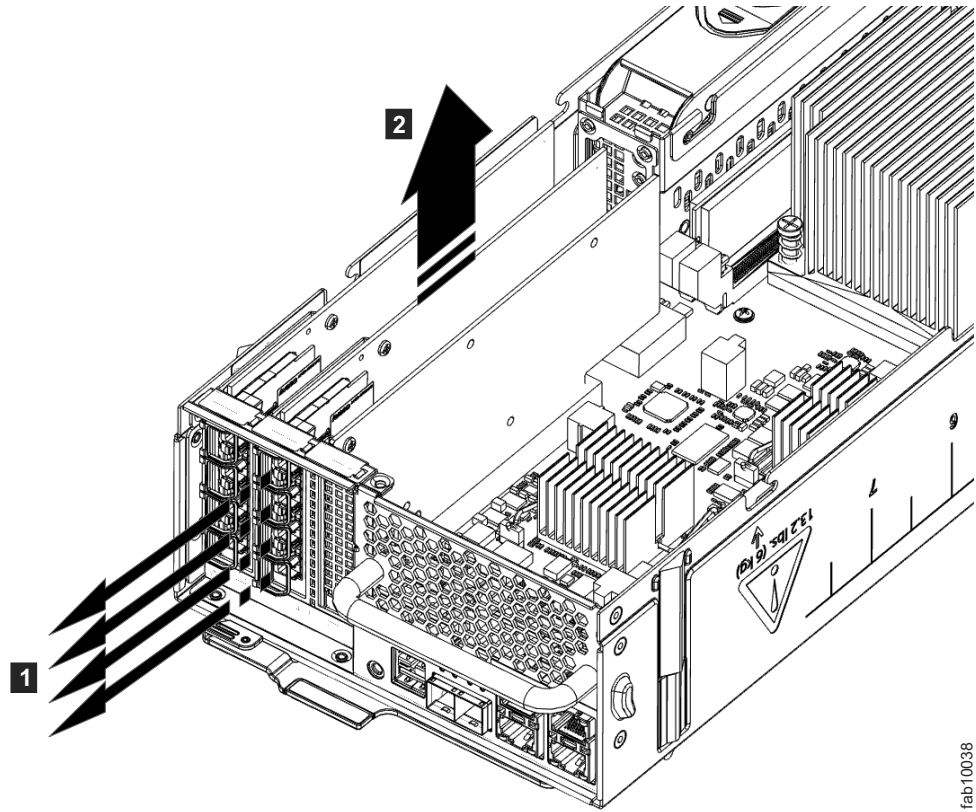


Figure 117. Removing the host interface adapter

6. Remove the replacement host interface adapter **1** from its package. Figure 118 on page 374 displays installing the host interface adapter.
7. Set the connecting edge of the replacement host interface adapter **3** on the host interface adapter connector so that the connectors are aligned.
8. Ensure that the adapter is perpendicular to the canister main board so that the small tab on the top of the bracket **2** is aligned with the alignment hole in the top edge of the slot.
9. Maintain alignment while applying pressure to the top edge of the host interface adapter opposite the connecting edge to push the host interface adapter into the connector **4** and **5**.

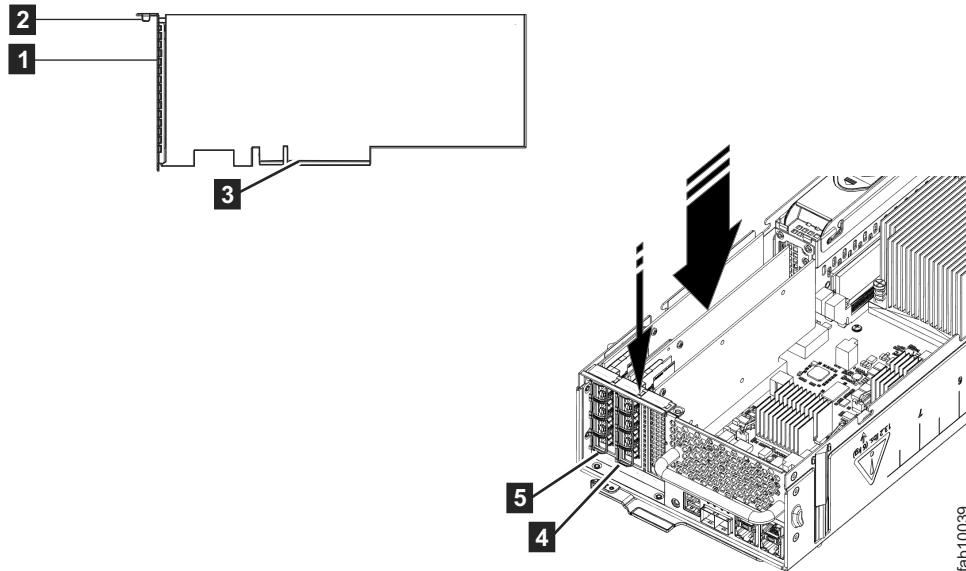


Figure 118. Installing the host interface adapter

10. Check that the host interface adapter is installed squarely in its slot. If the small tab of the mounting bracket is not positioned correctly, repeat steps 5 on page 372 onward to install the adapter correctly.
11. Replace the canister lid, as described in “Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister” on page 286.
12. If any SFP transceivers were removed from the rear-facing ports of the host interface adapter at step 2 on page 372, ensure each one is reinstalled by following the installation steps of “Replacing an SFP transceiver in a Storwize V7000 2076-524 control enclosure” on page 303.
13. Reinstall the canister into the enclosure from which it was removed in step 1 on page 372 following “Replacing a Storwize V7000 Gen2 node canister” on page 296. The node canister starts.
14. Reconnect the cables to the canister, ensuring cables go into the same ports from which they were removed in step 1 on page 372.
15. When the canister is back online, check the event log for any new events relating to hardware changes.

Replacing Storwize V7000 Gen2 host interface adapters in two control enclosures concurrently

It is possible to reconfigure one node canister of each control enclosure at the same time. During the procedure, both I/O groups (control enclosures) are online with no redundancy, but the total maintenance period is reduced.

To replace host interface adapters in both control enclosures concurrently, use the procedure for replacing the host interface adapter in a single enclosure, but complete each step in both enclosures before continuing to the next step. Table 116 on page 375 shows how to sequence the step in each node. Work your way down the table, completing each row before starting the next row.

For the procedure for replacing the host interface adapter in a single enclosure, refer to the “Replacing a Storwize V7000 Gen2 host interface adapter” on page 372.

Table 116. Replacing host interface adapters in two control enclosures concurrently

Control enclosure 1		Control enclosure 2	
Node canister 1	Node canister 2	Node canister 1	Node canister 2
Step 1		Step 1	
Step 2		Step 2	
...		...	
Final step		Final step	
	Step 1		Step 1
	Step 2		Step 2

	Final step		Final step

Replacing a CMOS battery

Remove and replace the complementary metal-oxide semiconductor (CMOS) battery.

Replacing a Storwize V7000 Gen2 CMOS battery

The complementary metal-oxide semiconductor (CMOS) battery is a coin-shaped power cell that is mounted inside a node canister. It is used to keep the system time when there is no power to the node canister. The lithium battery must be handled correctly to avoid possible danger. If you replace the battery, you must adhere to all safety instructions.

About this task

Use this procedure to replace a CMOS battery. Dispose of the faulty battery properly.

CAUTION: If your system has a module containing a lithium battery, replace it only with the same module type made by the same manufacturer. The battery contains lithium and can explode if not properly used, handled, or disposed of.

Do not:

- Throw or immerse into water
- Heat to more than 100°C (212°F)
- Repair or disassemble

Dispose of the battery as required by local ordinances or regulations. (C045)

Procedure

1. Complete “Procedure: Removing a Storwize V7000 Gen2 node canister” on page 279
2. Open the canister and remove the lid as described in “Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister” on page 286.
3. Locate the CMOS battery inside the node canister.. See Figure 119 on page 376

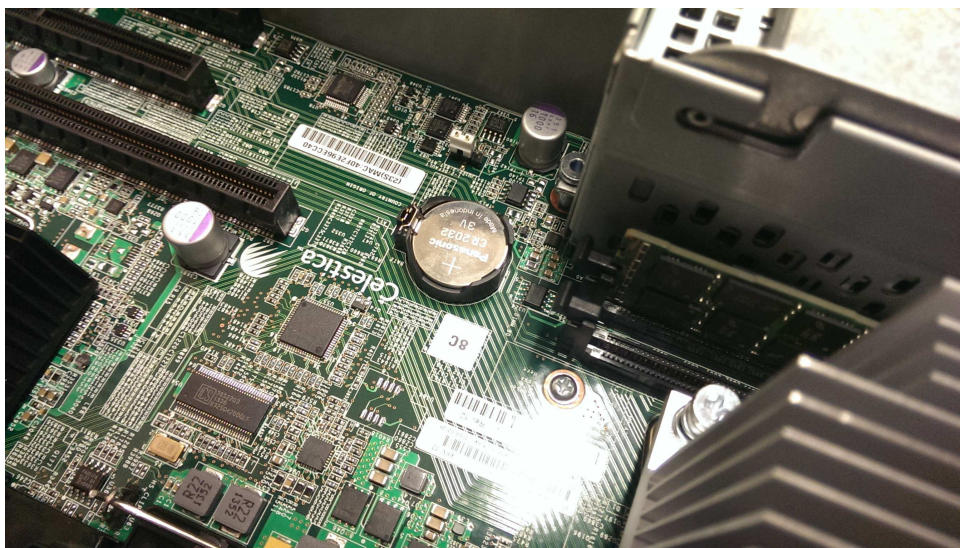


Figure 119. Replacing a CMOS Gen2 battery

4. Push the coin cell latch to the side to release the coin cell from its holder, then remove the expired coin cell.
5. Orient the replacement coin cell with the flat side upwards and place it down onto the coin cell holder.
6. Gently push the coin cell down into the holder so that it clicks under the latch and sits parallel with the canister main board.
7. Replace the canister lid as described in “Procedure: Removing and replacing the lid of a Storwize V7000 Gen2 node canister” on page 286.
8. Reinsert the canister into the slot from which it came.
9. Reconnect all cables.
10. Open the management GUI.
11. Use the management GUI to check that the time and date settings of the system are correct.
12. In the event log view, if a CMOS battery error is present, run the fix procedure.

General storage system procedures

This section provides general information about hardware and Fibre Channel link issues.

Procedure: SAN problem determination

About this task

SAN failures might cause Storwize V7000 Unified volumes to be inaccessible to host systems. Failures can be caused by SAN configuration changes or by hardware failures in SAN components.

The following list identifies some of the hardware that might cause failures:

- Power, fan, or cooling
- Application-specific integrated circuits
- Installed small form-factor pluggable (SFP) transceiver

- Fiber-optic cables

If error codes sent you here, complete the following steps:

Procedure

1. Verify that the power is turned on to all switches and storage controllers that the Storwize V7000 Unified system uses, and that they are not reporting any hardware failures. If problems are found, resolve those problems before you proceed further.
2. Verify that the Fibre Channel cables that connect the systems to the switches are securely connected.
3. If you have a SAN management tool, use that tool to view the SAN topology and isolate the failing component.

iSCSI performance analysis and tuning

This procedure provides a solution for Internet Small Computer Systems Interface (iSCSI) host performance problems while connected to a Storwize V7000 Unified system and its connectivity to the network switch.

About this task

Some of the attributes and host parameters that might affect iSCSI performance:

- Transmission Control Protocol (TCP) Delayed ACK
- Ethernet jumbo frame
- Network bottleneck or oversubscription
- iSCSI session login balance
- Priority flow control (PFC) setting and bandwidth allocation for iSCSI on the network

Procedure

1. Disable the TCP delayed acknowledgment feature.

To disable this feature, refer to OS/platform documentation.

- VMWare: <http://kb.vmware.com/selfservice/microsites/microsite.do>
- Windows: <http://support.microsoft.com/kb/823764>

The primary signature of this issue: read performance is significantly lower than write performance. Transmission Control Protocol (TCP) delayed acknowledgment is a technique that is used by some implementations of the TCP in an effort to improve network performance. However, in this scenario where the number of outstanding I/O is 1, the technique can significantly reduce I/O performance.

In essence, several ACK responses can be combined together into a single response, reducing protocol overhead. As described in RFC 1122, a host can delay sending an ACK response by up to 500 ms. Additionally, with a stream of full-sized incoming segments, ACK responses must be sent for every second segment.

Important: The host must be rebooted for these settings to take effect. A few platforms (for example, standard Linux distributions) do not provide a way to disable this feature. However, the issue was resolved with the version 7.1 release, and no host configuration changes are required to manage **TcpDelayedAck** behavior.

2. Enable jumbo frame for iSCSI.

Jumbo frames are Ethernet frames with a size in excess of 1500 bytes. The maximum transmission unit (MTU) parameter is used to measure the size of jumbo frames.

The Storwize V7000 Unified supports 9000-bytes MTU. Refer to the CLI command **cfgportip** to enable jumbo frame. This command is disruptive as the link flips and the I/O operation through that port pauses.

The network must support jumbo frames end-to-end for this to be effective; verify this by sending a ping packet to be delivered without fragmentation. For example:

- Windows:

```
ping -t <iscsi target ip> -S <iscsi initiator ip> -f -l <new mtu size - packet overhead (usually 36, might differ)>
```

The following command is an example of a command that is used to check whether a 9000-bytes MTU is set correctly on a Windows 7 system:

```
ping -t -S 192.168.1.117 192.168.1.217 -f -l 8964
```

The following output is an example of a successful reply:

```
192.168.1.217: bytes=8964 time=1ms TTL=52
```

- Linux:

```
ping -l <source iscsi initiator ip> -s <new mtu size> -M do <iscsi target ip>
```

- ESXi:

```
ping <iscsi target ip> -I <source iscsi initiator ip> -s <new mtu size - 28> -d
```

3. Verify the switch's port statistic where initiator/target ports are connected to make sure that packet drops are not high.

Review network architecture to avoid any bottlenecks and oversubscription.

The network needs to be balanced to avoid any packet drop; packet drop significantly reduces storage performance. Involve networking support to fix any such issues.

4. Optimize and utilize all iSCSI ports.

To optimize Storwize V7000 Unified resource utilization, all iSCSI ports must be used.

- Each port is assigned to one CPU, and by balancing the login, one can maximize CPU utilization and achieve better performance. Ideally, configure subnets equal to the number of iSCSI ports on the Storwize V7000 Unified node. Configure each port of a node with an IP on a different subnet and keep it the same for other nodes. The following example displays an ideal configuration:

```
Node 1
Port 1: 192.168.1.11
Port 2: 192.168.2.21
Port 3: 192.168.3.31
```

```
Node 2:
Port 1: 192.168.1.12
Port 2: 192.168.2.22
Port 3: 192.168.3.33
```

- Avoid situations where 50 hosts are logged in to port 1 and only five hosts are logged in to port 2.
- Use proper subnetting to achieve a balance between the number of sessions and redundancy.

5. Troubleshoot problems with PFC settings.

You do not need to enable PFC on the Storwize V7000 Unified system. Storwize V7000 Unified reads the data center bridging exchange (DCBx) packet and enables PFC for iSCSI automatically if it is enabled on the switch. In the **lspportip** command output, the fields `lossless_iscsi` and `lossless_iscsi6` show [on/off] depending on whether PFC is enabled or not for iSCSI on the system.

If the fields `lossless_iscsi` and `lossless_iscsi6` are showing off, it might be due to one of the following reasons:

- a. VLAN is not set for that IP. Verify the following checks:
 - For IP address type IPv4, check the `vlan` field in the **lspportip** output. It should not be blank.
 - For IP address type IPv6, check the `vlan_6` field in the **lspportip** output. It should not be blank.
 - If the `vlan` and `vlan_6` fields are blank, set the VLAN for the IP type using **Configuring VLAN for iSCSI**.
- b. Host flag is not set for that IP. Verify the following checks:
 - For IP address type IPv4, check the `host` field in the **lspportip** output. It should be yes.
 - For IP address type IPv6, check the `host_6` field in the **lspportip** output. It should be yes.
 - If the `host` and `host_6` fields are not yes, set the host flag for the IP type using the **cfgportip** CLI command.
- c. PFC is not properly set on the switch.

If the VLAN is properly set, and the host flag is also set, but the `lossless_iscsi` or `lossless_iscsi6` field is still showing off, some switch settings might be missing or incorrect.

Verify the following settings in the switch:

- Priority tag is set for iSCSI traffic.
- PFC is enabled for priority tag that is assigned to iSCSI CoS.
- DCBx is enabled on the switch.

Also check the appropriate documentation:

- Consult the documentation for enabling PFC on your specific switch.
- Consult the documentation for enabling PFC on Red Hat Enterprise Linux (RHEL) and Windows hosts specific to your configuration.

6. Ensure that proper bandwidth is given to iSCSI on the network.

You can divide the bandwidth among the various types of traffic. It is important to assign proper bandwidth for good performance. To assign bandwidth for iSCSI traffic, you need to first enable the priority flow control for iSCSI.

Fibre Channel link failures

When a failure occurs on a single Fibre Channel link, the small form-factor pluggable (SFP) transceiver might need to be replaced.

Before you begin

The following items can indicate that a single Fibre Channel link failed:

- The Fibre Channel status LEDs at the rear of the node canister
- An error that indicates a single port failed

Procedure

Attempt each of these actions, in the following order, until the failure is fixed.

1. Ensure that the Fibre Channel cable is securely connected at each end.
2. Replace the Fibre Channel cable.
3. Replace the SFP transceiver for the failing port on the node.

Note: Storwize V7000 Unified nodes are supported by both longwave SFP transceivers and shortwave SFP transceivers. You must replace an SFP transceiver with the same type of SFP transceiver. If the SFP transceiver to replace is a longwave SFP transceiver, for example, you must provide a suitable replacement. Removing the wrong SFP transceiver might result in loss of data access.

4. Contact the IBM Support Center for assistance in replacing the node canister.

Ethernet iSCSI host-link problems

If you are having problems attaching to the Ethernet hosts, your problem might be related to the network, the Storwize V7000 Unified system, or the host.

Note: Storwize V7000 Unified and Host IP should be on the same VLAN. Host and Storwize V7000 Unified nodes should not have same subnet on different VLANs.

For network problems, you can attempt any of the following actions:

- Test your connectivity between the host and Storwize V7000 Unified ports.
- Try to ping the Storwize V7000 Unified system from the host.
- Ask the Ethernet network administrator to check the firewall and router settings.
- Check that the subnet mask and gateway are correct for the Storwize V7000 Unified host configuration.

Using the management GUI for Storwize V7000 Unified problems, you can attempt any of the following actions:

- View the configured node port IP addresses.
- View the list of volumes that are mapped to a host to ensure that the volume host mappings are correct.
- Verify that the volume is online.

For host problems, you can attempt any of the following actions:

- Verify that the host iSCSI qualified name (IQN) is correctly configured.
- Use operating system utilities (such as Windows device manager) to verify that the device driver is installed, loaded, and operating correctly.
- If you configured the VLAN, check that its settings are correct. Ensure that Host Ethernet port, Storwize V7000 Unified Ethernet ports IP address, and Switch port are on the same VLAN ID. Ensure that on each VLAN, a different subnet is used. Configuring the same subnet on different VLAN IDs can cause network connectivity problems.

Recover system procedure

The recover system procedure recovers the entire storage system if the system state is lost from all control enclosure node canisters. The procedure re-creates the storage system by using saved configuration data. The saved configuration data is in the active quorum disk and the latest XML configuration backup file. The recovery might not be able to restore all volume data. This procedure is also known as Tier 3 (T3) recovery.

CAUTION:

If the system encounters a state where:

- **No nodes are active, and**
- **One or more nodes have node errors that require a node rescue, node canister replacement, or node firmware reinstallation**

STOP and contact IBM Remote Technical Support. Initiating this T3 recover system procedure while in this specific state can result in loss of the XML backup of the block volume storage configuration.

Before you recover the storage system, shut down the file modules:

- From a workstation with access to the management subnet, log on to the management CLI as an administrator. For example, the default admin password is admin0001.
- `ssh admin@<management IP>`
- `stopcluster`

After you complete the following storage system recovery procedure, refer to Turning on the system, located in the Information Center, to power the file modules back on.

Contact IBM Remote Technical support if the health indicator in the management GUI does not turn back to green within 30 minutes. They can assist you with recovering the file modules so that access to the file systems can be restored.

After you complete the storage system recovery procedure, contact IBM support. They can assist you with recovering the file modules so that access to the file systems can be restored.

Attention:

- Run service actions only when directed by the fix procedures. If used inappropriately, service actions can cause loss of access to data or even data loss. Before you attempt to recover a storage system, investigate the cause of the failure and attempt to resolve those issues by using other fix procedures. Read and understand all of the instructions before you complete any action.
- The recovery procedure can take several hours if the system uses large-capacity devices as quorum devices.

Do not attempt the recover system procedure unless the following conditions are met:

- All of the conditions have been met in “When to run the recover system procedure” on page 383.
- All hardware errors are fixed. See “Fix hardware errors” on page 383
- All node canisters have candidate status. Otherwise, see step 1.
- All node canisters must be at the same level of code that the storage system had before the system failure. If any node canisters were modified or replaced, use the service assistant to verify the levels of code, and where necessary, to reinstall the level of code so that it matches the level that is running on the other node canisters in the system.
- If the system is encrypted, insert the USB flash drive that contains the encryption key file into the node canister that will be running system recovery.
- System recovery will fail if the system has encryption enabled and a USB flash device containing the key is not found. Insert the USB flash drive that contains the encryption key file into the node canister that is running the system recovery, and then retry system recovery.

The system recovery procedure is one of several tasks that must be completed. The following list is an overview of the tasks and the order in which they must be completed:

1. Preparing for system recovery
 - a. Review the information regarding when to run the recover system procedure.
 - b. Fix your hardware errors and make sure that all nodes in the system are shown in service assistant or in the output from **sainfo lsservicenodes**.
 - c. Remove the system information for node canisters with error code 550 or error code 578 by using the service assistant, but only if the recommended user response for these node errors has already been followed. See “Removing system information for node canisters with error code 550 or error code 578 using the service assistant” on page 385.
 - d. For Virtual Volumes (VVols), shut down the services for any instances of Spectrum Control Base that are connecting to the system. Use the Spectrum Control Base command **service ibm_spectrum_control stop**.
2. Running the system recovery. After you prepared the system for recovery and met all the pre-conditions, run the system recovery.

Note: Run the procedure on one system in a fabric at a time. Do not run the procedure on different node canisters in the same system. This restriction also applies to remote systems.

3. Completing actions to get your environment operational.
 - Recovering from offline volumes by using the CLI.

- Checking your system, for example, to ensure that all mapped volumes can access the host.

When to run the recover system procedure

Attempt a recover procedure only after a complete and thorough investigation of the cause of the system failure. Attempt to resolve those issues by using other service procedures.

Attention: If you experience failures at any time while running the recover system procedure, call the IBM Support Center. Do not attempt to do further recovery actions, because these actions might prevent support from restoring the system to an operational status.

Certain conditions must be met before you run the recovery procedure. Use the following items to help you determine when to run the recovery procedure:

1. Check that no node in the system is active and that the management IP is not accessible. If any node has active status, it is not necessary to recover the system.
2. Resolve all hardware errors in nodes so that only node errors 578 or 550 are present. If this is not the case, go to “Fix hardware errors.”
3. Ensure all backend storage that is administered by the system is present before you run the recover system procedure.
4. If any nodes have been replaced, ensure that the WWNN of the replacement node matches that of the replaced node, and that no prior system data remains on this node.

Fix hardware errors

Before running a system recovery procedure, it is important to identify and fix the root cause of the hardware issues.

Identifying and fixing the root cause can help recover a system, if these are the faults that are causing the system to fail. The following are common issues which can be easily resolved:

- The node has been powered off or the power cords were unplugged.
- Check the node status of every node canister that is part of this system. Resolve all hardware errors except node error 578 or node error 550.
 - All nodes must be reporting either a node error 578 or a node error 550. These error codes indicate that the system has lost its configuration data. If any nodes report anything other than these error codes, do not perform a recovery. You can encounter situations where non-configuration nodes report other node errors, such as a 550 node error. The 550 error can also indicate that a node is not able to join a system.
 - If any nodes show a node error 550, record the error data that is associated with the 550 error from the service assistant.
 - In addition to the node error 550, the report can show data that is separated by spaces in one of the following forms:
 - Node identifiers in the format: `<enclosure_serial>-<canister slot ID>` (7 characters, hyphen, 1 number), for example, 01234A6-2
 - Quorum drive identifiers in the format: `<enclosure_serial>:<drive slot ID>[<drive 11S serial number>]` (7 characters, colon, 1 or 2 numbers, open square bracket, 22 characters, close square bracket), for example, 01234A9:21[11S1234567890123456789]

- Quorum MDisk identifier in the format: *WWPN/LUN* (16 hexadecimal digits followed by a forward slash and a decimal number), for example, 1234567890123456/12
- If the error data contains a node identifier, ensure that the node that is referred to by the ID is showing node error 578. If the node is showing a node error 550, ensure that the two nodes can communicate with each other. Verify the SAN connectivity, and if the 550 error is still present, restart one of the two nodes from the service assistant by clicking **Restart Node**.
- If the error data contains a quorum drive identifier, locate the enclosure with the reported serial number. Verify that the enclosure is powered on and that the drive in the reported slot is powered on and functioning. If the node canister that is reporting the fault is in the I/O group of the listed enclosure, ensure that it has SAS connectivity to the listed enclosure. If the node canister that is reporting the fault is in a different I/O group from the listed enclosure, ensure that the listed enclosure has SAS connectivity to both node canisters in the control enclosure in its I/O group. After verification, restart the node by clicking **Restart Node** from the service assistant.
- If the error data contains a quorum MDisk identifier, verify the SAN connectivity between this node and that WWPN. Check the storage controller to ensure that the LUN referred to is online. After verification, if the 550 error is still present, restart the node from the service assistant by clicking **Restart Node**.
- If there is no error data, the error is because there are insufficient connections between nodes over the Fibre Channel network. Each node must have at least two independent Fibre Channel logical connections, or logins, to every node that is not in the same enclosure. An independent connection is one where both physical ports are different. In this case, there is a connection between the nodes, but there is not a redundant connection. If there is no error data, wait 3 minutes for the SAN to initialize. Next, verify:
 - There are at least two Fibre Channel ports that are operational and connected on every node.
 - The SAN zoning allows every port to connect to every port on every other node
 - All redundant SANs (if used) are operational.

After verification, if the 550 error is still present, restart the node from the service assistant by clicking **Restart Node**.

Note: If after resolving all these scenarios, half or greater than half of the nodes are reporting node error 578, it is appropriate to run the recovery procedure. Call the IBM Support Center for further assistance.

- For any nodes that are reporting a node error 550, ensure that all the missing hardware that is identified by these errors is powered on and connected without faults.
- If you have not been able to restart the system, and if any node other than the current node is reporting node error 550 or 578, you must remove system data from those nodes. This action acknowledges the data loss and puts the nodes into the required candidate state.

Removing system information for node canisters with error code 550 or error code 578 using the service assistant

The system recovery procedure works only when all node canisters are in candidate status. Ensure that the service assistant displays all of the node canisters with the 550 error code. The 550 error code is the expected node error when more than half of the nodes in the system are missing or when the active quorum disk cannot be found. If the service assistant displays any node canisters with error codes 550 or 578 and all the recommended actions have been completed on these nodes, you must remove their system data.

About this task

Before performing this task, ensure that you have read the introductory information in the overall recover system procedure.

To remove system information from a node canister with an error 550 or 578, follow this procedure using the service assistant:

Procedure

1. Point your browser to the service IP address of one of the nodes, for example, https://node_service_ip_address/service/.
2. Log on to the service assistant.
3. Select **Manage System**.
4. Click **Remove System Data**.
5. Confirm that you want to remove the system data when prompted.
6. Remove the system data for the other nodes that display a 550 or a 578 error.
All nodes previously in this system must have a node status of Candidate and have no errors listed against them.
7. Resolve any hardware errors until the error condition for all nodes in the system is **None**.
8. Ensure that all nodes in the system display a status of candidate.

Results

When all nodes display a status of candidate and all error conditions are **None**, you can run the recovery procedure.

Running system recovery using the service assistant

Start recovery when all node canisters that were members of the system are online and have candidate status. Use the service assistant to verify the status. For any nodes that display error code 550 or 578, ensure that all nodes in the system are visible and all the recommended actions have been completed before placing them into candidate status. To place a node into candidate status, remove system information for that node canister. Do not run the recovery procedure on different node canisters in the same system.

Before you begin

Note: Ensure that the web browser is not blocking pop-up windows. If it does, progress windows cannot open.

Before you begin this procedure, read the recover system procedure introductory information; see “Recover system procedure” on page 381.

About this task

Attention: This service action has serious implications if not completed properly. If at any time an error is encountered not covered by this procedure, stop and call the support center.

Run the recovery from any node canisters in the system; the node canisters must not have participated in any other system.

Note: Each individual stage of the recovery procedure can take significant time to complete, depending on the specific configuration.

Procedure

1. Point your browser to the service IP address of one of the node canisters.
If the IP address is unknown or is not configured, assign an IP address using the initialization tool; see “Procedure: Changing the service IP address of a node canister” on page 271.
2. Log on to the service assistant.
3. Check that all node canisters that were members of the system are online and have candidate status.
If any nodes display error code 550 or 578, remove their system data to place them into candidate status; see “Procedure: Removing system data from a node canister” on page 270.
4. Select **Recover System** from the navigation.
5. Follow the online instructions to complete the recovery procedure.
 - a. Verify the date and time of the last quorum time. The time stamp must be less than 30 minutes before the failure. The time stamp format is *YYYYMMDD hh:mm*, where *YYYY* is the year, *MM* is the month, *DD* is the day, *hh* is the hour, and *mm* is the minute.
Attention: If the time stamp is not less than 30 minutes before the failure, call the support center.
 - b. Verify the date and time of the last backup date. The time stamp must be less than 24 hours before the failure. The time stamp format is *YYYYMMDD hh:mm*, where *YYYY* is the year, *MM* is the month, *DD* is the day, *hh* is the hour, and *mm* is the minute.
Attention: If the time stamp is not less than 24 hours before the failure, call the support center.
Changes that are made after the time of this backup date might not be restored.

Results

Any one of the following categories of messages might be displayed:

- T3 successful

The volumes are back online. Use the final checks to get your environment operational again.

- T3 recovery completed with errors

T3 recovery completed with errors: One or more of the volumes are offline because there was fast write data in the cache. To bring the volumes online, see “Recovering from offline volumes using the CLI” for details.

- T3 failed

Call the support center. Do not attempt any further action.

Verify that the environment is operational by completing the checks that are provided in “What to check after running the system recovery” on page 388.

If any errors are logged in the error log after the system recovery procedure completes, use the fix procedures to resolve these errors, especially the errors that are related to offline arrays.

If the recovery completes with offline volumes, go to “Recovering from offline volumes using the CLI.”

After you complete the storage system recovery procedure, contact support for assistance with recovering the file modules, so access to the file systems can be restored.

Recovering from offline volumes using the CLI

If a Tier 3 recovery procedure completes with offline volumes, then it is likely that the data which was in the write-cache of the node canisters was lost during the failure that caused all of the node canisters to lose the block storage system cluster state. You can use the command-line interface (CLI) to acknowledge that there was lost data lost from the write-cache, and bring the volume back online to attempt to deal with the data loss.

About this task

If you have run the recovery procedure but there are offline volumes, you can complete the following steps to bring the volumes back online. Any volumes that are offline and are not thin-provisioned (or compressed) volumes are offline because of the loss of write-cache data during the event that led all node canisters to lose their cluster state. Any data lost from the write-cache cannot be recovered. These volumes might need additional recovery steps after the volume is brought back online.

Note: If you encounter errors in the error log after running the recovery procedure that are related to offline arrays, use the fix procedures to resolve the offline array errors before fixing the offline volume errors.

Example

Complete the following steps to recover an offline volume after the recovery procedure has completed:

1. Delete all IBM FlashCopy function mappings and Metro Mirror or Global Mirror relationships that use the offline volumes.
2. Run the **recovervdisk** or **recovervdiskbysystem** command. (This will only bring the volume back online so that you can attempt to deal with the data loss.)

Contact IBM Remote Technical Support to help you with recovering from file volumes that have been corrupted by data lost from the write-cache. They might ask you to refer to “Recovering a GPFS file system” on page 196 and help you with interpreting the results from the **chkfs** CLI command.

3. Refer to “What to check after running the system recovery” for what to do with volumes that have been corrupted by the loss of data from the write-cache.
4. Recreate all FlashCopy mappings and Metro Mirror or Global Mirror relationships that use the volumes.

What to check after running the system recovery

Several tasks must be completed before you use the system.

The recovery procedure recreates the old system from the quorum data. However, some things cannot be restored, such as cached data or system data managing in-flight I/O. This latter loss of state affects RAID arrays managing internal storage. The detailed map about where data is out of synchronization has been lost, meaning that all parity information must be restored, and mirrored pairs must be brought back into synchronization. Normally this results in either old or stale data being used, so only writes in flight are affected. However, if the array had lost redundancy (such as syncing, or degraded or critical RAID status) prior to the error requiring system recovery, then the situation is more severe. Under this situation you need to check the internal storage:

- Parity arrays will likely be syncing to restore parity; they do not have redundancy when this operation proceeds.
- Because there is no redundancy in this process, bad blocks might have been created where data is not accessible.
- Parity arrays could be marked as corrupt. This indicates that the extent of lost data is wider than in-flight I/O, and in order to bring the array online, the data loss must be acknowledged.
- RAID-6 arrays that were actually degraded prior the system recovery might require a full restore from backup. For this reason, it is important to have at least a capacity match spare available.

Be aware of these differences regarding the recovered configuration:

- FlashCopy mappings are restored as “idle_or_copied” with 0% progress. Both volumes must have been restored to their original I/O groups.
- The management ID is different. Any scripts or associated programs that refer to the system-management ID of the clustered system (system) must be changed.
- Any FlashCopy mappings that were not in the “idle_or_copied” state with 100% progress at the point of disaster have inconsistent data on their target disks. These mappings must be restarted.
- Intersystem remote copy partnerships and relationships are not restored and must be re-created manually.
- Consistency groups are not restored and must be re-created manually.
- Intrasystem remote copy relationships are restored if all dependencies were successfully restored to their original I/O groups.
- If hardware was replaced before the recovery, the SSL certificate might not be restored. If it is not restored, then a new self-signed certificate is generated with a validity of 30 days. Follow the associated Directed Maintenance Procedures (DMP) for a permanent resolution.
- The system time zone might not have been restored.
- The GPFS system quorum state held on the control enclosure might not have been restored.
- Any Global Mirror secondary volumes on the recovered system might have inconsistent data if there was replication I/O from the primary volume cached

on the secondary system at the point of the disaster. A full synchronization is required when recreating and restarting these remote copy relationships.

- Immediately after the T3 recovery process runs, compressed disks do not know the correct value of their used capacity. The disks initially set the capacity as the entire real capacity. When I/O resumes, the capacity is shrunk down to the correct value.

Similar behavior occurs when you use the `-autoexpand` option on vdisks. The real capacity of a disk might increase slightly, caused by the same kind of behavior that affects compressed vdisks. Again, the capacity shrinks down as I/O to the disk is resumed.

Before using the block volumes that are accessed by the SAN or with iSCSI, complete the following tasks:

- Start the block host systems.
- Manual actions might be necessary on the hosts to trigger them to rescan for devices. You can complete this task by disconnecting and reconnecting the Fibre Channel cables to each host bus adapter (HBA) port.
- Verify that all mapped volumes can be accessed by the hosts.
- Run file system consistency checks on the block hosts.
- Run the application consistency checks.

Before using the file volumes that are used by GPFS on the file modules to provide Network Attached Storage (NAS), complete the following task:

- Contact IBM support for assistance with recovering the GPFS quorum state so that access to files as NAS can be restored.

For Virtual Volumes (VVols), complete the following tasks.

- After you confirm that the T3 completed successfully, restart Spectrum Control Base (SCB) services. Use the Spectrum Control Base command **`service ibm_spectrum_control start`**.
- Refresh the storage system information on the SCB GUI to ensure that the systems are in sync after the recovery.
 - To complete this task, login to the SCB GUI.
 - Hover over the affected storage system, select the menu launcher, and then select **Refresh**. This step repopulates the system.
 - Repeat this step for all Spectrum Control Base instances.
- Rescan the storage providers from within the vSphere Web Client.
 - Select **vCSA > Manage > Storage Providers > select Active VP > Re-scan icon**.

For Virtual Volumes (VVols), also be aware of the following information.

FlashCopy mappings are not restored for VVols. The implications are as follows.

- The mappings that describe the VM's snapshot relationships are lost. However, the Virtual Volumes that are associated with these snapshots still exist, and the snapshots might still appear on the vSphere Web Client. This outcome might have implications on your VMware back-up solution.
 - Do not attempt to revert to snapshots.
 - Use the vSphere Web Client to delete any snapshots for VMs on a VVol data store to free up disk space that is being used unnecessarily.

- The targets of any outstanding 'clone' FlashCopy relationships might not function as expected (even if the vSphere Web Client recently reported clone operations as complete). For any VMs, which are targets of recent clone operations, complete the following tasks.
 - Perform data integrity checks as is recommended for conventional volumes.
 - If clones do not function as expected or show signs of corrupted data, take a fresh clone of the source VM to ensure that data integrity is maintained.

Backing up and restoring the system configuration

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

Configuration data for the system provides information about your block system and the objects that are defined in it. The backup and restore functions of the **svconfig** command can back up and restore only your configuration data for the Storwize V7000 system. You must regularly back up your file systems and your application data by using the appropriate backup methods.

You can maintain your configuration data for the system by completing the following tasks:

- Backing up the configuration data
- Restoring the configuration data
- Deleting unwanted backup configuration data files

Before you back up your configuration data, the following prerequisites must be met:

- No independent operations that change the configuration for the system can be running while the backup command is running.
- No object name can begin with an underscore character (_).

Note:

- The default object names for controllers, I/O groups, and managed disks (MDisks) do not restore correctly if the ID of the object is different from what is recorded in the current configuration data file.
- All other objects with default names are renamed during the restore process. The new names appear in the format *name_r* where *name* is the name of the object in your system.

Contact the IBM support center to help you prepare the Storwize V7000 Unified system to do the restoring of the system configuration on the control enclosure.

The configuration restore procedure is designed to restore the information about your block storage configuration, such as block volumes, local Metro Mirror information, local Global Mirror information, storage pools, and nodes. All the data that is written to the block volumes is not restored.

To restore the data on the block volumes, you must restore the application data separately from any application that uses the volumes on the clustered system as storage. The file volumes are not restored. You must restore the file module configuration and the file systems separately. Therefore, you must have a backup of this data before you follow the configuration recovery process.

Before you restore your configuration data, the following prerequisites must be met:

- You have the Security Administrator role associated with your user name and password.
- You have a copy of your backup configuration files on a server that is accessible to the system.
- You have a backup copy of your application data that is ready to load on your system after the restore configuration operation is complete.
- You know the current license settings for your system.
- You did not remove any hardware since the last backup of your configuration.
- No zoning changes were made on the Fibre Channel fabric which would prevent communication between the Storwize V7000 Unified and any storage controllers which are present in the configuration.
- For configurations with more than one I/O group, if a new system is created on which the configuration data is to be restored, the I/O groups for the other control enclosures must be added.
- You have at least 3 USB flash drives if encryption was enabled on the system when its configuration was backed up. The USB flash drives are used for generation of new keys as part of the restore process or for manually restoring encryption if the system has less than 3 USB ports.

Use the following steps to determine how to achieve an ideal T4 recovery:

- Open the appropriate `svc.config.backup.xml` (or `svc.config.cron.xml`) file with a suitable text editor or browser and navigate to the **node section** of the file.
- For each node entry, make a note of the value of the following properties: `IO_group_id`, `canister_id`, `enclosure_serial_number`.
- Use the CLI **sainfo lsservicenodes** command and the data to determine which node canisters previously belonged in each I/O group.

Restoring the system configuration should be performed via one of the nodes previously in IO group zero. For example, **property name="IO_group_id" value="0"**. The remaining enclosures should be added, as required, in the appropriate order based on the previous **IO_group_id** of its node canisters.

Note: It is not currently possible to determine which canister within the identified enclosure was previously used for cluster creation. Typically the restoration should be performed via canister 1.

Before you begin, hardware recovery must be complete. The following hardware must be operational: hosts, Storwize V7000 Unified enclosures, internal flash drives and expansion enclosures (if applicable), the Ethernet network, the SAN fabric, and any external storage systems (if applicable).

Backing up the system configuration using the CLI

You can back up your configuration data using the command-line interface (CLI).

Before you begin

Before you back up your configuration data, the following prerequisites must be met:

- No independent operations that change the configuration can be running while the backup command is running.
- No object name can begin with an underscore character (`_`).

About this task

The backup feature of the **svcconfig** CLI command is designed to back up information about your system configuration, such as volumes, local Metro Mirror information, local Global Mirror information, storage pools, and nodes. All other data that you wrote to the volumes is *not* backed up. Any application that uses the volumes on the system as storage, must use the appropriate backup methods to back up its application data.

You must regularly back up your configuration data and your application data to avoid data loss, such as after any significant changes to the system configuration.

Note: The system automatically creates a backup of the configuration data each day at 1 AM. This backup is known as a **cron** backup and is written to `/dumps/svc.config.cron.xml_serial#` on the configuration node.

Use the these instructions to generate a manual backup at any time. If a severe failure occurs, both the configuration of the system and application data might be lost. The backup of the configuration data can be used to restore the system configuration to the exact state it was in before the failure. In some cases, it might be possible to automatically recover the application data. This backup can be attempted with the Recover System Procedure, also known as a Tier 3 (T3) procedure. To restore the system configuration without attempting to recover the application data, use the Restoring the System Configuration procedure, also known as a Tier 4 (T4) recovery. Both of these procedures require a recent backup of the configuration data.

Note: The configuration applies to the Storwize V7000 only, and not the file modules.

Complete the following steps to back up your configuration data:

Procedure

1. Use your preferred backup method to back up all of the application data that you stored on your volumes.
2. Issue the following CLI command to back up your configuration:
`svcconfig backup`

The following output is an example of the messages that might be displayed during the backup process:

```
CMMVC6155I SVCCONFIG processing completed successfully
```

The **svcconfig backup** CLI command creates three files that provide information about the backup process and the configuration. These files are created in the `/dumps` directory of the configuration node canister.

Table 117 describes the three files that are created by the backup process:

Table 117. Files created by the backup process

File name	Description
<code>svc.config.backup.xml_serial#</code>	Contains your configuration data.
<code>svc.config.backup.sh_serial#</code>	Contains the names of the commands that were issued to create the backup of the system.

Table 117. Files created by the backup process (continued)

File name	Description
svc.config.backup.log_<serial#>	Contains details about the backup, including any reported errors or warnings.

3. Check that the **svcconfig backup** command completes successfully, and examine the command output for any warnings or errors. The following output is an example of the message that is displayed when the backup process is successful:

```
CMMVC6155I SVCCONFIG processing completed successfully.
```

If the process fails, resolve the errors, and run the command again.

4. Keep backup copies of the files outside the system to protect them against a system hardware failure. Copy the backup files off the system to a secure location; use either the management GUI or SmartCloud Provisioning command line. For example:

```
pscp -unsafe superuser@cluster_ip:/dumps/svc.config.backup.*
/offclusterstorage/
```

The `cluster_ip` is the IP address or DNS name of the system and **offclusterstorage** is the location where you want to store the backup files.

Tip: To maintain controlled access to your configuration data, copy the backup files to a location that is password-protected.

Restoring the system configuration

Use this procedure in the following situations: only if the recover procedure failed or if the data that is stored on the volumes is not required. Use this procedure in the following situations: only if the recover procedure failed, if the data that is stored on the volumes is not required, or if the files that are stored in the file volumes by the file modules are not required. For directions on the recover procedure, see “Recover system procedure” on page 381.

Before you begin

This configuration restore procedure is designed to restore information about your configuration, such as volumes, local Metro Mirror information, local Global Mirror information, storage pools, and nodes. The data that you wrote to the volumes is not restored. To restore the data on the volumes, you must restore application data from any application that uses the volumes on the clustered system as storage separately. You must restore the file module configuration and the file systems separately. Therefore, you must have a backup of this data before you follow the configuration recovery process.

If encryption was enabled on the system when its configuration was backed up, then at least 3 USB flash drives need to be present in the node canister USB ports for the configuration restore to work. The USB flash drives do not need to contain any keys. They are for generation of new keys as part of the restore process.

About this task

You must regularly back up your configuration data and your application data to avoid data loss. If a system is lost after a severe failure occurs, both configuration

for the system and application data is lost. You must restore the system to the exact state it was in before the failure, and then recover the application data.

During the restore process, the nodes and the storage enclosure will be restored to the system, and then the MDisks and the array will be re-created and configured. If there are multiple storage enclosures involved, the arrays and MDisks will be restored on the proper enclosures based on the enclosure IDs.

If you do not understand the instructions to run the CLI commands, see the command-line interface reference information.

To restore your configuration data, follow these steps:

Procedure

1. Verify that all nodes are available as candidate nodes before you run this recovery procedure. You must remove errors 550 or 578 to put the node in candidate state.
2. Create a system.
 - If your system is a Storwize V7000 Gen2 system, use the technician port.
 - If your system is a Storwize V7000 Gen1 system, use the initialization tool that is available on the USB flash drive. Select the **Initialize a new** Storwize V7000 Unified (block system only) option from the **Welcome** panel of the initialization tool.
3. In a supported browser, enter the IP address that you used to initialize the system and the default superuser password (passw0rd).
4. If the clustered system was previously configured as replication layer, then use the **chsystem** command to change the layer setting.
5. Identify the configuration backup file from which you want to restore.

The file can be either a local copy of the configuration backup XML file that you saved when you backed-up the configuration or an up-to-date file on one of the nodes.

Configuration data is automatically backed up daily at 01:00 system time on the configuration node.

Download and check the configuration backup files on all nodes that were previously in the system to identify the one containing the most recent complete backup

- a. From the management GUI, click **Settings > Support**.
- b. Click **Show full log listing**.
- c. For each node (canister) in the system, complete the following steps:
 - 1) Select the node to operate on from the selection box at the top of the table.
 - 2) Find all the files with names that match the pattern `svc.config.*.xml*`.
 - 3) Double-click the files to download them to your computer.

The XML files contain a date and time that can be used to identify the most recent backup. After you identify the backup XML file that is to be used when you restore the system, rename the file to `svc.config.backup.xml`.

6. Copy onto the system the XML backup file from which you want to restore.

```
pscp full_path_to_identified_svc.config.file
superuser@cluster_ip:/tmp/svc.config.backup.xml
```
7. If the system was originally configured as a replication layer system, change the layer of the system to replication by running the following command:

```
svctask chsystem -layer replication
```

If the command fails with the following error, make sure that no other controllers are visible to the system, and then go to step 7 on page 394.CMMVC7143E The command cannot be initiated because nodes from another cluster are visible.

8. Issue the following CLI command to compare the current configuration with the backup configuration data file:

```
svcconfig restore -prepare
```

This CLI command creates a log file in the /tmp directory of the configuration node. The name of the log file is svc.config.restore.prepare.log.

Note: It can take up to a minute for each 256-MDisk batch to be discovered. If you receive error message CMMVC6200W for an MDisk after you enter this command, all the managed disks (MDisks) might not be discovered yet. Allow a suitable time to elapse and try the **svcconfig restore -prepare** command again.

9. Issue the following command to copy the log file to another server that is accessible to the system:

```
pscp superuser@cluster_ip:/tmp/svc.config.restore.prepare.log  
full_path_for_where_to_copy_log_files
```

10. Open the log file from the server where the copy is now stored.

11. Check the log file for errors.

- If you find errors, correct the condition that caused the errors and reissue the command. You must correct all errors before you can proceed to step 12.
- If you need assistance, contact the IBM Support Center.

12. Issue the following CLI command to restore the configuration:

```
svcconfig restore -execute
```

This CLI command creates a log file in the /tmp directory of the configuration node. The name of the log file is svc.config.restore.execute.log.

13. Issue the following command to copy the log file to another server that is accessible to the system:

```
pscp superuser@cluster_ip:/tmp/svc.config.restore.execute.log  
full_path_for_where_to_copy_log_files
```

14. Open the log file from the server where the copy is now stored.

15. Check the log file to ensure that no errors or warnings occurred.

Note: You might receive a warning that states that a licensed feature is not enabled. This message means that after the recovery process, the current license settings do not match the previous license settings. The recovery process continues normally and you can enter the correct license settings in the management GUI later.

When you log in to the CLI again over SSH, you see this output:

What to do next

You can remove any unwanted configuration backup and restore files from the /tmp directory on your configuration by issuing the following CLI command:

```
svcconfig clear -all
```

Deleting backup configuration files using the CLI

You can use the command-line interface (CLI) to delete backup configuration files.

About this task

Complete the following steps to delete backup configuration files:

Procedure

1. Issue the following command to log on to the system:

```
plink -i ssh_private_key_file superuser@control_enclosure_management_ip
```

where *ssh_private_key_file* is the name of the SSH private key file for the superuser and *control_enclosure_management_ip* is the IP address or DNS name of the system from which you want to delete the configuration.
2. Issue the following CLI command to erase all of the files that are stored in the /tmp directory:

```
svcconfig clear -all
```

Chapter 6. Call home and remote support

This section includes information for testing call home connections.

The call home function enables the system to automatically notify IBM Support about the hardware configuration and status of the system. Using this information, IBM Support can contact the system administrator in case of any issues.

Testing a call home connection

Use this information to test a call home connection to the IBM support.

From the block-level storage system

If call home actions fail, perform the following steps:

1. Go to **Settings > Support > Call Home** on the management GUI.
2. To check for any email connection problem, in the **General** group, click **Test Email Server Connection** and type an email ID to receive a call home test email.
3. If the test fails, contact your network administrator and verify that the email server settings are correct.

From file modules

If call home actions fail, perform the following steps:

1. Go to **Settings > Support > Call Home Log** on the management GUI and review the status of the attempted call home actions.
2. For a call home entry with **Failed** status, review information under the Details column. There can be issues with the configuration or outbound connectivity between this system and the IBM call home server.
3. To check for any call home connection problem, go to the **Outbound Connectivity** group under the **Call Home** tab and click **Test Server Connection**.
4. If the test fails, contact your network administrator and verify that the proxy server settings are correct.

Establishing an AOS connection

Use this information to establish an AOS connection with IBM remote support for diagnosing and reviewing issues and problems on your system.

Before establishing a connection, be sure that you configure the system for AOS by following the tasks in **Installing > Adding file modules to an existing Storwize V7000 system > Post configuration of the Storwize V7000 Unified system** located in the Information Center,

Establishing a lights-out AOS connection

Use this information to establish a lights-out AOS connection with IBM remote support for diagnosing and reviewing issues and problems on your Storwize V7000 Unified system.

About this task

Configure the system for a lights-out connection using the Enable IBM Tivoli® Assist On-Site (AOS) task.

After you configure the system, no other tasks are needed. The remote support contact might ask you for machine information, such as machine type and models, serial numbers, and your machine name. This information helps them locate the system within the back-end AOS connection point repository. The repository is an internal list that shows all available systems that are configured for lights-out connectivity.

Establishing a lights-on AOS connection

Use this information to establish a lights-on AOS connection with IBM remote support for diagnosing and reviewing issues and problems on your IBM Storwize V7000 Unified system.

About this task

This procedure requires that a keyboard, video, and mouse is attached to the local IBM Storwize V7000 Unified file module and that a customer representative is physically present at the connection for the duration of the remote support session.

To establish the AOS connection, perform the following steps:

Note: Each step at the beginning identifies whether the **remote IBM support representative** or the **customer** in the customer data center performs the step.

Procedure

1. **Remote IBM support representative:** Start the connection process from your remote location.
 - a. Establish telephone or Sametime® communications to the IBM authorized servicer at the customer site to find out the problem maintenance request (PMR) number if you do not know it already, and the customer name and geography.
 - b. Open the AOS console and click the connect icon (the plug icon).
 - c. Enter your AOS user ID and password.
 - d. Select the HTTP link type of connection.
 - e. Enter the customer name, the case number (use the PMR number), and the geography.
 - f. Talk to the IBM authorized servicer at the customer site to make sure that the servicer is ready to establish the link before you submit the form.
 - g. Submit the form to the AOS server.
2. **Remote IBM support representative:** Wait for the AOS console to display the connection code when the AOS server returns the code.
3. **Remote IBM support representative:** Communicate the connection code to the IBM authorized servicer at the customer site.

Note: The connection code has a default timeout of 5 minutes. If the IBM authorized servicer at the customer site takes longer than 5 minutes to link to the AOS server, you can extend it for 5 minutes (twice). After the link is established, the link stays active until either you or the authorized servicer breaks the connection.

4. **Customer:** From the file module, log in as root and run **cnrslaunchaos**.
5. **Customer:** Enter the connection code that the IBM support representative gave you.
The script launches the Firefox browser and downloads the executable for establishing the AOS session. Confirm the file download. The file is stored in the /home/root/desktop directory.
6. **Customer:** When the executable file finishes downloading, close the Firefox download window and close the browser.
The launch script runs the AOS binary executable file that it downloaded.
7. **Customer:** Grant the IBM support representative the appropriate level of access (Active, Monitor, or Chat) according to customer security for conducting the maintenance action. For example, click **Active**.
Active mode gives full remote access.
Monitor mode restricts the IBM support representative to a view of the console, where the representative can offer guidance on what actions you might take to analyze and correct the problem.
Chat mode opens a chat window with no view of the console.

Chapter 7. Recovery procedures

This section covers the recovery procedures for the file modules and the control enclosure.

User ID and system access

This section covers the recovery procedures for the topics that support the user ID and system access.

Accessing a file module as root

Some procedures require that you log on to a file module as root.

About this task

You can use the following methods to access a file module as root.

Procedure

Access a file module as root.

- Type the following command in an X terminal, for example, a Windows or a Linux operating system:
`ssh -p 1602 root@<file module IP>`
- Use a Windows application like PuTTY to ssh to port 1602 of a file module service IP and log in as root with the root password that you recorded in your access information. See “Record the access information” on page 37.

Recovering from losing the root password

Some recovery procedures require the root password to be entered for the file module.

Before you begin

If you have forgotten the file module root password, you can follow the procedure to change it from any file module user ID that has sufficient authority to run the **chrootpwd** command successfully.

About this task

To recover a lost root password, perform the following steps:

Procedure

Log in to the management CLI as admin:

Issue the **chrootpwd** command to change the password. Changing the password is the easiest way to recover the root password. If the **chrootpwd** command does not work, continue to the next step and finish the lost root password recovery procedure.

```
chrootpwd [-c { clusterID | clusterName }] [-p password] [-v]
```

Results

The chrootpwd program prompts you for the new root password.

The chrootpwd program sets the new root password on both file modules in the cluster.

Resetting the NAS ssh key for configuration communications

The configuration communications between the Storwize V7000 file modules and the control enclosure are done by using SSH over the site 1 Gbps Ethernet LAN; whereas the file data traffic is passed over the direct connect Fibre Channel links by using the SCSI protocol.

Before you begin

During the USB initialization of the Storwize V7000 Unified system, one of the node canisters in the control enclosure creates a public/private key pair to use for ssh. The node canister stores the public key and writes the private key to the USB flash drive memory.

One of the file modules then takes the private key from the USB flash drive memory to use for ssh. The file module passes it to the other file module over the direct connect Ethernet link and then deletes the private key from the USB flash drive memory so that it cannot be used on the wrong system.

It might be necessary to reset the NAS SSH key in the following circumstances:

- When communications between the Storwize V7000 file module and the Storwize V7000 control enclosure is not authorized because of a bad key.
- When both Storwize V7000 file modules have lost the original NAS ssh key.
- When the Storwize V7000 control enclosure has lost the NAS ssh key.

About this task

Reset the NAS SSH key so that the communications between the file modules and the Storwize V7000 control enclosure resume. To reset the ssh key, generate the NAS.ppk file on the Storwize V7000, and transfer it via **SCP** to the management node to import it.

Procedure

1. Log on to the Storwize V7000 control enclosure management CLI as superuser and run the following command to generate the new NAS SSH key:

```
satask chnaskey -privkeyfile NAS.ppk
```

The private key is left in the /dumps directory.
2. Use the management GUI to see which of the file modules is the active management node and find the IP address for that file module.
3. Log on to the file module that is the active management node as CLI user via the file module IP address that you identified in the previous step. For example:

```
ssh -p 1602 admin@file module IP address
```

You are prompted for the admin password twice.
4. Use **SCP** to copy the private key file from the /dumps directory on the Storwize V7000 to the /tmp directory on this file module using the following command:

```
scp superuser@system IP address:/dumps/NAS.ppk /tmp
```

You are prompted for the Storwize V7000 superuser password.

5. Log on to the Storwize V7000 Unified management CLI as admin via the management IP and run the following command to register the new NAS SSH key:

```
chstoragesystem --sonasprivkey /tmp/NAS.ppk
```

Working with NFS clients that fail to mount NFS shares after a client IP change

Use this information to resolve a refused mount or Stale NFS file handle response to an attempt to mount Network File System (NFS) shares after a client IP change.

About this task

After a client IP change, a **df -h** command can return no results, as shown in the following example:

Filesystem	Size	Used	Avail	Use%	Mounted on
machinename: filename:	-	-	-	-	/sharename

The **ls** command can return the following error:

```
ls: .: Stale NFS file handle
```

The Storwize V7000 Unified system hosting file module might display the following error:

```
mgmt002st001 mountd[3055867]: refused mount  
request from hostname for sharename (/): not exported
```

Note: The following steps do not apply if you are using NFSv4.0 server.

If one of these errors occurs, complete the following steps.

Procedure

1. Access the file module CLI as a privileged user.
2. Issue the **sc onnode all /usr/sbin/exportfs -a** command to flush the NFS cache in each file module.
3. Verify that the NFS mount is successful. If the problem persists, restart the NFS service on the file module that is refusing the mount requests from that client.
4. Verify that the NFS share mount is successful.

Working with file modules that report a stale NFS file handle

To recover from the “Stale NFS file handle” file system state on a file module, you must suspend, reboot, and resume the file module.

About this task

Note: If the “Stale NFS file handle” message was displayed after a client IP change, refer to “Working with NFS clients that fail to mount NFS shares after a client IP change.”

Because of errors or conditions related to this file module, the file module disconnected itself from the file system that was shared with the other nodes. All of the file descriptors that were opened to the file system through this file module have become “stale”, as indicated by command output or a Stale NFS file handle

error message, and cannot access their corresponding files. When this occurs, all affected file modules enter an unhealthy state, and a CIM similar to the following is sent to the alert log:

```
GPFS Error - check stale file handle failed with error code 1:  
see stale file handle on /ibm/gpfs0 on file module: mgmt001st001
```

If you receive the error above, complete the following steps:

Procedure

1. Open the CLI with a privileged user, and issue **sc /usr/sbin/exportfs -a t** to flush the NFS cache in each file module. Verify that the state of each affected file module is healthy and that no new “Stale NFS file handle” CIMs are displayed in the alert log after you resume the file module. If the problem persists, continue with the following steps.
2. Review the event log to identify the affected file system and all of the nodes where the file system displays the state “Stale NFS file handle”.
3. Suspend each affected file module.
4. Reboot each affected file module.
5. Resume each affected file module.
6. Verify that the state of each affected file module is healthy after you resume the file module and that no new “Stale NFS file handle” CIMs appear in the alert log.

Recovering the GPFS

1. Enter **lsnode -r**. This displays GPFS and CTDB status.
2. Check for the GPFS that are in stale mount by running **lsmount** or **mm1smount gpfs1 -L**.
3. Enter **sc onnode all df**. This displays **df: 'ibm/gpfsX': Stale NFS handle** where X is the gpfs number such as gpfs0.

To recover gpfs, follow the procedure below:

Note: You need to perform these steps on the active management node.

1. Enter **lsnode -r**. Note down the problematic node.
2. Enter **initnode -r -n <node>** to reboot the affected node.

Ping the node or check uptime or wait for the node to come up.

1. **lsnode -r**
2. **resumenode <node>**

Perform the above procedure for all the nodes in the cluster.

File module-related issues

This section covers the recovery procedures related to file module issues.

Restoring System x firmware (BIOS) settings

During critical repair actions such as the replacement of a system planar in an IBM Storwize V7000 Unified file module, you might have to reset the System x firmware.

Before you begin

The firmware and software code package for the Storwize V7000 Unified microcode can automatically configure the default settings for the System x firmware to the required Storwize V7000 Unified settings. However, to enable the automatic configuration, you must reset the System x firmware from its current state to the default configuration.

About this task

Use the following procedure to set the System x firmware to the default state and start the automatic Storwize V7000 Unified configuration.

Note: After the power-on, the installation of the firmware can take up to 70 minutes in some cases.

Procedure

1. Attach the USB and VGA cable from the KVM switch to the video and USB port of the Storwize V7000 Unified file module that is to be reset.
2. Open the KVM unit in the base rack.
3. Turn on the KVM unit if it is off.
4. Press **PrtSc** to display the KVM selector screen.
5. Scroll down to select the USB cable, then press **Enter**.
6. Turn on the affected file module.
7. From the IBM System x Server Firmware screen, press **F1** to set up the firmware.

A few seconds after the IBM System x Server Firmware screen is displayed, F1, and other options are displayed at the bottom of the screen:

- F1 - Setup
- F2 - Diagnostics
- F12 - Select Boot Device

8. From the System Configuration and Boot Management screen, scroll down to click **Load Default Settings**, and then press **Enter**.
The screen goes blank for a few seconds and then returns to the System Configuration and Boot Management screen.
9. Click **Save**.
10. A window displays a prompt to ask to reset the IMM now. Select **Y**.
11. Press **ESC** twice to return to the System Configuration and Boot Management screen.
12. Scroll down to click **Boot Manager**, and then press **Enter**.
13. Scroll down to click **Add Boot Option**, and then press **Enter**.
14. Scroll down to click **Legacy Only**, and then press **Enter**.
The option is not visible until you scroll down. Selecting the option removes it from the list of available options.
15. Press **ESC** twice to return to the System Configuration and Boot Management screen.
16. Scroll down to click **Save Settings**, and then press **Enter**.
17. Press **ESC** or click **Exit Setup**, and then press **Enter**.
18. When prompted, click **Y** to exit the setup menu.

The system now reboots. During the reboot, the Storwize V7000 Unified code automatically modifies the configuration of the System x firmware (BIOS) to change the default settings to the required settings.

Recovering from file systems that are offline after the volumes came back online

The problem that caused the file volumes to go offline long enough to cause the file systems to become unmounted may have caused disks to be marked as failed which will prevent the file system from being automatically mounting after the volume comes back online.

About this task

The file systems will usually be automatically mounted as soon as the file volumes come back online. However, if GPFS experienced IO errors as the volume went offline then it may mark the disks as failed.

If this happens then the automatic mounting of the file system will not work and the disks must be started using the **Start All Disks** action against the file system in the management GUI before they are mounted using the management GUI.

Procedure

To re mount any file system that was not automatically remounted when the file volumes came back online:

1. Go to the **files > file systems** page in the management GUI to see if any file systems are offline.
2. Hover over the Status indicator of any file system which does not have an OK status.
3. If **The <pool name> file system pool contains failing disks** is displayed then select the action to **Start All Disks** used by this file system.
4. If hovering over the Status indicator of the file system shows that the file system is not mounted on any node, or on one of the nodes, then select the action to mount the file system.

Results

If the health status indicator is still red after completing all recovery procedures then refer to “Health status and recovery” on page 62 to help you return the health status indicator back to green.

Recovering from a multipath event

Use this procedure to recover a node from a **multipathd** failure.

Before you begin

Use this procedure after completing the procedure in Fibre Channel connectivity between file modules and control enclosure.

The Storwize V7000 Unified system can experience problems where the **multipathd** failures occur. If the paths are not automatically restored, a system reboot can recover the paths.

Important: Perform this procedure against the passive management node only.

Procedure

1. Verify that the node that the **multipathd** event occurred against is the passive management node. If the node that experiences the **multipathd** problems is the active node, then perform the management node failover procedure. See “Performing management node role failover on a “good” system” on page 184.
2. Reboot the file module. See “Rebooting a file module” on page 91.

Diagnosing a multipath event

The **multipath -ll** command verifies that all storage devices are either active or not active.

The following output shows that all storage devices are active.

```
[root@yourmachine.mgmt001st001 ~]# multipath -ll
array1_sas_89360007 (360001ff070e9c00000000001989360007) fm-0 IBM,2073-720
[size=3.1T][features=1 queue_if_no_path][hwhandler=0][rw]
\_ round-robin 0 [prio=50][active]
\_ 6:0:0:0 sdb 8:16 [active][ready]
\_ round-robin 0 [prio=10][enabled]
\_ 8:0:0:0 sdg 8:96 [active][ready]
array1_sas_89380009 (360001ff070e9c00000000001b89380009) fm-1 IBM,2073-720
[size=3.1T][features=1 queue_if_no_path][hwhandler=0][rw]
\_ round-robin 0 [prio=50][active]
\_ 6:0:0:2 sdd 8:48 [active][ready]
\_ round-robin 0 [prio=10][enabled]
\_ 8:0:0:2 sdi 8:128 [active][ready]
```

The following output shows that the storage devices are not active.

```
[root@kd271f6.mgmt002st001 ~]# multipath -ll
mpathq (360050768029180b060000000000000007) dm-8 IBM,2145
size=2.5G features='1 queue_if_no_path' hwhandler='0' wp=rw
| ~- 5:0:0:7 sdr 65:16 failed ready running
| ~- 6:0:0:7 sdi 8:128 failed ready running
mpathp (360050768029180b060000000000000005) dm-3 IBM,2145
size=2.5G features='1 queue_if_no_path' hwhandler='0' wp=rw
| ~- 5:0:0:5 sdp 8:240 failed ready running
| ~- 6:0:0:5 sdg 8:96 failed ready running
```

The output [active][ready] identifies an active device. The output failed ready running identifies a device that is not active.

Recovering from an NFSD service error

Use this procedure to recover from an NFSD service error.

About this task

This recovery procedure starts the NFSD when it is down.

Procedure

1. Log in as a CLI user with privileged authority.
2. Issue the **sc service nfsd start** command.
3. If the problem persists, restart the node.
4. If the restart action does not resolve the issue, contact the next level of support.

Recovering from an SCM error

Use this procedure to recover from a service configuration management (SCM) error.

About this task

Complete the following procedure if output from the **lshealth -r** CLI command contains a line similar to the following:

```
SCM                ERROR   SCM system has found some errors
```

Note: This procedure involves analyzing various logs depending on the errors displayed by the initial SCM error log.

Procedure

1. If an error is displayed, run the **lshealth -i SCM** command to show the details of the component with the error. SCM is a component, that monitors other components. Ensure to note the details shown by the **Message** and **Value** columns.
2. To know the error code, run the **lslog** command or open the graphical user interface (GUI) Eventlog page.
3. Compare the results returned by **lslog** command with **lshealth -i SCM** command. This procedure helps you in mapping the error. If you are not able to link the **lshealth -i SCM** output with the **lslog** output, continue to the next step.
4. Open the CNSCM log located at **/var/log/cnlog/cnscm** for the file module that reported the error.
5. Review the error entries around the listed time stamp and then check the log for issues that seem related that occurred before the listed time stamp. For example, you might find GPFS-related issues appearing earlier and later, too.
6. Review the log entries and try to match the entries with the **lslog** output. If you are not able to match the entries, continue to the next step.
7. Based on the log entries, check the appropriate corresponding log. If the issue appeared to be related to GPFS, for example, you could look for the root cause in **/var/adm/ras/mmfs.log**.
8. If the log entries do not help resolve the error, contact the next level of support.

Recovering from an httpd service error

Use this procedure to recover from an httpd service error when the service is reported as unhealthy or off.

About this task

Procedure

To fix the httpd error, perform the following steps:

1. Attempt to start the http service manually.
 - a. Log in as CLI user with privileged authority.
 - b. Issue the **sc service http start** command.
2. When you complete the service action, refer to “Health status and recovery” on page 62.

Recovering from an sshd_data service error

Use this procedure to recover from an sshd_data service error.

About this task

This recovery procedure starts the `sshd_data` when it is down.

Procedure

1. Log in as a CLI user with privileged authority.
2. Issue the service **`sc sshd_data start`** command.
3. If the problem persists, restart the node.
4. If the restart action does not resolve the issue, contact the next level of support.

Recovering from an `sshd_int` service error

Use this procedure to recover from an `sshd_int` service error.

About this task

This recovery procedure starts the `sshd_int` when it is down.

Procedure

1. Log in as a CLI user with privileged authority.
2. Issue the service **`sc sshd_int start`** command.
3. If the problem persists, restart the node.
4. If the restart action does not resolve the issue, contact the next level of support.

Recovering from an `sshd_mgmt` service error

Use this procedure to recover from an `sshd_mgmt` service error.

About this task

This recovery procedure starts the `sshd_mgmt` when it is down.

Procedure

1. Log in as a CLI user with privileged authority.
2. Issue the service **`sc sshd_mgmt start`** command.
3. If the problem persists, restart the node.
4. If the restart action does not resolve the issue, contact the next level of support.

Recovering from an `sshd_service` service error

Use this procedure to recover from an `sshd_service` service error.

About this task

This recovery procedure starts the `sshd_service` when it is down.

Procedure

1. Log in as a CLI user with privileged authority.
2. Issue the service **`sc sshd_service start`** command.
3. If the problem persists, restart the node.
4. If the restart action does not resolve the issue, contact the next level of support.

Recovering from the 1710 bus error due to the /var directory being full

If the /var directory is full and a node is rebooted, the 1710 bus error message might appear on the console screen.

During code upgrade or troubleshooting, or during normal system usage, when a node is rebooted, the operating system might be unable to boot and you might get the above error message. In such cases, check the size of the /var directory as follows:

1. On the management node, use the following steps to boot into the single-user mode using GRUB:
 - a. At the **GRUB splash** screen at boot time, press any key to enter the **GRUB interactive** menu.
 - b. Select **Red Hat Enterprise Linux** with the version of the kernel that you want to boot, and type a to append the line.
 - c. Go to the end of the line, and type single as a separate word (press the **Spacebar** and then type single). Press **Enter** to exit edit mode.
2. Once in the single-user mode, verify the /var size by:
 - a. Running the **df -h** command.
 - b. If the /var directory is 100% full, move the files to /ftdc/core/, and reboot the file module.

Note: If you still get the 1710 bus error message, contact IBM support.

Note: IBM Software Upgrade Test Utility that is run prior to code upgrades, checks the / and /var partitions and displays the following warning or error message, if they are close to being full:

- For file system:
File system / is xx% full. Please clear some disk space.
- For /var:
File system /var is xx% full. Please clear some disk space.

In all the above cases, you get the actual disk percentage in place of xx.

Note:

- When the value of xx is 95, upgrade test utility reports an error.
- When the value of xx is 85, upgrade test utility reports a warning.

Control enclosure-related issues

This section covers the recovery procedures that involve control enclosure issues.

Recovering when file volumes come back online

Use this procedure to recover a file system after all the file volumes are back online after a repair or recovery action.

About this task

Each fix procedure that brings the file volumes back online also suggests that you run this procedure. This procedure checks that the file systems have also come back online.

Perform the following steps to check that the file systems are back online after their file volumes are back online following an outage.

Procedure

1. In the management GUI, check that all volumes are back online.
2. Go to **Monitoring > Events** and click the **Block** tab.
3. Run any **Next recommended action**.
4. When all volumes are back online, go to **Filesystems** in the management GUI.
5. If any file systems are not online, recover them by using the recover a GPFS file system procedure. See “Recovering a GPFS file system” on page 196.
6. If there are file systems that have not come back online, go to **Monitoring > Events** and click the **File** tab to fix any errors.
7. If there are any stale NFS handle errors for the offline file systems, follow the “Working with file modules that report a stale NFS file handle” on page 403.

Recovering when a file volume does not come back online

An offline volume can normally be fixed by performing the fix procedure for the event in the management GUI.

About this task

Procedure

To run the fix procedures, perform the following steps:

1. Log in to the Storwize V7000 Unified management GUI.
2. Go to **Monitoring > Events** and click the **Block** tab.
3. Run any **Next recommended action**.

Results

If the fix procedures do not bring back a file system volume online, contact your service provider for assistance.

Recovering from offline compressed volumes

Recovering from offline compressed volumes. Getting them back online.

When a Storwize V7000 storage pool (MDisk Group) runs out of space:

- Any volume that tries to expand (such as new data being written to a compressed volume) is taken offline.
- Taking a file volume offline takes the Network Shared Disk (NSD) offline because each NSD is made from one file volume.
- Taking meta data (NSD) offline takes the whole file system offline (but putting meta data in a compressed volume is not allowed).
- The file system is unmounted if it is offline for more than 30 seconds.
- This is different from the file system filling up, which causes the file system to enter read only mode.

There are two options to recover from this:

- Increase the storage pool capacity.
- Free the unusable blocks in the compressed volumes.

Table 118. Recovering from offline compressed volumes.

Scenario	Recovery Procedure	Who does it?
Storage pool warning (80% full)	Provision more storage to pool	You You (Storwize V7000 fix procedure)
Compression ratio wrong (file system still online)	Increase the storage to pool size Or Free the unusable blocks in the compressed volumes.	You (with help from this page) You (with help from IBM Remote Technical support)
Storage pool full (file system offline)	Provision more storage to pool	You (Storwize V7000 fix procedure)
Storage pool full (file system offline) No available storage	Borrow hot spare disks, bring file system online, free up space, shrink filesystem, return hot spare disks	You (with help fro IBM Remote Technical Support)

Increasing the Storage pool capacity

To increase storage pool capacity, add more RAID arrays to the storage pool using the management GUI.

Storage can be taken or borrowed from the block allocation to resolve out of space conditions in the file allocation. Point in time block copies are a good candidate for deletion.

Storwize V7000 Unified can virtualize external block storage controllers. If spare capacity is available on other block storage controllers then you can virtualize those and add the mdisks to the volume storage pool.

Free the unusable blocks in the compressed volumes

If you cannot increase the storage pool capacity then contact IBM Remote Technical Support to help you.

Recovering from a 1001 error code

A 1001 error code indicates that the Storwize V7000 control enclosure has automatically performed a recovery. The control enclosure CLI is now restricted to make sure that there are no more block storage configuration changes until IBM Remote Technical Support has checked that it is safe for block storage configuration changes to be allowed again.

About this task

The file volumes presented by the control enclosure for GPFS to use as disks for file systems may have been offline long enough to cause the file systems to be unmounted. The file systems will usually be automatically mounted as soon as the file volumes come back online after the control enclosure recovery. You can immediately remount any remaining unmounted file systems without waiting for IBM support to tell you that it is safe for you to re-enable the control enclosure CLI.

Note: The management GUI can become very slow when the control enclosure CLI is restricted, so the following procedure shows how to use the management CLI to check if the file systems are mounted. However, it is better to use the management GUI if that is working fine.

Procedure

To check if the file systems were automatically mounted following the recovery of the control enclosure:

1. Log on to the management CLI with your administrator credentials. For example:

```
ssh admin@<management_IP address>
```

2. Use the `lsnode -r` CLI command to check the status of CTDB and GPFS on each file module. For example:

```
lsnode -r
```

3. Use the `lsmount` CLI command to check if all of your file systems that should be mounted are mounted. For example:

```
[kd52v6h.ibm]$ lsmount
File system Mount status Last update
gpfs0      not mounted  10/17/12 10:44 AM
gpfs1      not mounted  10/17/12 10:44 AM
gpfs2      not mounted  10/17/12 10:44 AM
EFSSG1000I The command completed successfully.
```

4. If all of the required file systems are mounted on both nodes then there is no need to continue going through this procedure because network users should be able to access files on GPFS. Otherwise use the `lsdisk` CLI command to check if all of the disks are available. For example:

```
[kd52v6h.ibm]$ lsdisk
Name File system Failure group Type Pool Status Availability Timestamp
Block properties
IFS1350385068630 gpfs0 1 metadataOnly system ready up 10/17/12 10:27 AM
IFS1350385068630,io_grp0,,easytier,6005076802AD80227800000000000000
IFS1350385068806 gpfs0 1 metadataOnly system ready up 10/17/12 10:27 AM
IFS1350385068806,io_grp0,,easytier,6005076802AD80227800000000000001
IFS1350385089739 gpfs0 2 metadataOnly system ready up 10/17/12 10:27 AM
IFS1350385089739,io_grp0,,easytier,6005076802AD80227800000000000002
IFS1350385089889 gpfs0 2 metadataOnly system ready up 10/17/12 10:27 AM
IFS1350385089889,io_grp0,,easytier,6005076802AD80227800000000000003
IFS1350385108175 gpfs0 0 dataOnly system ready up 10/17/12 10:27 AM
IFS1350385108175,io_grp0,,easytier,6005076802AD80227800000000000004
```

5. If all disks are up then you can use the `mountfs` CLI command to mount each file system that is not mounted. For example

```
mountfs <file system name>
```

6. Otherwise if none of the disks are up or some of the disks are not up then use the `lsdisk` CLI command to check if all of your file volumes that should be online are online. Note that the names of file volumes are the same as the names of the disks. For example

```
[kd52v6h.ibm]$ lsvdisk
id name IO_group_id IO_group_name status mdisk_grp_id mdisk_grp_name capacity
type
FC_id FC_name RC_id RC_name vdisk_UID fc_map_count copy_count fast_write_state
0 IFS1350385068630 0 io_grp0 online 1 metal 100.00GB striped
6005076802AD80227800000000000000 0 1 not_empty
1 IFS1350385068806 0 io_grp0 online 1 metal 100.00GB striped
6005076802AD80227800000000000001 0 1 not_empty
2 IFS1350385089739 0 io_grp0 online 2 meta2 100.00GB striped
6005076802AD80227800000000000002 0 1 not_empty
3 IFS1350385089889 0 io_grp0 online 2 meta2 100.00GB striped
6005076802AD80227800000000000003 0 1 not_empty
4 IFS1350385108175 0 io_grp0 online 0 mdiskgrp0 341.00GB striped
6005076802AD80227800000000000004 0 1 not_empty
```

7. If any file volumes are offline then refer to Recovering when a file volume does not come back online.

8. If none of the disks are up but all file volumes are online then the multi pathing driver in the file modules may have failed and the best way to recover is to reboot the file modules one after the other using the procedure below.
9. If some of the disks are not up but the volumes are online then restart all disks used by a file system before you continue to mount it.
10. Use the `chdisk` CLI command to restart all disks used by the file system. For example:
`chdisk <comma separated list of disk names> --action start`
11. Use the `mountfs` CLI command to mount the file system. For example:
`mountfs <file system name>`

What to do next

Rebooting the file modules if none of the disks are up but all file volumes are online:

To reboot the file modules if the multi-pathing driver may have failed following a recovery of the control enclosure:

1. Identify the passive and the active management nodes from the Description column in the output from the CLI command:
`lsnode -r`
 Reboot the file module that is the passive management node using the CLI command:
`stopcluster -node <node name> -reboot`
2. Wait until both nodes show **OK** in the Connection status column of the output from the CLI command:
`lsnode -r`
3. Resume the file module back into the cluster using the CLI command:
`resumefile <node name>`
4. Reboot the file module that is the active management node using the CLI command. The active management node fails over to the file module that you rebooted first.
`stopcluster -node <node name> -reboot`
5. Log back on to the Storwize V7000 Unified CLI. Wait until both nodes show OK in the Connection status column of the output from the CLI command:
`lsnode -r`
6. Resume the file module back into the cluster using the CLI command:
`resumefile <node name>`
7. Then wait for GPFS to be active on both file modules in the output of the CLI command:
`lsnode -r`
8. Check that the file systems are mounted by using the `lsmount -r` management CLI command:
`lsmount -r`
9. See Checking the GPFS file system mount on each file module if any file systems are not mounted.

Note that the management GUI can become very slow when the Storwize V7000 CLI is restricted. When you log on to the management GUI, it issues a warning that the Storwize V7000 CLI is restricted. The management GUI runs the fix

procedure to direct you to send logs to IBM. The fix procedure directs you back to this procedure to make the file systems accessible again.

To collect the Storwize V7000 logs, select the **Collect Logs** option from the navigation in the service assistant. Choose the **With statesave** option.

The fix procedure re-enables the control enclosure CLI, provided that IBM support approved of this procedure.

After completing this procedure the health status indicator could still be red because the Fibre Channel links may not have sent an event showing that they have recovered. Refer to Connectivity issues to help you see if this is the case and refer to Health status and recovery to help you return the health status indicator back to green.

Restoring data

This section covers the recovery procedures that relate to restoring data.

Note: While the restore is in progress, do not unlink the fileset, unmount the file system, or delete the fileset, fileset snapshot, or file system.

Restoring asynchronous data

Recovering a file system with asynchronous replication requires that you configure and start a replication relationship from the target site to the source site.

Before you begin

After the source site (site A) has failed, set the target site (Site B) as the new source site and replicate back to Site A. To restore asynchronous data, perform the following steps:

Procedure

1. Where the previous replication relationship was Site A replicating to Site B, configure the asynchronous replication by reversing the source and target site information. Site B replicates to Site A. See “Configuring asynchronous replication” and transpose the source and target information.
2. Start the replication that was configured in step 1 by using the **startrepl -fullsync** CLI command. See “Starting and stopping asynchronous replication” for more information.
3. If the amount of data that is to be replicated back to Site A is large, multiple replications from Site B to Site A might be required. Multiple replications are required until modifications to Site B can be suspended to perform a final replication to Site A to enable Site A to synch up.

Note: Do not use the **fullsync** option for these incremental replications.

4. After you verify that the data on Site A has been replicated accurately, you can reconfigure Site A as the primary site. Remove any replication tasks from Site B to Site A by using the **rmtask** CLI command.

Restoring IBM Spectrum Protect data

The Storwize V7000 Unified system contains a IBM Spectrum Protect client that works with your IBM Spectrum Protect server system to perform high-speed data backup and recovery operations.

Before you begin

Before restoring a file system, determine whether a backup is running and when backups were completed. To restore the data, perform the following steps:

Procedure

1. Determine whether a backup is running and when backups were completed by running the **lsbackup** CLI command. Specify the file system.

For example, the command to display the gpfs0 file system backup listing shows the output in the following format: # lsbackup gpfs0 Filesystem Date Message gpfs0 20.01.2010 02:00:00.000 G0300IEFSSG0300I The filesystem gpfs0 backup started. gpfs0 19.01.2010 06:10:00.123 G0702IEFSSG0702I The filesystem gpfs0 backup was done successfully. gpfs0 15.01.2010 02:00:00.000 G0300IEFSSG0300I The filesystem gpfs0 backup started.

2. Restore the backup by using the **startrestore** CLI command. Specify a file system name pattern.

You cannot restore two file systems at the same time; therefore, the file pattern cannot match more than one file system name.

Use the **-t** option to specify a date and time in the format "dd.MM.yyyyHH:mm:ss.SSS" to restore files as they existed at that time. If a time is not specified, the most recently backed up versions are restored. For example, to restore the /ibm/gpfs0/temp/* file pattern to its backed up state as of January 19, 2010 at 12:45 PM, enter the following command:

```
# startrestore "/ibm/gpfs0/temp/*" -t "19.01.2010 12:45:00.000"
```

See the **startrestore** CLI command for additional command information, default options, and file pattern examples.

Attention: The **-R** option overwrites files and has the potential to overwrite newer files with older data.

3. Use the **lsbackupfs** CLI command to determine whether a restore is running. The **Message** field displays RESTORE_RUNNING if a restore is running on a file system.

4. Monitor the progress of the restore process by using the **QUERY SESSION** command in the IBM Spectrum Protect CLI client.

Run this command twice and compare the values in the Bytes Sent column of the output. Incremental values indicate that the process is in progress; whereas, identical values indicate that the restore process has stopped.

Note: The following error message can occur while restoring millions of files: ANS1030E The operating system refused a TSM request for memory allocation. 2010-07-09 15:51:54-05:00 dsmc return code: 12

What to do next

If the file system is managed by IBM Spectrum Protect for Space Management, break down the restore into smaller file patterns or subdirectories that contain fewer files.

If the file system is not managed by IBM Spectrum Protect for Space Management, try to force a no-query-restore (NQR) by altering the path that is specified for the restore. To do this action, include all files by putting a wildcard ("*") after the file system path:

```
# startrestore "ibm/gpfs0/*"
```

This example attempts a no query restore, which minimizes memory issues with the IBM Spectrum Protect client because the IBM Spectrum Protect server does the optimization of the file list. If you are still unable to restore a larger number of files at the same time, break down the restore into smaller file patterns or subdirectories that contain fewer files.

Upgrade recovery

This section covers the recovery procedures that relate to upgrade.

Error codes and recommendations when running the **applysoftware** command

If any errors are posted after you issue the **applysoftware** command, see Table 119 and take the described course of action. Follow these guidelines:

1. Follow the actions in the order presented.
2. After each recommended fix, restart the upgrade by issuing the **applysoftware** command again. If the action fails, try the next recommended action.
3. If the recommended actions fail to resolve the issue, call the IBM Support Center.

*Table 119. Upgrade error codes from using the **applysoftware** command and recommended actions*

Error Code	The applysoftware command explanation	Action
EFSSG1000I	The command completed successfully.	None.
EFSSG4100	The command completed successfully.	None.
EFSSG4101	The required parameter was not specified.	Check the command and verify that the parameters are entered correctly.
EFSSG4101A	The applysoftware command returned required parameter not specified.	
EFSSG4102	The software package does not exist.	Verify that the file actually exists where specified. Also verify that the command is passing the correct location parameters.
EFSSG4102A	The applysoftware command returned software package does not exist	
EFSSG4103	The software package is not valid.	The package might be corrupt. If this problem persists, download a new package and try again.
EFSSG4103A	The applysoftware command returned invalid software package return code.	
EFSSG4104	An unexpected return code.	Call your next level of support.

Table 119. Upgrade error codes from using the **applysoftware** command and recommended actions (continued)

Error Code	The applysoftware command explanation	Action
EFSSG4105	Unable to mount the USB flash drive.	Run <code>umount /media/usb</code> , then remove the USB flash drive. Reinsert the USB flash drive. If the error persists, remove the USB flash drive and reboot. After the system reboots, reinsert the USB flash drive.
EFSSG4105C	The applysoftware command returned unable to mount USB.	
EFSSG4106A	The applysoftware command returned that there is insufficient system file system space.	
EFSSG4153	The required parameter was not specified.	Verify that the file actually exists where specified. Also verify that the command is passing the correct location parameters.
EFSSG4154	You must start on primary management node <code>mgmt001st001</code> .	Switch to the other node and try the command again.
EFSSG4154A	The applysoftware returned must start on primary management node <code>mgmt001st001</code> .	
EFSSG4155	Unable to mount USB flash drive.	Back up to a USB flash drive. Enter <code># backupmanagementnode --unmount /media/usb</code> . Remove the USB flash drive and insert again. If the error persists, remove the USB flash drive and reboot. When the system is running, insert the USB flash drive again.
EFSSG4155I	The applysoftware command returned upgrade is already running.	
EFSSG4156	The specified International Organization for Standardization (ISO) does not exist.	Verify that the file actually exists where specified. Also verify that the command is passing the correct location parameters.
EFSSG4156A	The applysoftware command returned the specified ISO does not exist.	

Table 119. Upgrade error codes from using the **applysoftware** command and recommended actions (continued)

Error Code	The applysoftware command explanation	Action
EFSSG4157	The specific upgrade International Organization for Standardization (ISO) content is not valid.	The package might be corrupt. If this problem persists, download a new package and try again.
EFSSG4157I	The applysoftware command returned the specific upgrade ISO invalid content.	
EFSSG4158	The specific upgrade cannot be installed over the current version.	Check the upgrade documentation and verify that the level you are coming from is compatible with the level you are going to. If the upgrade level is not compatible, download the correct level and try again. If the upgrade level is compatible and the error persists, call the IBM Support Center.
EFSSG4158I	The applysoftware command returned the specific upgrade cannot be installed over the current version.	
EFSSG4159	The system is in an unhealthy state and the upgrade cannot start.	See Chapter 3, "Getting started troubleshooting," on page 47. Determine if the system has issues.
EFSSG4159I	The applysoftware command returned that the system is in an unhealthy state and upgrade cannot start.	
EFSSG4160	The system has insufficient file system space.	At least 3 GB of space is required. Remove unneeded files from the /var file system.
EFSSA0201C	The license agreement has not been accepted.	

General upgrade error codes and recommended actions

If any errors are posted during the upgrade process, see Table 120 on page 420 and take the described course of action. If the error you see is not listed in this table, call the IBM Support Center. Follow these guidelines:

1. Follow the actions in the order presented.
2. After each recommended fix, restart the upgrade by issuing the **applysoftware** command again. If the action fails, try the next recommended action.
3. If the recommended actions fail to resolve the issue, call the IBM Support Center.

Table 120. Upgrade error codes and recommended actions

Error Code	Explanation	Action
018C	Unable to determine active management node.	<ol style="list-style-type: none"> 1. Check to see if management service is running on active node. If it is not use startmgtsrv to start. 2. Contact IBM Remote Technical Support.
018E	Internal error - cluster or node not provided.	Contact IBM Remote Technical Support.
019A	Yum update failed.	Contact IBM Remote Technical Support.
019B	Unable to remove StartBackupTSM task.	<ol style="list-style-type: none"> 1. Check to see if management service is running on active node. If it is not, use startmgtsrv to start. 2. Contact IBM Remote Technical Support.
019C	Unable to determine active management node.	<ol style="list-style-type: none"> 1. Check to see if management service is running on active node. If it is not use startmgtsrv to start. 2. Contact IBM Remote Technical Support.
019D	Check the system health.	<ol style="list-style-type: none"> 1. Use lsnode to check the status of CTDB and GPFS for the nodes. Reboot the unhealthy node and wait for the node to be up again. Then, again check the health of the node with lsnode. 2. Contact IBM Remote Technical Support.
019E	Internal error - cluster or node not provided.	Contact IBM Remote Technical Support.
019F	CIM restart failed.	Contact IBM Remote Technical Support.
01A0	Failed to reboot.	<p>Determine the cause of the failed reboot:</p> <ol style="list-style-type: none"> 1. Check console of system if able. See if the system is hung in BIOS or during boot. 2. Check cabling of system. 3. Check light path diagnostic for error indications. . 4. Reboot the system from console and restart upgrade. 5. Contact IBM Remote Technical Support.
01A1	Internal upgrade error.	Contact IBM Remote Technical Support.
01A2	A GPFS command has failed.	Contact IBM Remote Technical Support.
01A3	Unable to uninstall CNCSM callbacks.	Contact IBM Remote Technical Support.

Table 120. Upgrade error codes and recommended actions (continued)

Error Code	Explanation	Action
01A4	Unable to stop backup jobs.	<ol style="list-style-type: none"> 1. Check the status of the backups by typing <code>lsjobstatus -j backup</code>. 2. Attempt to stop backups by typing <code>stopbackup --all</code>. 3. Contact IBM Remote Technical Support.
01A5	Backup cron jobs are running.	<ol style="list-style-type: none"> 1. Check the condition of tasks by typing <code>lstask -t cron</code>. 2. Attempt to remove the backup by typing <code>rmtask StartBackupTSM</code>. 3. Contact IBM Remote Technical Support.
01A6	Unable to install CNCSM callbacks.	Contact IBM Remote Technical Support.
01A7	Internal vital product data (VPD) error.	Contact IBM Remote Technical Support.
01A8	Check the health of management service.	<ol style="list-style-type: none"> 1. Attempt to start the management service with <code>startmgtsrv</code> on active management node 2. Contact IBM Remote Technical Support.
01A9	Unable to stop performance collection daemon.	Contact IBM Remote Technical Support.
01AB	Internal upgrade error in <code>node_setup_system</code> .	Contact IBM Remote Technical Support.
01B1	Management node replication failed.	<ol style="list-style-type: none"> 1. Follow the replication recovery procedure. See Resolving issues reported by <code>lshealth</code> for resolving the management node replication failure. 2. Contact IBM Remote Technical Support.
01B2	Unable to start performance collection daemon.	Contact IBM Remote Technical Support.
01B3	Failed to copy upgrade package to Storwize V7000.	This could be caused by a number of issues. Check Monitoring > Events under both the block tab and the file tab on the management GUI for an event that could have caused this error and follow the recommended action. If there is no obvious event that could have caused this error, refer to “Ethernet connectivity from file modules to the control enclosure” on page 68.

Table 120. Upgrade error codes and recommended actions (continued)

Error Code	Explanation	Action
01B4	Failed to start upgrade on Storwize V7000 with the applysoftware command.	This could be caused by a number of issues. Check Monitoring > Events under both the block tab and the file tab on the management GUI for an event that could have caused this error and follow the recommended action. If there is no obvious event that could have caused this error, refer to “Ethernet connectivity from file modules to the control enclosure” on page 68.
01B5	Storwize V7000 multipaths are unhealthy.	Check the Fibre Channel connections to the system. Reseat Fibre Channel cables. For more information, see Fibre Channel connectivity between file modules and control enclosure.
01B6	Storwize V7000 vdisks are unhealthy as indicated by using the lsvdisk command.	See Chapter 5, “Control enclosure,” on page 203.
01B7	Failed to query status of Storwize V7000 upgrade by using the svcinfolupdate command.	This could be caused by a number of issues. Check Monitoring > Events under both the block tab and the file tab on the management GUI for an event that could have caused this error and follow the recommended action. If there is no obvious event that could have caused this error, refer to “Ethernet connectivity from file modules to the control enclosure” on page 68.
01B8	Failed to query status of Storwize V7000 nodes by using the svcinfolnode command.	See Chapter 5, “Control enclosure,” on page 203.
01B9	Failed to check the Storwize V7000 version.	This could be caused by a number of issues. Check Monitoring > Events under both the block tab and the file tab on the management GUI for an event that could have caused this error and follow the recommended action. If there is no obvious event that could have caused this error, refer to “Ethernet connectivity from file modules to the control enclosure” on page 68.
01BA	Unable to verify the correct software version.	<ol style="list-style-type: none"> 1. Check the health of the storage controllers. 2. Contact IBM Remote Technical Support.
01BC	Check the health of storage controllers.	Contact IBM Remote Technical Support.
01BD	Unable to update software repository.	<ol style="list-style-type: none"> 1. Ensure that the system is not under a heavy load. Restart the upgrade. 2. Contact IBM Remote Technical Support.

Table 120. Upgrade error codes and recommended actions (continued)

Error Code	Explanation	Action
01BE	Unable to distribute upgrade callbacks.	<ol style="list-style-type: none"> 1. Check on health of the cluster using lshealth. 2. Contact IBM Remote Technical Support.
01BF	Upgrade callback failed	<ol style="list-style-type: none"> 1. Contact your customer advocate. Upgrade callbacks are custom steps placed on a system before the start of upgrade. 2. Contact IBM Remote Technical Support.
01C0	Asynchronous replication is running. Stop asynchronous replication and continue with the upgrade.	<ol style="list-style-type: none"> 1. Stop asynchronous replication by typing <code>stoprepl gpfs0 --kill</code>. Asynchronous replication is considered active if in RUNNING or KILLING state. 2. Contact IBM Remote Technical Support.
01C1	Asynchronous replication failed to stop. Stop asynchronous replication and continue with the upgrade.	<ol style="list-style-type: none"> 1. Stop asynchronous replication by typing <code>stoprepl gpfs0 --kill</code>. Asynchronous replication is considered active if in RUNNING or KILLING state. 2. Contact IBM Remote Technical Support.
01C2	Failed while checking for current running asynchronous jobs.	<ol style="list-style-type: none"> 1. Attempt to check status of <code>lsrepl</code>. If this command is working restart upgrade. 2. Contact IBM Remote Technical Support.
01C3	Could not stop CTDB.	Contact IBM Remote Technical Support.
01C4	Unable to remove callbacks	Contact IBM Remote Technical Support.
01C5	Could not reinstall Lib_Utils.	Contact IBM Remote Technical Support.
01C6	Failed while running <code>sonas_update_yum</code> .	Contact IBM Remote Technical Support.
01C7	Unable to get list of cluster nodes.	Contact IBM Remote Technical Support.
01C8	Failed while running <code>cnrssconfig</code> .	Contact IBM Remote Technical Support.
01C9	Unable to install CIM configuration.	Contact IBM Remote Technical Support.
01CA	Unable to get name of cluster.	Contact IBM Remote Technical Support.
01CB	Unable to install GPFS packages.	Contact IBM Remote Technical Support.
01CC	Could not install platform. Upgrade on target system.	Contact IBM Remote Technical Support.

Table 120. Upgrade error codes and recommended actions (continued)

Error Code	Explanation	Action
01CD	Unable to mount GPFS file systems.	<ol style="list-style-type: none"> 1. See “Checking the GPFS file system mount on each file module” on page 189 2. Restart upgrade and see if this was a transient issue. 3. Follow GPFS troubleshooting documentation. 4. Contact IBM Remote Technical Support.
01CE	Unable to update system security.	<ol style="list-style-type: none"> 1. Restart upgrade and see if this was a transient issue. 2. Contact IBM Remote Technical Support.
01CF	Unable to configure node.	<ol style="list-style-type: none"> 1. Pull both power supply cables from subject node. Wait 10 seconds, then plug back in. After the system restarts, try again. 2. Contact IBM Remote Technical Support.
01D0	Unable to disable call home.	Contact IBM Remote Technical Support.
01D1	Unable to enable call home.	Contact IBM Remote Technical Support.
01D2	Failed to stop GPFS.	<ol style="list-style-type: none"> 1. Follow GPFS troubleshooting documentation. 2. Contact IBM Remote Technical Support.
01D3	Could not determine if backups are running.	<ol style="list-style-type: none"> 1. Attempt to stop backups. 2. Type <code>lsjobstatus -j backup;echo \$?</code>. If the return code is 0, start the upgrade again. 3. If the return code is any other number, contact IBM Remote Technical Support.
01D5	Storwize V7000 stalled_non_redundant.	Refer to Storwize V7000 documentation.
01D6	Storwize V7000 System stalled.	Refer to Storwize V7000 documentation.
01D8	CTDB cluster is unhealthy.	<ol style="list-style-type: none"> 1. See “Checking CTDB health” on page 187. 2. Use <code>lshealth</code> or RAS procedures to determine unhealthy components. 3. Contact IBM Remote Technical Support.

Table 120. Upgrade error codes and recommended actions (continued)

Error Code	Explanation	Action
01DA	GPFS system is unhealthy.	<ol style="list-style-type: none"> 1. See “Checking the GPFS file system mount on each file module” on page 189. 2. Use lsnode -r to confirm GPFS is unhealthy. If node GPFS is healthy, restart the upgrade. 3. Contact IBM Remote Technical Support.
01DB	Failed to stop performance center.	Contact IBM Remote Technical Support.
01DC	Failed to configure performance center.	Contact IBM Remote Technical Support.
01DD	Failed to start performance center.	Contact IBM Remote Technical Support.
01DE	Unable to communicate with passive management node.	<ol style="list-style-type: none"> 1. Ensure that the active mgmt node can communicate with the passive management node before restarting the upgrade. 2. Contact IBM Remote Technical Support.
01DF	Upgrade must be resumed from the other management node.	Restart upgrade from other management node. This might require that a failover be issued first.
01E0	HSM upgrade failed.	Contact IBM Remote Technical Support.
01E1	mmchconfig Failed	Contact IBM Remote Technical Support.
01E2	mmauth Failed	Contact IBM Remote Technical Support.
01E3	mmlsfs Failed	Contact IBM Remote Technical Support.
01E3	mmchfs Failed	Contact IBM Remote Technical Support.
01E4	Disable HSM failed	Contact IBM Remote Technical Support.
01E5	Enable HSM failed	Contact IBM Remote Technical Support.
01E7	Unable to ping node	<ol style="list-style-type: none"> 1. Verify that the node is powered on. If yes, ping the node. 2. Restart the node and then, ping the node. 3. Check the network connections and correct them, if required. 4. Contact IBM Remote Technical Support.
01E8	Unable to apply firmware to Emulex adapters.	Contact IBM Remote Technical Support.

Table 120. Upgrade error codes and recommended actions (continued)

Error Code	Explanation	Action
01FF	NAS upgrade cannot continue because a system upgrade is in progress.	<p>Enter <code>svcinfolupdate</code> to check the state of the system upgrade.</p> <ul style="list-style-type: none"> • If the status is <code>system_updating</code> or <code>system_completing</code>, allow the upgrade to complete. Then restart the NAS upgrade from the file modules. • If the status is <code>system_prepared</code>, stop the current upgrade by entering <code>svctask applysoftware -abort</code>. Then restart the NAS upgrade from the file modules. <p>If <code>lupdate</code> is not available enter <code>svcinfolupdate</code> to check the state of the system upgrade.</p> <ul style="list-style-type: none"> • If the status is <code>upgrading</code>, allow the upgrade to complete. Then restart the NAS upgrade from the file modules. • If the status is <code>prepared</code>, stop the current upgrade by entering <code>svctask applysoftware -abort</code>. Then restart the NAS upgrade from the file modules.
0200	A failure has occurred while reading the registry.	<ol style="list-style-type: none"> 1. Use lsnode to check the status of CTDB and GPFS for the nodes. Reboot the unhealthy node and wait for the node to be up again. Then, again check the health of the node with lsnode. 2. Contact IBM Remote Technical Support.
0201	Unable to open CIM configuration file.	Contact IBM Remote Technical Support.
0202	Unable to update CIM configuration.	Contact IBM Remote Technical Support.
0511	Unable to disable CNSCM monitoring.	Contact IBM Remote Technical Support.
0513	Management service could not be stopped	Contact IBM Remote Technical Support.
0514	Database backup failed	Contact IBM Remote Technical Support.
0515	Database backup failed	Contact IBM Remote Technical Support.
0516	Database package install failed	Contact IBM Remote Technical Support.
0517	Database initialization failed	Contact IBM Remote Technical Support.
0518	Database service not running	Contact IBM Remote Technical Support.
0520	Database restore failed	Contact IBM Remote Technical Support.
0521	Database replication suspend or resume error.	Contact IBM Remote Technical Support.
0522	Unable to clean the CTDB configuration file.	Contact IBM Remote Technical Support.

Table 120. Upgrade error codes and recommended actions (continued)

Error Code	Explanation	Action
0523	Unable to upgrade Samba packages.	Contact IBM Remote Technical Support.

Chapter 8. Troubleshooting compressed file systems

To ensure that the capacity demands are not exceeded, the underlying block storage pools that provide the compression mechanism for file systems need to be monitored and maintained.

When a compressed pool is created, the administrator must estimate the compression ratio that will be achieved on that pool. Provided that the compression ratio is achieved or exceeded, there will be no issues. However, if the overall achieved compression ratio is lower than the predicted ratio, the actual capacity that is allocated will not be sufficient for storing the file system data.

When a block storage pool that is used for compressed file systems runs out of capacity, any compressed volume using that pool that expands is taken offline. If a volume used by a file system is offline for more than 30 seconds then the file system is unmounted, and all I/O to the file system fails. This behavior is different from uncompressed file systems. When a file system runs out of capacity, the file system enters read-only mode.

To ensure that capacity demands are met for a compressed file systems, monitor capacity usage for the block storage pools and volumes that provide the underlying compression mechanism for file systems. For details on setting thresholds and monitoring capacity for both the block storage and compressed file systems, see “Monitoring file system compression” on page 434

However, there can be cases where the storage demands of the data that is being compressed exceeds the capacity, and additional capacity needs to be added to the system. The following table provides an overview of typical recovery scenarios that are related to running out of capacity for compressed file systems.

Table 121. Capacity failure scenarios

Failure Scenario	Recovery Procedure
Storage pool warning indicates the pool is at the specified capacity threshold. The default threshold is 80%.	“Recovery procedure: Increase capacity of the storage pool” on page 430
Estimated compression savings for file system is not achieved. (File system is still online)	One of these options: <ol style="list-style-type: none">1. “Recovery procedure: Increase capacity of the storage pool” on page 4302. Free the unusable blocks in the compressed volumes. Note: Contact IBM Remote Technical Support or your service representative to complete this recovery procedure.
Storage pool is full and the file system pool is offline.	“Recovery procedure: Adding additional capacity for offline compressed file systems” on page 431
Storage pool is full and the file system pool is offline, but no additional storage is available to add to the pool.	Contact IBM Remote Technical Support or your service representative.

Recovery procedure: Increase capacity of the storage pool

If the allocated capacity of the block storage pool exceeds the specified capacity threshold then its compressed volumes can go offline. The default threshold is 80% of capacity; however, the value can be set lower or higher depending on the environment. If a compressed file volume is offline for 30 seconds, the file system is unmounted. Proper monitoring of storage pool thresholds is essential to ensure capacity consumption is not exceeding expectations. If the used capacity does exceed the specified threshold, you can recover by adding more storage to the block storage pool.

The most important metric to monitor is the physical capacity that is used in the storage pool. Make sure the physical allocation does not exceed the specified threshold. The default threshold is set at 80%. To reduce the current utilization of the used capacity, more physical capacity needs to be added to the storage pool or NSDs removed from the file system. To view the current level of utilization for block storage pools that are used for file system compression, select **Files > File Systems** and ensure that the **Storage pools** filter is selected. The management GUI displays all the file systems and their associated storage pools. Select the file system and expand the file system pool to display the block storage pool that is used for that file system. The Capacity column displays the current used capacity for the file and the underlying block storage pools. To view specific thresholds for individual volumes, select the **NSDs** filter to display the block volumes that are used in the file system. To view specific thresholds for individual volumes, right-click a volume and select **Properties**. In the upper right of the Properties panel, an allocation bar is displayed with the current threshold indicated by a red vertical bar.

Add any available MDisk: If an MDisk has already been created but not assigned to a pool, complete these steps:

1. In the management GUI, select **Pools > MDisk by Pools**.
2. Select **Not in pool** to display all the available MDisk that are not currently allocated to a storage pool.
3. Right-click the MDisk that you want to add to the storage pool and select **Add to Pool**.
4. On the **Add to Pool** dialog, select the pool and click **Add to Pool**.
5. Verify that the MDisk was added to the selected pool by expanding the pool and ensuring that the added MDisk is displayed.

Add any available drives: If MDisk have not been configured from available internal drives, you can provision the available drives into existing storage pools by completing these steps:

1. In the management GUI, select **Pools > Internal Storage**.
2. Select **Configure Storage**.
3. On the **Configure Internal Storage** dialog, select **Select a different configuration** and complete the following steps:
 - a. In the **Drive Class** field, select the drive class that is available based on the installed storage on the system.
 - b. In the **Preset** field, select the RAID configuration for the storage you are configuring.
 - c. Select **Optimize for capacity** to configure all available capacity.
 - d. Verify the configuration and click **Next**.

- e. Click **Expand an existing pool** and select the storage pool that is used for compression.
4. Click **Finish**.

Allocate storage from available external storage: The system supports adding external storage systems to provide additional capacity and virtualization. If your environment has external storage systems, you can increase capacity to the storage pool by completing these steps:

1. In the management GUI, select **Pools > External Storage**.
2. Select the storage system to view a list of MDisks that are currently detected on the external storage system. If there are no MDisks that are displayed, click **Detect MDisks**. If the Storwize V7000 Unified system attached to external storage systems, you can allocate additional LUNs.
3. Right-click an unmanaged MDisk and select **Add to Pool**.
4. On the **Add to Pool** dialog, select the pool and click **Add to Pool**.
5. Verify that the MDisk was added to the selected pool by expanding the pool and ensuring that the added MDisk is displayed.

Recovery procedure: Adding additional capacity for offline compressed file systems

In this situation, the storage pool has run out of capacity. As a result, the file system is unmounted and has gone offline, which makes all I/O to the file system fail.

To recover from this situation, you can either add available MDisks to the pool, or if free MDisks are not available, you can make spare drives available to build a new array (MDisk) to add to the pool. However, because spare drives are automatically used as backup drives when other drives fail on the system, using a spare drive to recover an offline file system can prevent an automated recovery if another drive fails on the system. After the file system is brought back online and capacity deficiencies have been addressed, return the drive to use as a spare or add another drive to replace it as a spare. If you add a new drive, new drives must be added to the system.

Increasing capacity to the storage pool

If MDisks are available to provide extra capacity to the storage pool that the compressed file system uses, you can add MDisks to the pool or create more MDisk (arrays).

Add any available MDisks: If an MDisk has already been created but not assigned to a pool, complete these steps:

1. In the management GUI, select **Pools > MDisk by Pools**.
2. Select **Not in pool** to display all the available MDisks that are not currently allocated to a storage pool.
3. Right-click the MDisks that you want to add to the storage pool and select **Add to Pool**.
4. On the **Add to Pool** dialog, select the pool and click **Add to Pool**.
5. Verify that the MDisk was added to the selected pool by expanding the pool and ensuring that the added MDisk is displayed.

Add any available drives: If MDisks have not been configured from available internal drives, you can provision the available drives into existing storage pools by completing these steps:

1. In the management GUI, select **Pools > Internal Storage**.
2. Select **Configure Storage**.
3. On the **Configure Internal Storage** dialog, select **Select a different configuration** and complete the following steps:
 - a. In the **Drive Class** field, select the drive class that is available based on the installed storage on the system.
 - b. In the **Preset** field, select the RAID configuration for the storage you are configuring.
 - c. Select **Optimize for capacity** to configure all available capacity.
 - d. Verify the configuration and click **Next**.
 - e. Click **Expand an existing pool** and select the storage pool that is used for compression.
4. Click **Finish**.

Using spare drives to add capacity to the storage pool

If drives are not available, you need to make spare drives available to add capacity to the storage pool, bring the file system back online, ensure capacity for the storage pool does not run out again, and return spare drives to the system.

Note: If you are unfamiliar with managing spare goals and spare disks, contact IBM support for guidance. Increasing capacity in this way is meant only as a short term solution to this problem. Further provisioning to permanently resolve capacity constraints can be conducted with the help of IBM service personnel who might recommend that additional drives be added to your system. To use spare drives to add capacity to the storage pool and bring file systems back online, complete these steps:

1. **Mark a spare drive as a candidate drive:** When block storage is configured on the system, available drives are categorized based on their drive class and drive type. To provide for drive redundancy, some drives are mark as spares, which provide backup drives in the event of a drive failure. Other drives are marked as candidates, which means they can be used as capacity for block storage pools. To mark a spare drive as a candidate and make it available to the block storage pool, complete these steps:
 - a. In the management GUI, select **Pools > Internal Storage**.
 - b. From the list of drives that display, right-click a drive that is marked as a spare drive and select **Mark as... > Candidate**.

Note: The **Use** column displays how a specific drive is used on the system.

- c. Click **OK**.
2. **Expand the storage pool:** After the spare drive has been marked as a candidate drive, you can expand the capacity of the block storage pool that is used for the offline file system.
 - a. In the management GUI, select **Pools > Internal Storage**.
 - b. Select **Configure Storage**.
 - c. On the **Configure Internal Storage** dialog, select **Select a different configuration** and complete the following steps:

- 1) In the **Drive Class** field, select the drive class of the candidate drive (former spare) that is available based on the installed storage on the system. Verify the correct number of disks is displayed.
 - 2) In the **Preset** field, select the RAID configuration for the storage you are configuring. If you are adding only one disk, the only RAID option is RAID0 which does not provide any data protection.
 - 3) Select **Optimize for capacity** to configure all available capacity.
 - 4) Verify the configuration and click **Next**.
 - 5) Click **Expand an existing pool** and select the storage pool that is used for compression.
3. **Check event logs to ensure all underlying volumes are back online.** Before bringing the file system back online ensure that all the errors for both the block volumes and the file system have been resolved by completing these steps:
 - a. In the management GUI, select **Monitoring > Events** and select **Block**.
 - b. Run the fix procedures in the recommended order for all events that are related to the block volume that is used by the file system.
 - c. Select **File** and fix all errors that are related to the offline file systems.
 4. **Bring file systems back online:** After the capacity has been added to the storage pool, bring the file system back online by completing these steps:
 - a. In the management GUI, select **Files > File Systems**.
 - b. Right-click the compressed file system that is offline and select **Mount**. If the file system does not come back online you may need to restart all of the disks that the file system uses Right-click the compressed file system that went offline and select **Start All Disks**.
 5. **Prevent the file system from running out of capacity again:**

First ensure that you have free capacity at least the size of the real capacity of the temporary drive that you are adding.

To decrease the file system capacity, you can remove the disks (NSD) and the corresponding mapping to block volumes to force migration of the data to other NSDs, thus freeing up space on the file system. To remove an NSD, contact IBM Remote Technical Support.
 6. **Return spare drives to the system:** To ensure that drive redundancy is not compromised, spare drives that were used to bring the offline file systems back online need be replaced either by returning the original drive back to its spare use or by adding a new drive to the system. Ensure that the file system capacity has been decreased accordingly before returning the spare drives to the systems. To return the drive back to its spare use, complete these steps:
 - a. In the management GUI, select **Pools > Internal Storage**.
 - b. From the list of drives that display, ensure that no MDisk are associated with the drive. If the drive is associated with MDisk, select **Pools > MDisk by Pool**. Right-click the MDisk and select **Remove from Pool**.
 - c. In the management GUI, select **Pools > Internal Storage**.
 - d. From the list of drives that display, right-click a drive you marked as a candidate in Step 1 and select **Mark as... > Spare**.
 - e. Click **OK**.

To add additional drives to the system, complete these steps:

 - a. Acquire additional drives from IBM or vendor.
 - b. Install drives into available drive slots on the enclosure. See "Installing a hot-swap hard disk drive" on page 149.
 - c. After the drives are available, select **Pools > Internal Storage**.

- d. From the list of drives that display, right-click the new drive and select **Mark as... > Spare**.

Monitoring file system compression

You can use the management GUI to monitor file and file system pool capacity metrics in a single view by selecting **Monitoring > Capacity** and **Files > File Systems > Storage pools**.

You can use two views to monitor the capacity usage on the system. Select **Monitoring > Capacity** to display a consolidated view of all information needed to monitor capacity-related information on the system. In addition, you can create alerts on capacity where you are notified when a specified capacity threshold has been reached for file system or storage pool capacity. The Capacity View shows system-wide compression savings and thin provisioning efficiency on storage pool level.

The most important metric to monitor is the physical capacity that is used in the storage pool. Make sure the physical allocation does not exceed the specified threshold. The default threshold is set at 80%. To reduce the current utilization of the used capacity, more physical capacity needs to be added to the storage pool or NSDs removed from the file system. To view the current level of utilization for block storage pools that are used for file system compression, select **Files > File Systems** and ensure that the **Storage pools** filter is selected. The management GUI displays all the file systems and their associated storage pools. Select the file system and expand the file system pool to display the block storage pool that is used for that file system. The Capacity column displays the current used capacity for the file and the underlying block storage pools. To view specific thresholds for individual volumes, select the **NSDs** filter to display the block volumes that are used in the file system. To view specific thresholds for individual volumes, right-click a volume and select **Properties**. In the upper right of the Properties panel, an allocation bar is displayed with the current threshold indicated by a red vertical bar.

Whenever a threshold is reached and an alert is issued, the system suggests actions that correspond to the specific scenario. If action is not taken and the storage pool reaches 100% utilization, volumes and their related network shared disks (NSDs) can go offline, which causes the file system to go offline. To see an overview of recovery scenarios, go to Chapter 8, "Troubleshooting compressed file systems," on page 429.

Theoretically, the total virtual capacity for all volumes in a pool can exceed the actual physical capacity that is available to the storage pool. For example, an administrator creates a 10 TB file system from a storage pool that has 10 TB of capacity. In this example, one volume is used and is allocated the full 10 TB of capacity to store this data. On average, the data that is stored in this file system has 60% compression savings. After the file system is full with 10 TB of data that gets 60% compression savings, it has actually used only 4 TB of physical capacity from the pool to store the compressed data. To use the remaining 6 TB of unused capacity, virtual capacity can be added for the volumes in the pool.

However, in reality, you need contingency capacity on the storage pool that remains unallocated and available to minimize impact to capacity utilization when data changes affects compression rates. In most cases, data does not have the same compression rate because it is constantly changing over the course of life cycle. Incompressible data or data that does not compress well can be added to a file

system, which impacts compression rates. The system default for the contingency threshold at 80% of the physical capacity which provides 20% contingency capacity for the storage pool, which is adequate for most environment. For example, if an administrator has a storage pool with 10 TB of physical storage and sets the threshold to 80%, only 8 TB out of the physical 10 TB are available in the pool. However, if the data in the pool receives 60% compression savings, the administrator can store approximately 20 TB of uncompressed user data in 8 TB of physical space. In this way, the maximum amount of virtual capacity exceeds the physical capacity for the compressed storage pool. To calculate the recommended virtual capacity, you can use the following equation:

Recommended maximum virtual capacity (in TB) = $(CT * PC) * (1 / (1 - CR))$

Contingency threshold (CT)

0.8 to represent 80% contingency threshold.

Physical capacity in TB (PC)

10 TB physical capacity that is available in the pool.

Compression savings (CR)

0.6 represents 60% compression savings.

File System Capacity Management

Additionally you must also monitor file capacity utilization to ensure that the file system does not reach 100% utilization and run out of capacity. The capacity utilization of a file system issued physical capacity in the compressed pool. The system uses the same threshold and alerting system and suggests corrective actions when thresholds are reached. If based on the original, uncompressed capacity that the system presents to users and applications of the file system.

To free up capacity in a file system, you can either remove NSDs from the file system or increase the current capacity of the storage pool, which can be used to expand the volumes that are related to the NSDs from the unused physical capacity. If corrective action to reduce utilization is not performed before the file system reaches 100% utilization, the file system goes offline and no longer handles read and write requests.

Appendix. Accessibility features for IBM Storwize V7000 Unified

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

Accessibility features

These are the major accessibility features for the Storwize V7000 Unified:

- You can use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. HTML documents have been tested using JAWS version 15.0.
- This product uses standard Windows navigation keys.
- Interfaces are commonly used by screen readers.
- Keys are discernible by touch, but do not activate just by touching them.
- Industry-standard devices, ports, and connectors.
- You can attach alternative input and output devices.

The Storwize V7000 Unified online documentation and its related publications are accessibility-enabled. The accessibility features of the online documentation are described in [Viewing information in the information center](#).

Keyboard navigation

You can use keys or key combinations to perform operations and initiate menu actions that can also be done through mouse actions. You can navigate the Storwize V7000 Unified online documentation from the keyboard by using the shortcut keys for your browser or screen-reader software. See your browser or screen-reader software Help for a list of shortcut keys that it supports.

IBM and accessibility

See the IBM Human Ability and Accessibility Center for more information about the commitment that IBM has to accessibility.

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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 grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.*

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

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "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 jurisdictions 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those

websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

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

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US*

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

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

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

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.

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

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

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.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application

programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

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

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.

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

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux and the Linux logo is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other product and service names might be trademarks of IBM or other companies.

Electronic emission notices

This section contains the electronic emission notices or statements for the United States and other countries.

Federal Communications Commission (FCC) statement

This explains the Federal Communications Commission's (FCC's) statement.

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, might cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors, or by

unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device might not cause harmful interference, and (2) this device must accept any interference received, including interference that might cause undesired operation.

Industry Canada compliance statement

This Class A digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe A est conform à la norme NMB-003 du Canada.

Australia and New Zealand Class A Statement

Attention: This is a Class A product. In a domestic environment this product might cause radio interference in which case the user might be required to take adequate measures.

European Union Electromagnetic Compatibility Directive

This product is in conformity with the protection requirements of European Union (EU) Council Directive 2004/108/EC on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM option cards.

Attention: This is an EN 55022 Class A product. In a domestic environment this product might cause radio interference in which case the user might be required to take adequate measures.

Responsible Manufacturer:

International Business Machines Corp.
New Orchard Road
Armonk, New York 10504
914-499-1900

European community contact:

IBM Deutschland GmbH
Technical Regulations, Department M372
IBM-Allee 1, 71139 Ehningen, Germany
Tele: +49 (0) 800 225 5423 or +49 (0) 180 331 3233
Email: halloibm@de.ibm.com

Germany Electromagnetic Compatibility Directive

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse A EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 2004/108/EG zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EU-Mitgliedsstaaten und hält die Grenzwerte der EN 55022 Klasse A ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung der IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung der IBM gesteckt/eingebaut werden.

EN 55022 Klasse A Geräte müssen mit folgendem Warnhinweis versehen werden:

“Warnung: Dieses ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funk-Störungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Maßnahmen zu ergreifen und dafür aufzukommen.”

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem “Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG).” Dies ist die Umsetzung der EU-Richtlinie 2004/108/EG in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) (bzw. der EMC EG Richtlinie 2004/108/EG) für Geräte der Klasse A

Dieses Gerät ist berechtigt, in übereinstimmung mit dem Deutschen EMVG das EG-Konformitätszeichen - CE - zu führen.

Verantwortlich für die Einhaltung der EMV Vorschriften ist der Hersteller:

International Business Machines Corp.
New Orchard Road
Armonk, New York 10504
Tel: 914-499-1900

Der verantwortliche Ansprechpartner des Herstellers in der EU ist:

IBM Deutschland GmbH
Technical Regulations, Abteilung M372
IBM-Allee 1, 71139 Ehningen, Germany
Tele: +49 (0) 800 225 5423 or +49 (0) 180 331 3233
Email: halloibm@de.ibm.com

Generelle Informationen:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55022 Klasse A.

People's Republic of China Class A Statement

中华人民共和国“A类”警告声明

声明

此为A级产品，在生活环境中，该产品可能会造成无线电干扰。在这种情况下，可能需要用户对其干扰采取切实可行的措施。

Taiwan Class A compliance statement

警告使用者：
這是甲類的資訊產品，在
居住的環境中使用時，可
能會造成射頻干擾，在這
種情況下，使用者會被要
求採取某些適當的對策。

taitemi

Taiwan Contact Information

This topic contains the product service contact information for Taiwan.

IBM Taiwan Product Service Contact Information:

IBM Taiwan Corporation

3F, No 7, Song Ren Rd., Taipei Taiwan

Tel: 0800-016-888

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

f2c00790

Japan VCCI Council Class A statement

This explains the Japan Voluntary Control Council for Interference (VCCI) statement.

この装置は、クラス A 情報技術装置です。この装置を家庭環境で使用する
と電波妨害を引き起こすことがあります。この場合には使用者が適切な対策
を講ずるよう要求されることがあります。

VCCI-A

Japan Electronics and Information Technology Industries Association Statement

This statement explains the Japan JIS C 61000-3-2 product wattage compliance.

(一社) 電子情報技術産業協会 高調波電流抑制対策実施
要領に基づく定格入力電力値：Knowledge Center を参照

This statement explains the Japan Electronics and Information Technology Industries Association (JEITA) statement for products less than or equal to 20 A per phase.

高周波電流規格 JIS C 61000-3-2 適合品

This statement explains the JEITA statement for products greater than 20 A, single phase.

高周波電流規格 JIS C 61000-3-2 準用品

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。

- ・回路分類 : 6（単相、PFC回路付）
- ・換算係数 : 0

This statement explains the JEITA statement for products greater than 20 A per phase, three-phase.

高周波電流規格 JIS C 61000-3-2 準用品

本装置は、「高圧又は特別高圧で受電する需要家の高調波抑制対策ガイドライン」対象機器（高調波発生機器）です。

- ・回路分類 : 5（3相、PFC回路付）
- ・換算係数 : 0

Korean Communications Commission Class A Statement

This explains the Korean Communications Commission (KCC) statement.

이 기기는 업무용(A급)으로 전자파적합기기로서 판매자 또는 사용자는 이 점을 주의하시기 바라며, 가정외의 지역에서 사용하는 것을 목적으로 합니다.

Russia Electromagnetic Interference Class A Statement

This statement explains the Russia Electromagnetic Interference (EMI) statement.

ВНИМАНИЕ! Настоящее изделие относится к классу А.
В жилых помещениях оно может создавать радиопомехи, для
снижения которых необходимы дополнительные меры



Printed in USA

GA32-1057-17

