

PRIMECLUSTER Global Disk Services

A decorative horizontal band with a blue background. It features abstract network graphics, including glowing nodes, connecting lines, and wireframe spheres, suggesting a global or data network theme.

Configuration and Administration Guide 4.2

Oracle Solaris

J2S1-6602-02ENZ0(02)
March 2015

Preface

This manual describes the setting up and managing of GDS (Global Disk Services) and GDS Snapshot, an optional product of GDS, discussing the supported functions.

Purpose

This manual aims to help users understand GDS and GDS Snapshot, including environment configuration, operation and maintenance.

Target Reader

This manual is devoted to all users who will operate and manage GDS and GDS Snapshot and programmers who will create applications running in GDS environments.

Organization

This manual consists of 6 chapters, appendices and a glossary. The contents are as follows:

"YES" in the Operator or Administrator column indicates who would benefit from reading that portion.

Headers	Contents	Operator	Administrator
Chapter 1 Functions	Explains the features of GDS and GDS Snapshot.	YES	YES
Chapter 2 Objects	Explains the objects used by GDS and GDS Snapshot.	YES	YES
Chapter 3 Starting and Exiting GDS Management View	Explains how to start and exit GDS Management and the required browser environment.	YES	YES
Chapter 4 Management View Screen Elements	Explains the contents of Management View.	YES	YES
Chapter 5 Operation	Explains the details about GDS operation such as operation flow, settings, maintenance and management.	YES	YES
Chapter 6 Backing Up and Restoring	Explains how to back up and restore data on disks managed by GDS.	YES	YES
Appendix A General Notes	Explains guidelines, general precautions, and configuration tips necessary for using GDS.	-	YES
Appendix B Log Viewing with Web-Based Admin View	For details, see the supplementary "Web-Based Admin View Operation Guide."	-	YES
Appendix C Web-Based Admin View Operating Environment Setting	For details, see the supplementary "Web-Based Admin View Operation Guide."	YES	YES
Appendix D Command Reference	Explains the commands available in GDS and GDS Snapshot.	-	YES
Appendix E GDS Messages	Explains the contents, possible causes, and resolutions for GDS messages that appear when setting or operating GDS and GDS Snapshot.	YES	YES

Headers	Contents	Operator	Administrator
Appendix F Troubleshooting	Explains resolutions for abnormality of objects and physical disks managed by GDS.	-	YES
Appendix G Frequently Asked Questions (FAQ)	Lists frequently asked questions regarding GDS and GDS Snapshot.	YES	YES
Glossary	Explains GDS and GDS Snapshot terminology. Please refer to it as necessary.	YES	YES

Related documentation

Please refer to the following documents according to need.

- PRIMECLUSTER Concepts Guide
- PRIMECLUSTER Installation and Administration Guide
- PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide
- PRIMECLUSTER Reliant Monitor Services (RMS) with Wizard Tools Configuration and Administration Guide
- PRIMECLUSTER Web-Based Admin View Operation Guide
- PRIMECLUSTER Global File Services Configuration and Administration Guide
- INSTALLATION GUIDE PRIMECLUSTER(TM) GDS
- INSTALLATION GUIDE PRIMECLUSTER(TM) GDS Snapshot

Manual Printing

Use the PDF file to print this manual.

Adobe Reader is required to read and print this PDF file. To get Adobe Reader, see Adobe Systems Incorporated's website.

Symbol

The following conventions are used in this manual:



Point

Main points are explained.



Note

Items that require attention are explained.



Information

Useful information is given.



See

Manual names and sections of this manual you should refer to are given.

Abbreviated name

- ETERNUS SF AdvancedCopy Manager is abbreviated as ACM.
- Oracle Solaris might be described as Solaris, Solaris Operating System, or Solaris OS.
- Oracle Solaris 9 is abbreviated as Solaris 9.
- Oracle Solaris 10 is abbreviated as Solaris 10.

Export Controls

Exportation/release of this document may require necessary procedures in accordance with the regulations of your resident country and/or US export control laws.

Trademarks

Microsoft and Windows are registered trademarks of Microsoft Corporation, in the United States and other countries.

Window images are used according to the guidelines given by Microsoft Corporation.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

EMC, PowerPath, SRDF, Symmetrix and TimeFinder are registered trademarks of EMC Corporation.

SAN Manager is a trademark of EMC Corporation.

PRIMECLUSTER is a registered trademark of Fujitsu LIMITED.

Other product names that appear in this manual are product names, trademarks, or registered trademarks of respective companies.

Notice

- No part of this manual may be reproduced without permission.
- This manual is subject to change without advance notice.

Date of publication and edition

May 2006, First edition
June 2009, Second edition
October 2011, 2.1 edition
March 2015, 2.2 edition

All Rights Reserved, Copyright(C) 2006-2015 Fujitsu LIMITED.

Revision History

Revision	Location	Edition
Described the system boot in case of a disk failure.	1.2.1.1 System Disk Mirroring	2.1

Revision	Location	Edition
Modified description of disk array units.	1.5.3 Instant Snapshot by OPC A.2.19 Instant Snapshot by OPC	
Added explanation of "Pop-up menu."	4.1 Screen Configuration	
Added Note.	4.2.1 General 5.1.4.2 Backup (by Synchronization) 5.1.7 Disk Migration 5.3.4 Disk Swap 5.5.1 Removing a File System 5.5.5 Unmirroring the System Disk A.2.6 Disk Size F.1.3 Volume Status Abnormality	
Modified description of "Update Physical Disk Information."	4.2.3 Operation	
Added explanation of "Check Status."	4.2.3 Operation	
Modified the mirroring conditions of the system disk.	5.2.1 System Disk Settings	
Changed the title.	5.3.1 Viewing Configurations/Statuses and Monitoring Statuses A.2.31 Swapping Internal Disks (FC-AL)	
Added "5.3.1.4 Viewing Object Statuses."	5.3.1 Viewing Configurations/Statuses and Monitoring Statuses	
Added the description in Information.	6.1.1 Backing Up 6.1.2 Restoring (When the System Can Be Booted)	
Modified the description in Information.	6.1.1 Backing Up 6.1.2 Restoring (When the System Can Be Booted)	
Modified the description in See.	6.1.1 Backing Up 6.1.2 Restoring (When the System Can Be Booted) 6.1.3 Restoring (When the System Cannot Be Booted) F.1.11 Disk Unit Error	
Deleted "5) Create a special file for a tape device."	6.1.2 Restoring (When the System Can Be Booted)	
Modified Notes.	6.1.2 Restoring (When the System Can Be Booted) 6.1.3 Restoring (When the System Cannot Be Booted) A.2.9 System Disk Mirroring D.20 Volume Expansion Using Commands	
Modified "1) Boot the system from the CD-ROM unit (or network)."	6.1.3 Restoring (When the System Cannot Be Booted)	
Deleted "2) Create a special file for a tape device."	6.1.3 Restoring (When the System Cannot Be Booted)	
Added "3-5) Restart the volume on the node where the restore was performed."	6.3.3 Restoring	
Modified the execution example in "11) Deleting the class."	6.5.3 Restoring from Backup Disks	
Classified the preconditions into three.	A.1.8 Proxy Configuration Preconditions	

Revision	Location	Edition
Added Notes. - "A.2.32 Swapping Internal Disks (SPARC T3-1/T3-2/T3-4)." - "A2.34 Emergency Startup Disk for Restoration"	A.2 Important Points	
Added "Required time for Synchronization Copying Invoked by Hot Spare."	A.2.8 Hot Spare	
Modified the conditions for mirroring while holding the data.	A.2.10 Keep Disk	
Modified the conditions when the Advanced Copy Function is not used.	A.2.18 Using the Advanced Copy Function in a Proxy Configuration	
Modified the lead sentence.	A.2.20 To Use EMC Symmetrix	
Modified "Relate" of "PRIMARY OPTIONS."	D.14 sdxproxy - Proxy object operations	
Added the description in Information.	D.20 Volume Expansion Using Commands	
Modified the resolution of messages. - 44028 - 60039 - 60257	E.3.3 Warning Messages (44000-44099) E.4.1 Error Messages (60000-60499)	
Modified description of the message and resolution. - 60057	E.4.1 Error Messages (60000-60499)	
Added messages. - 60368 - 60369 - 60372	E.4.1 Error Messages (60000-60499)	
Modified the message. - 5014	E.5.1 Error Messages (5000-5099)	
Modified Resolution of "(2) The copy destination slice was made INVALID due to an I/O error generated on the copy source slice during synchronization copying."	F.1.1 Slice Status Abnormality	
Added description to Resolution c) in "(1) Disk is in DISABLE status."	F.1.2 Disk Status Abnormality	
Modified the title of Resolution (3).	F.1.6 GDS Management View Abnormality	
Modified the cause and resolution in "(1) The Advanced Copy function cannot be used in master-proxy copying."	F.1.7 Proxy Object Abnormality	
Modified the resolution in '(1) The error message "ERROR: class: cannot operate in cluster environment, ..." is output, and the operation cannot be conducted on the class class.'	F.1.9 Cluster System Related Error	
Added description for the -d option.	F.2.2 Collecting Initial Investigation Material with sdxsnap.sh Script	
Deleted terms. - ETERNUS3000 series	Glossary	

Revision	Location	Edition
- ETERNUS6000 series		
Added the error message 60418.	E.4.1 Error Messages (60000-60499)	2.2

Contents

Chapter 1 Functions.....	1
1.1 GDS Features.....	1
1.2 Functions for High Availability.....	2
1.2.1 Disk Mirroring.....	3
1.2.1.1 System Disk Mirroring.....	3
1.2.1.2 Mirroring between Disk Array Unit.....	4
1.2.1.3 Shared Disk Mirroring.....	5
1.2.2 Hot Spare.....	6
1.2.3 Hot Swap.....	8
1.2.4 Just Resynchronization Mechanism (JRM).....	9
1.3 Functions for High Manageability.....	10
1.3.1 Operation Management Interface.....	10
1.3.2 Centralized Disk Management.....	11
1.3.3 Name Management.....	12
1.3.4 Single System Image Environment.....	13
1.3.5 Access Control.....	14
1.3.6 Realizing Large Capacity and I/O Load Balancing.....	16
1.3.6.1 Logical Partitioning.....	16
1.3.6.2 Disk Concatenation.....	16
1.3.6.3 Disk Striping.....	17
1.3.6.4 Combining Disk Striping with Mirroring.....	18
1.3.7 Online Volume Expansion.....	18
1.3.8 Snapshots by Slice Detachment.....	20
1.4 GDS Snapshot Features.....	22
1.5 Proxy Volume.....	22
1.5.1 Snapshot by Synchronization.....	24
1.5.2 Snapshot Function with No Load to the Server/SAN.....	28
1.5.3 Instant Snapshot by OPC.....	29
1.5.4 Instant Restore.....	30
1.5.5 Online Disk Migration.....	32
1.5.6 Creating an Alternative Boot Environment.....	33
1.6 Shadow Volume.....	34
1.6.1 Accessing Volumes from an External Server.....	34
1.6.2 Using the Copy Function of a Disk Unit.....	35
Chapter 2 Objects.....	37
2.1 SDX Object.....	37
2.2 Disk Class.....	37
2.3 SDX Disk.....	40
2.4 Disk Group.....	43
2.5 Logical Volume.....	45
2.6 Logical Slice.....	50
2.7 GDS Snapshot Objects.....	54
2.7.1 Proxy Object.....	54
2.7.2 Shadow Object.....	58
2.7.2.1 Shadow Class.....	59
2.7.2.2 Shadow Disk.....	61
2.7.2.3 Shadow Group.....	63
2.7.2.4 Shadow Volume.....	65
2.7.2.5 Shadow Slice.....	70
Chapter 3 Starting and Exiting GDS Management View.....	71
3.1 Preparation for Starting GDS Management View.....	71
3.1.1 Deciding the User Group.....	71
3.1.1.1 User Group Types.....	71

3.1.1.2 Creating User Groups.....	71
3.1.1.3 Registering to a User Group.....	71
3.1.2 Setting up the Client Environment.....	72
3.1.3 Setting up the Web Environment.....	72
3.2 Starting the GDS Management View.....	72
3.2.1 Starting Web-Based Admin View Operation Menu.....	72
3.2.2 Web-Based Admin View Operation Menu Functions.....	73
3.2.3 Starting GDS Management View.....	74
3.3 Exiting GDS Management View.....	75
3.4 Changing the Web-Based Admin View Settings.....	76
Chapter 4 Management View Screen Elements.....	77
4.1 Screen Configuration.....	77
4.2 Menu Configuration and Functions.....	80
4.2.1 General.....	80
4.2.2 Settings.....	81
4.2.3 Operation.....	89
4.2.4 View.....	91
4.2.5 Help.....	92
4.3 Icon Types and Object Status.....	92
4.4 Object Information.....	96
Chapter 5 Operation.....	98
5.1 Operation Outline.....	98
5.1.1 System Disk Settings.....	98
5.1.2 Automatic Configuration.....	99
5.1.3 Configuration Settings.....	99
5.1.3.1 Single Volume Configuration Settings.....	99
5.1.3.2 Other Volume Configuration Settings.....	100
5.1.4 Backup.....	101
5.1.4.1 Backup (by Slice Detachment).....	101
5.1.4.2 Backup (by Synchronization).....	102
5.1.4.3 Backup (by OPC).....	104
5.1.5 Restore.....	104
5.1.6 Disk Swap.....	105
5.1.7 Disk Migration.....	106
5.1.8 Configuration Change.....	106
5.1.9 Unmirroring the System Disk.....	107
5.1.10 Operations from GDS Management View.....	108
5.2 Settings.....	111
5.2.1 System Disk Settings.....	111
5.2.2 Automatic Configuration Wizard.....	118
5.2.3 Operating from the Settings Menu.....	123
5.2.3.1 Class Configuration.....	123
5.2.3.2 Cluster System Class Configuration.....	127
5.2.3.3 Group Configuration.....	128
5.2.3.4 Volume Configuration.....	131
5.2.4 File System Configuration.....	134
5.2.5 Proxy Configuration.....	136
5.2.5.1 Join.....	136
5.2.5.2 Relate.....	141
5.3 Operation in Use.....	144
5.3.1 Viewing Configurations/Statures and Monitoring Statures.....	144
5.3.1.1 Confirming SDX Object Configuration.....	144
5.3.1.2 Viewing Proxy Object Configurations.....	150
5.3.1.3 Monitoring Object Status.....	155
5.3.1.4 Viewing Object Statures.....	157
5.3.2 Backup.....	158

5.3.2.1 Backup (by Slice Detachment).....	158
5.3.2.2 Backup (by Synchronization).....	165
5.3.2.3 Backup (by OPC).....	170
5.3.3 Restore.....	173
5.3.4 Disk Swap.....	177
5.3.5 Disk Migration.....	182
5.3.6 Copying Operation.....	185
5.4 Changes.....	189
5.4.1 Class Configuration.....	189
5.4.2 Group Configuration.....	194
5.4.3 Volume Configuration.....	197
5.5 Removals.....	198
5.5.1 Removing a File System.....	198
5.5.2 Removing a Volume.....	199
5.5.3 Removing a Group.....	201
5.5.4 Removing a Class.....	202
5.5.5 Unmirroring the System Disk.....	203
5.5.6 Breaking a Proxy.....	205
Chapter 6 Backing Up and Restoring.....	208
6.1 Backing Up and Restoring a System Disk.....	208
6.1.1 Backing Up.....	208
6.1.2 Restoring (When the System Can Be Booted).....	212
6.1.3 Restoring (When the System Cannot Be Booted).....	216
6.2 Backing Up and Restoring a System Disk through an Alternative Boot Environment.....	221
6.2.1 System Configuration.....	222
6.2.2 Summary of Backup.....	223
6.2.3 Summary of Restore.....	224
6.2.4 Summary of Procedure.....	226
6.2.5 Configuring an Environment.....	226
6.2.6 Backing Up.....	229
6.2.7 Restoring.....	234
6.3 Backing Up and Restoring Local Disks and Shared Disks.....	239
6.3.1 Offline Backup.....	239
6.3.2 Online Backup (by Slice Detachment).....	241
6.3.3 Restoring.....	243
6.4 Online Backup and Instant Restore through Proxy Volume.....	244
6.4.1 Online Backup (by Synchronization).....	246
6.4.2 Online Backup (Snapshot by OPC).....	254
6.4.3 Instant Restore.....	260
6.5 Backing Up and Restoring through Disk Unit's Copy Functions.....	265
6.5.1 Configuring an Environment.....	265
6.5.2 Backing Up.....	266
6.5.3 Restoring from Backup Disks.....	270
6.6 Backing Up and Restoring through an External Server.....	272
6.6.1 Backing Up and Restoring a Logical Volume with No Replication.....	274
6.6.1.1 System Configuration.....	275
6.6.1.2 Summary of Backup.....	276
6.6.1.3 Summary of Restore.....	277
6.6.1.4 Summary of Procedure.....	278
6.6.1.5 Configuring an Environment.....	278
6.6.1.6 Backing Up.....	279
6.6.1.7 Restoring.....	284
6.6.2 Backing Up and Restoring through Snapshot by Slice Detachment.....	290
6.6.2.1 System Configuration.....	291
6.6.2.2 Summary of Backup.....	292
6.6.2.3 Summary of Restore.....	293

6.6.2.4 Summary of Procedure.....	294
6.6.2.5 Configuring an Environment.....	295
6.6.2.6 Backing Up.....	296
6.6.2.7 Restoring.....	302
6.6.3 Backing Up and Restoring Using Snapshots from a Proxy Volume.....	309
6.6.3.1 System Configuration.....	309
6.6.3.2 Summary of Backup.....	310
6.6.3.3 Summary of Restore from a Proxy Volume.....	311
6.6.3.4 Summary of Restore from Tape.....	312
6.6.3.5 Summary of Procedure.....	314
6.6.3.6 Configuring an Environment.....	315
6.6.3.7 Backing Up.....	316
6.6.3.8 Restoring from a Proxy Volume.....	321
6.6.3.9 Restoring from Tape.....	322
6.6.4 Backing Up and Restoring by the Disk Unit's Copy Function.....	327
6.6.4.1 System Configuration.....	328
6.6.4.2 Summary of Backup.....	329
6.6.4.3 Summary of Restore from a BCV.....	331
6.6.4.4 Summary of Restore from Tape.....	332
6.6.4.5 Summary of Procedure.....	334
6.6.4.6 Configuring an Environment.....	335
6.6.4.7 Backing Up.....	336
6.6.4.8 Restoring form a BCV.....	343
6.6.4.9 Restoring from Tape.....	344
6.7 Backing Up and Restoring Object Configurations.....	351
6.7.1 Backing Up.....	351
6.7.2 Restoring.....	352
Appendix A General Notes.....	354
A.1 Rules.....	354
A.1.1 Object Name.....	354
A.1.2 Number of Classes.....	355
A.1.3 Number of Disks.....	355
A.1.4 Number of Groups.....	356
A.1.5 Number of Volumes.....	356
A.1.6 Number of Keep Disks.....	356
A.1.7 Creating Group Hierarchy.....	356
A.1.8 Proxy Configuration Preconditions.....	358
A.1.9 Number of Proxy Volumes.....	358
A.1.10 Proxy Volume Size.....	358
A.1.11 Proxy Group Size.....	359
A.2 Important Points.....	359
A.2.1 Physical Disk Driver.....	359
A.2.2 Managing System Disks.....	360
A.2.3 Restraining Access to Physical Special File.....	360
A.2.4 Booting from a CD-ROM Device (or Network).....	361
A.2.5 Initializing Disk.....	361
A.2.6 Disk Size.....	361
A.2.7 Volume Size.....	362
A.2.8 Hot Spare.....	363
A.2.9 System Disk Mirroring.....	364
A.2.10 Keep Disk.....	365
A.2.11 GFS Partition on a Keep Disk.....	366
A.2.12 Creating a Snapshot by Slice Detachment.....	366
A.2.13 The Difference between a Mirror Slice and a Proxy Volume.....	366
A.2.14 Just Resynchronization Mechanism (JRM).....	366
A.2.15 Online Volume Expansion.....	367

A.2.16 Swapping Physical Disks.....	369
A.2.17 Object Operation When Using Proxy.....	370
A.2.18 Using the Advanced Copy Function in a Proxy Configuration.....	370
A.2.19 Instant Snapshot by OPC.....	371
A.2.20 To Use EMC Symmetrix.....	372
A.2.21 Using EMC TimeFinder or EMC SRDF in a Proxy Configuration.....	376
A.2.22 Ensuring Consistency of Snapshot Data.....	377
A.2.23 Data Consistency at the time of Simultaneous Access.....	378
A.2.24 Volume Access Mode.....	378
A.2.25 Operation in Cluster System.....	378
A.2.26 Changing Over from Single Nodes to a Cluster System.....	379
A.2.27 Disk Switch.....	379
A.2.28 Shadow Volume.....	379
A.2.29 Configuration Tables.....	382
A.2.30 GDS Management View.....	382
A.2.31 Swapping Internal Disks (FC-AL).....	383
A.2.32 Swapping Internal Disks (SPARC T3-1/T3-2/T3-4).....	387
A.2.33 Use of GDS in Solaris 10 Environments.....	392
A.2.34 Emergency Startup Disk for Restoration.....	393
A.3 General Points.....	393
A.3.1 Guidelines for Mirroring.....	393
A.3.2 Guidelines for Striping.....	394
A.3.3 Guidelines for Concatenation.....	394
A.3.4 Guidelines for Combining Striping with Mirroring.....	395
Appendix B Log Viewing with Web-Based Admin View.....	396
Appendix C Web-Based Admin View Operating Environment Setting.....	397
Appendix D Command Reference.....	398
D.1 sdxclass - Class operations.....	398
D.2 sdxdisk - Disk operations.....	399
D.3 sdxgroup - Group operations.....	408
D.4 sdxvolume - Volume operations.....	412
D.5 sdxslice - Slice operations.....	417
D.6 sdxinfo - Display object configuration and status information.....	420
D.7 sdxattr - Set objects attributes.....	435
D.8 sdxswap - Swap disk.....	441
D.9 sdxfix - Restore a failed object.....	443
D.10 sdxcopy - Synchronization copying operation.....	445
D.11 sdxroot - Root file system mirroring definition and cancellation.....	447
D.12 sdxparam - Configuration parameter operations.....	450
D.13 sdxconfig - Object configuration operations.....	451
D.14 sdxproxy - Proxy object operations.....	455
D.15 sdxshadowdisk - Shadow disk operations.....	465
D.16 sdxshadowgroup - Shadow group operations.....	470
D.17 sdxshadowvolume - Shadow volume operations.....	474
D.18 Volume Creation Using Command.....	477
D.19 Snapshot Creation Using Command.....	481
D.20 Volume Expansion Using Commands.....	482
Appendix E GDS Messages.....	490
E.1 Web-Based Admin View Messages (0001-0099).....	490
E.2 Driver Messages.....	490
E.2.1 Panic Messages (20000).....	491
E.2.2 Warning Messages (22000-22099).....	491
E.2.3 Information Messages (24000-24099).....	495
E.2.4 Internal Error Messages (26000).....	498

E.3 Daemon Messages.....	498
E.3.1 Halt Messages (40000-40099).....	500
E.3.2 Error Messages (42000-42099).....	502
E.3.3 Warning Messages (44000-44099).....	515
E.3.4 Information Messages (46000-46199).....	525
E.3.5 Internal Error Messages (48000).....	544
E.4 Command Messages.....	545
E.4.1 Error Messages (60000-60499).....	546
E.4.2 Warning Messages (62000-62099).....	623
E.4.3 Information Messages (64000-64099).....	628
E.4.4 Fix Messages (66000).....	643
E.4.5 Internal Error Messages (68000).....	643
E.5 Operation Management View Messages.....	644
E.5.1 Error Messages (5000-5099).....	644
E.5.2 Warning Messages (5000,5100-5199).....	649
E.5.3 Information Messages (5200-5299).....	660
E.6 Error Codes (90000-90099).....	669
Appendix F Troubleshooting.....	671
F.1 Resolving Problems.....	671
F.1.1 Slice Status Abnormality.....	671
F.1.2 Disk Status Abnormality.....	678
F.1.3 Volume Status Abnormality.....	680
F.1.4 Class Status Abnormality.....	696
F.1.5 System Disk Abnormality.....	698
F.1.6 GDS Management View Abnormality.....	704
F.1.7 Proxy Object Abnormality.....	705
F.1.8 EMC Symmetrix Abnormality.....	707
F.1.9 Cluster System Related Error.....	711
F.1.10 Shadow Object Errors.....	717
F.1.11 Disk Unit Error.....	717
F.2 Collecting Investigation Material.....	718
F.2.1 Collecting with fjsnap Command.....	718
F.2.2 Collecting Initial Investigation Material with sdxsnap.sh Script.....	719
Appendix G Frequently Asked Questions (FAQ).....	720
G.1 Operation Design.....	720
G.2 Environment Configuration.....	723
G.3 Operation.....	724
Glossary.....	727

Chapter 1 Functions

This chapter describes the features and functions of GDS (Global Disk Services) and GDS Snapshot.

1.1 GDS Features

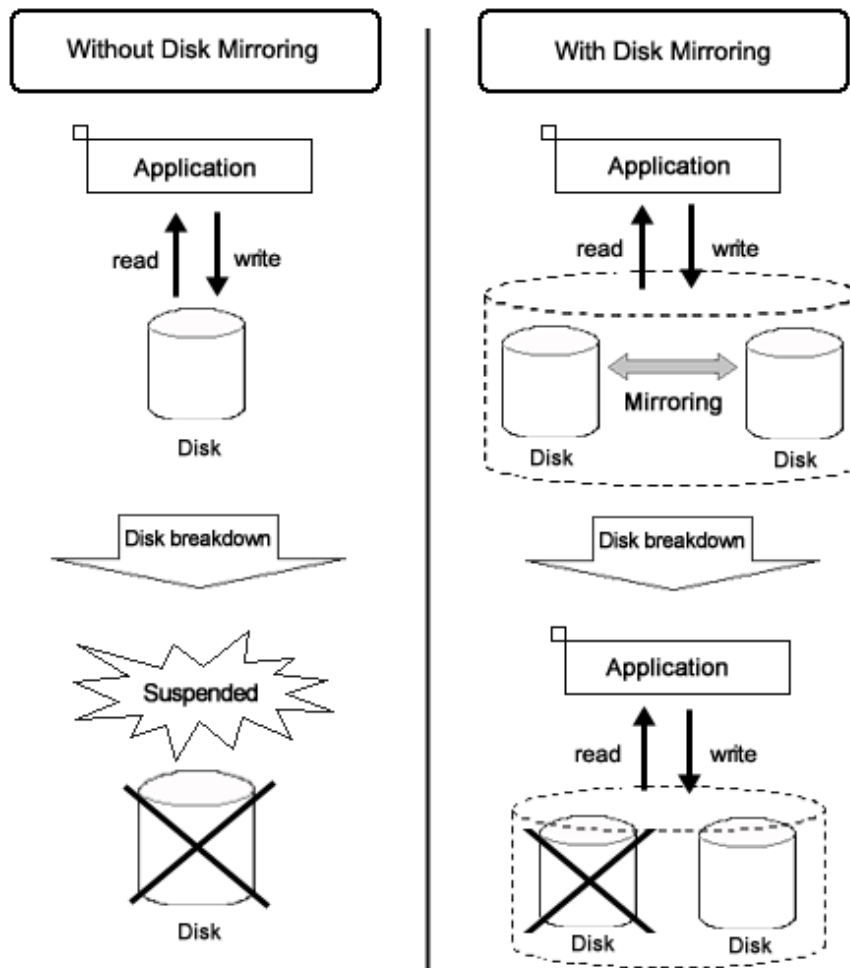
GDS is volume management software that improves the availability and manageability of disk-stored data. GDS protects disk data from hardware failures and operational mistakes and supports the management of disk units.

GDS has the following functions, which are closely related:

- To improve availability of disk data
- To improve manageability of disk data


GDS's mirroring function protects data from hardware failures by maintaining replicas of disk data on multiple disks. This allows users to continue to access disk data without stopping the application in the event of unexpected trouble.

Figure 1.1 Disk Mirroring



GDS allows users to integrate management of all disk units connected to a Solaris server. In a PRIMECLUSTER system, GDS also allows users to integrate management of disk units that are shared by multiple servers in a SAN (Storage Area Network) environment as well as local disk units that are connected to specific servers.

Figure 1.2 SAN (Storage Area Network)




Multiple servers can be connected to multiple disk units in a SAN environment. Disk-stored data can be accessed from those servers. This allows simultaneous access to file systems or databases and improves the efficiency of data duplication between the servers and backup procedures. On the other hand, it also carries the risk of data damage, as multiple servers will compete to access the shared disk. Therefore, volume management functions suitable for the SAN environment are essential.

Since GDS's management function is suitable for a SAN environment, advanced system operation for disk management can be performed with ease.

The user-friendly functions simplify management, and at the same time, prevent data corruption by operational mistakes.

Figure 1.3 Access Control Mechanism in a SAN Environment



1.2 Functions for High Availability

Failure of hardware related to disk unit interrupts data access, causing applications or systems to stop.

In this section, functions which protect data from unexpected hardware problems, and improve system availability are explained.

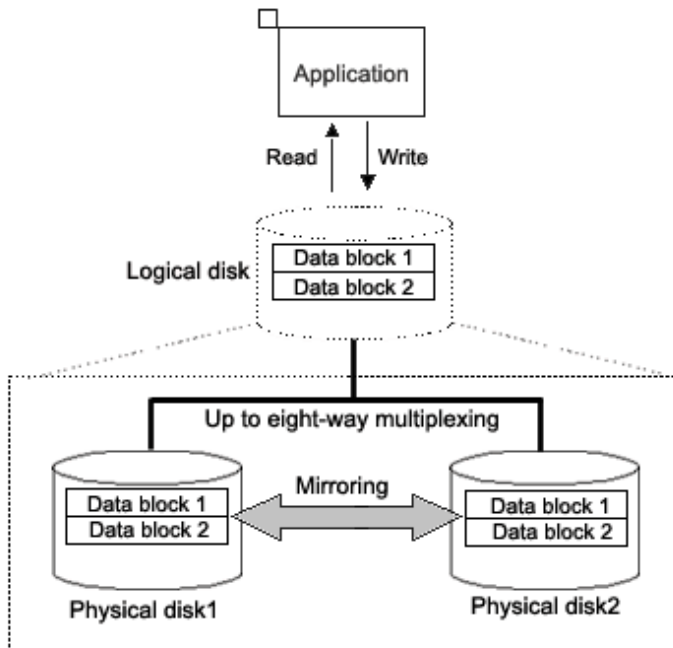
1.2.1 Disk Mirroring

Disk mirroring utilizes multiple disks in order to maintain data synchronization.

GDS realizes disk mirroring by creating one logical disk from multiple physical disks.

Disk mirroring allows applications to access a different disk even if one mirror disk (a disk that is mirrored) fails. Therefore, data will not be damaged and the applications will continue normal operation.

Figure 1.4 Disk Mirroring



GDS supports mirroring of a single disk unit (JBOD) and mirroring of disk arrays. It can mirror disk units of various usage and configuration, including a system disk installed with an operating system, and a disk connected to multiple servers.

Some examples of special mirroring implementation are explained below.

1.2.1.1 System Disk Mirroring

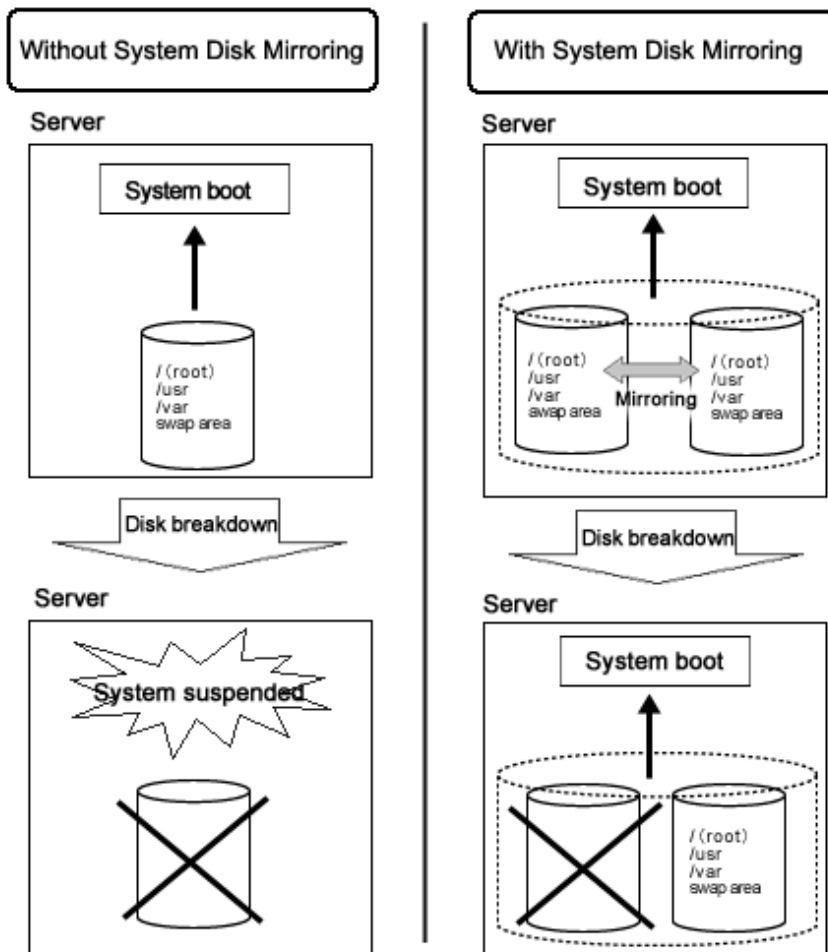
System disk mirroring is the function that mirrors the system disk on which the running Solaris has been installed.

If a system disk failure occurs, the entire system probably stops and even booting will become impossible, and operating system reinstallation will be required. As a result, service stop may remain for a long time.

System disk mirroring ensures continuous system operation even when a failure occurs in part of the disks. Additionally, it enables system boot in a situation where disk failure has occurred.

System disk mirroring in a SAN boot environment is also supported.

Figure 1.5 System Disk Mirroring




1.2.1.2 Mirroring between Disk Array Unit

GDS realizes high-performance, high-reliability mirroring of disk arrays.

Mirroring two disk arrays connected with Fibre Channel provides data protection from unexpected accidents and blackouts.

Also, by combining this feature with special software, disk units with redundant access paths can also be mirrored.

Figure 1.6 Mirroring between Disk Array Unit




1.2.1.3 Shared Disk Mirroring

GDS provides a function that mirrors shared disk units connected a cluster system composed of multiple servers (also referred to as nodes). This is called shared disk mirroring in distinction from local disk mirroring that mirrors disks connected to a single node.

The GDS's shared disk mirroring function can be used with applications, such as the GFS Shared File System, that provide simultaneous access from multiple servers to shared disks, not to mention switch-over or standby type applications.

Figure 1.7 Shared Disk Mirroring

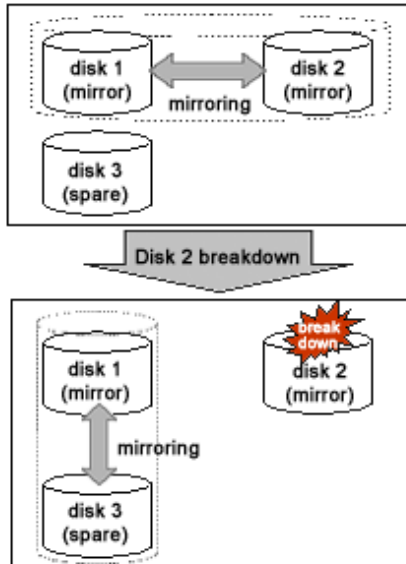


1.2.2 Hot Spare

Overview

"hot spare" is a feature that realizes automatic mirroring recovery using spare disks in the event of mirrored disk failure.

Figure 1.8 Hot Spare



Spare Disk Automatic Connection

If an I/O error occurs in a disk connected to a mirror group, a spare disk is automatically connected to the mirror group. Subsequently, synchronization copying for the spare disk takes place and the mirroring status will be restored.

Spare Disk Automatic Disconnection

After the disk where an I/O error occurred is recovered, the spare disk is automatically disconnected from the mirror group. For example, if disk failure causes an I/O error and a spare disk is automatically connected, after the failed disk is swapped with another disk and synchronization copying to the new disk is complete, the spare disk will automatically be disconnected.

Hot Spare Mode (Spare Disk Selection Mode)

A spare disk automatically connected in the event of I/O error in a mirrored disk is selected from spare disks that are registered with the failed disk's class. There are two modes of selecting spare disks: external mode and internal mode.

- External Mode (Default)

If an I/O error occurs in a disk of a disk array unit, this method selects a spare disk that belongs to a different disk case from that of the failed disk.

If an I/O error occurs in a disk irrelevant to a disk array unit (such as an internal disk), it selects a spare disk that is connected to a controller different from that the failed disk is connected to first.


When no applicable unconnected spare disk is found there, a spare disk that belongs to the same disk case or is connected to the same controller as that of the disk with the I/O error is selected.

Features:

If an I/O error occurs in a disk, the I/O cable of the disk case may cause an error, the entire disk case may be down, or the controller

may have a breakdown. By searching for a spare disk starting from a different disk case or controller from that of a disk with an I/O error, a normal spare disk can promptly be found and early mirroring recovery becomes possible.

Figure 1.9 Hot Spare in External mode



- Internal Mode

If an I/O error occurs in a disk of a disk array unit, this method selects a spare disk that belongs to the same disk case as that of the failed disk.

If an I/O error occurs in a disk irrelevant to a disk array unit (such as an internal disk), it selects a spare disk that is connected to the same controller as that of the failed disk.


When no applicable unconnected spare disk is found there, spare disk automatic connection is not performed.

Features:

A configuration that mirrors disks belonging to different disk unit cases or connected to different controllers is the configuration with high-availability, which secures continuous operations even if one of the disk cases or controllers fails. By selecting a spare disk that

belongs to the same disk case or is connected to the same controller as that of the disk with an I/O error, the configuration with high-availability can be maintained.

Figure 1.10 Hot Spare in Internal Mode



Note

Points of Concern for Hot Spare

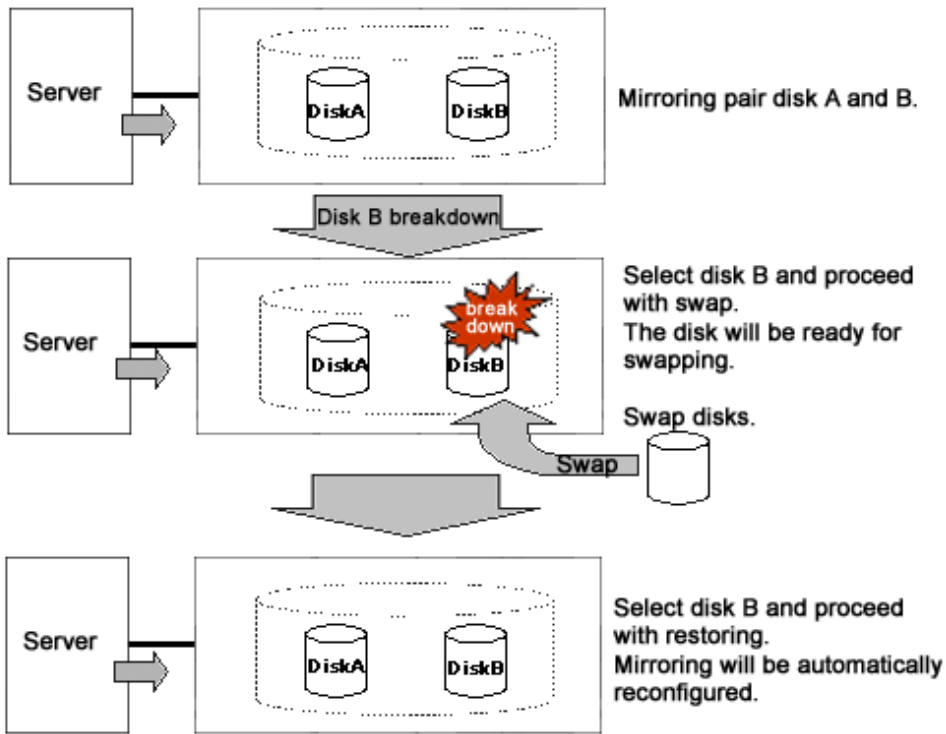
See "A.2.8 Hot Spare."

1.2.3 Hot Swap

Hot swap allows exchange of faulty disk unit parts without stopping the application in the event of a mirror disk failure. With GDS, administrator does not have to be aware of the system configuration. The system configuration is transparent to the administrator. Therefore, the administrator can simply select a failed disk displayed in the GUI window to conduct preparation before swapping disks and restoration through mirroring after the swapping.

If a spare disk has been operating in place of the defective disk unit, the spare disk will disconnect automatically, returning to its original state.

Figure 1.11 Hot Swap



Note

Conditions for Hot Swap

The GDS hot swap function is available only for exchanging disks of disk units supporting hot swap.

Do not remove or insert disks of disk units not supporting hot swap. It may cause breakdown or damage data.

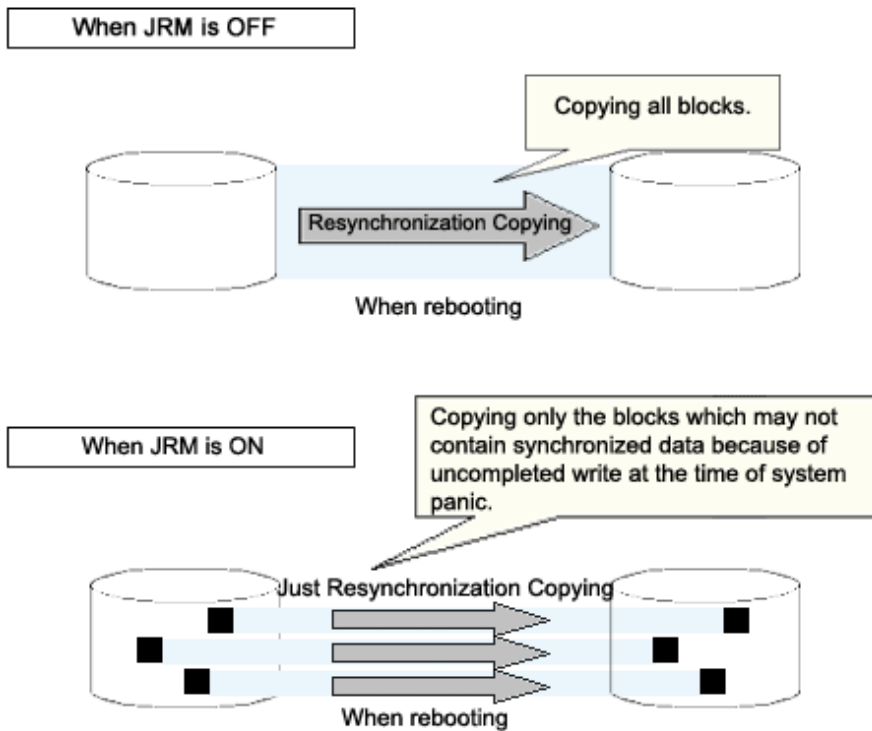
1.2.4 Just Resynchronization Mechanism (JRM)

When a system is down due to panic and so on, data must be resynchronized between the disk units after reboot (after cluster application failover for a cluster system).

Although the user can run an application during this copying process, redundancy will be lost and loads will be imposed during the process.

In order to resolve this problem GDS offers a mechanism called Just Resynchronization Mechanism (JRM). JRM completes copying, which usually takes a few minutes for every 1GB, in no time by copying only portions where mirroring synchronization is no longer retained.

Figure 1.12 Just Resynchronization Mechanism (JRM)



Information

What is JRM?

JRM stands for Just Resynchronization Mechanism, a feature that only allows duplication of the damaged part of data.

Information

Three Types of JRM

There are three types of Just Resynchronization Mechanism (JRM): for volumes, for slices and for proxy. For details, see "[A.2.14 Just Resynchronization Mechanism \(JRM\)](#)."

1.3 Functions for High Manageability

GDS supports various tasks related to disk management, such as installation, operation, and maintenance.

Efficient management functions are essential especially in a SAN environment, where one disk is directly shared with multiple servers.

Functions improving disk manageability are explained below.

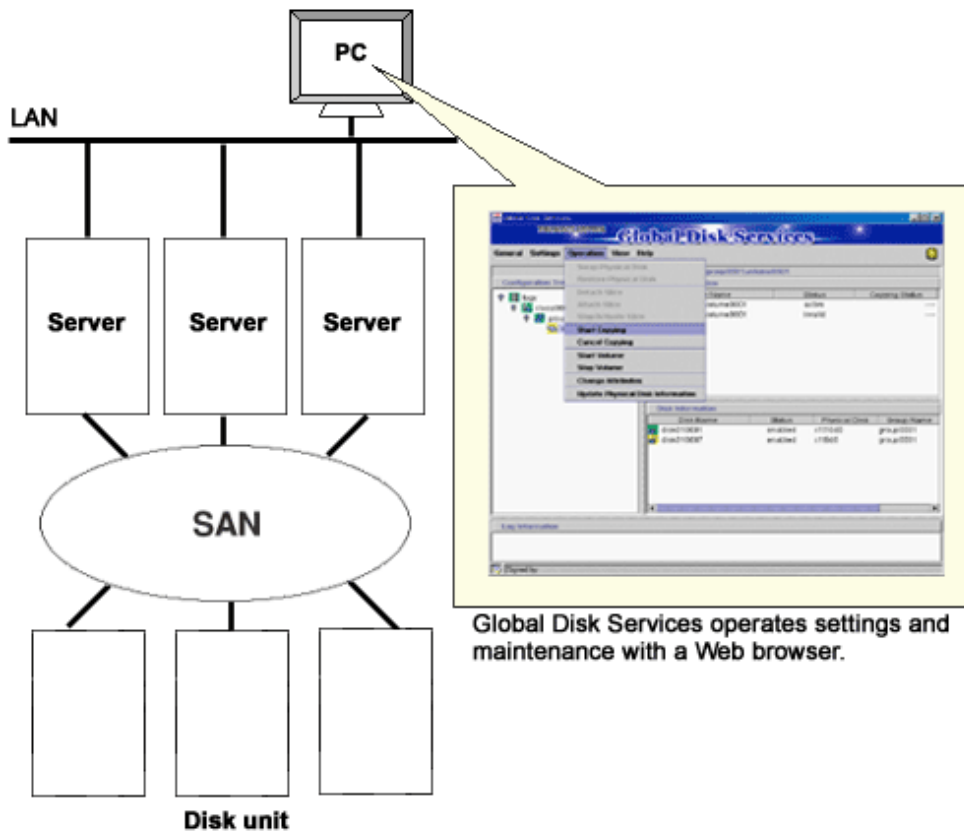
1.3.1 Operation Management Interface

By registering a disk unit with GDS, system administrators will be able to centrally manage all operation on the disk (configuration setting, configuration management, operation monitoring, backup, data migration, configuration change, maintenance) via GDS operation management interface.

Operation management interface provides the GUI (Graphical User Interface), automatic processing, operation log, and CLI (Command Line Interface) useful for liaising with other applications, allowing easy and intuitive operation even for inexperienced Solaris users.

Since the GUI is Web browser-based, system administrator will be able to centrally monitor and operate from any remote location.

Figure 1.13 Operation Management Interface



Global Disk Services operates settings and maintenance with a Web browser.

See

For the operation methods of the GUI, see "Chapter 3 Starting and Exiting GDS Management View," "Chapter 4 Management View Screen Elements," and "Chapter 5 Operation."

For the usage methods of CLI, see "Appendix D Command Reference."

For the operations available on GUI and on CLI, see "5.1.10 Operations from GDS Management View."

1.3.2 Centralized Disk Management


By registering all disk units connected to servers with GDS, it becomes possible to perform the centralized management of all operations on the disks from the GDS operation management interface.

Since disk units managed by GDS are virtualized as logical volumes, there is no need for the application to handle physical disks.

As to disk units, there are no limitations on physical configurations (single disks, disk arrays, multipath), connection configurations (local connection, shared connection), and intended use (system disks, local disks, cluster switch-over disks, cluster shared disks). Centralized management can be performed on any disk units.

Whether or not to mirror the managed disk can be determined as necessary. For example, a user may want to manage a disk array with sufficient availability, without carrying out mirroring.

Figure 1.14 Centralized Disk Management



1.3.3 Name Management

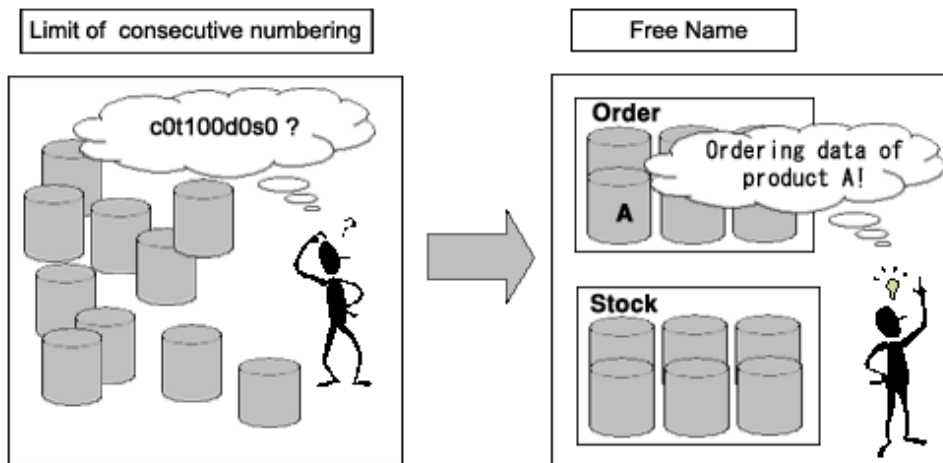
In Solaris, disks are named in "cXtXdX" format, where "X" is a number assigned by the OS, and the administrator differentiates the disks by the assigned consecutive numbers.

This was not a problem if the disk configuration was small, and if the disks were accessed from a certain server. However, when the environment is such that there are many disks connected, or in a SAN environment where a disk is shared by multiple servers, it is impossible to manage the disks using consecutive numbers.

With GDS, the administrator can freely name objects such as physical disks and logical volumes. Names that are easy to remember, for example, those associated with hardware configuration or data contents can be assigned.

Once an object is named, the name remains the same even if the physical configuration is changed. In other words, the user does not have to make any changes to applications that recognize the name.

Figure 1.15 Free Device Name



1.3.4 Single System Image Environment

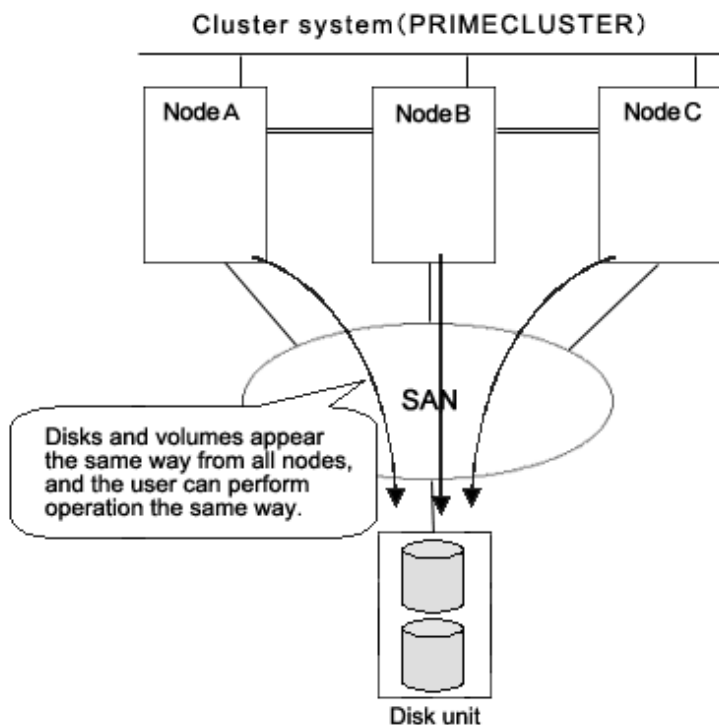
In a SAN environment, a disk unit can be accessed from multiple servers.

GDS provides a single system image environment, in which a cluster system with multiple servers (also referred to as nodes) appears as one system to the users or application.

This single system image environment can be utilized as explained below.

- Application programs sharing the same disk can lookup disks and volumes using the same name from all nodes.
- Application programs can access the disk simultaneously from all nodes. For details, see "Note."
- Application programs can perform operations from all nodes, such as changing the configuration on objects (such as disks and volumes). Changes are reflected on all nodes.

Figure 1.16 Single System Image Environment





Access Exclusion Control

When the same block is accessed simultaneously, data consistency is maintained by access exclusion control performed by the application that runs on multiple nodes simultaneously.

1.3.5 Access Control

GDS provides the following access control functions to prevent data damage from events such as an improper operation by the user.

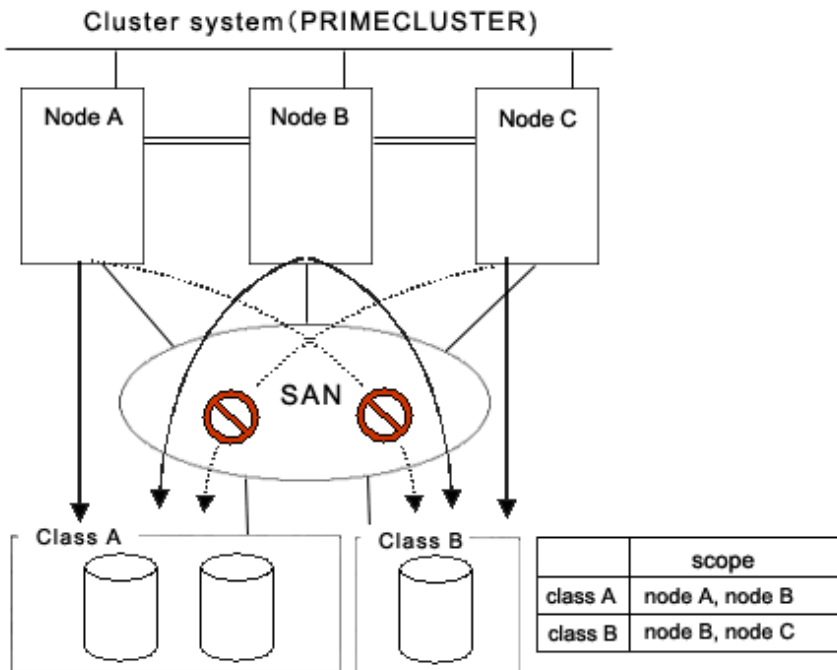
Class Scope

In order to manage a disk unit with GDS, it must be registered with a certain class. A class is like a container that holds multiple disks. By defining the class attribute called "scope", the user can specify which node set has access, or can perform operation on the disks registered to that class. Since operation on the disk is restricted to the nodes you specify, there is no risk of changing the configuration by mistake, or losing data consistency.

Suppose there are disks connected to two nodes; node A and B. The administrator wants node A to have access to some of the disks, while node B has access to the remaining disks.

In this case, two classes should be created to separately manage the disks shared by node A and B. This ensures that access by the node that is not selected in the class scope is restricted.

Figure 1.17 Class Scope




Starting and Stopping Volume

GDS's logical volume can be started or stopped for each node sharing the volume.

Since a stopped volume cannot be accessed from the node, there is no risk of losing data consistency by mistake.

Figure 1.18 Starting and Stopping Volume




Access Mode

Logical volume has an attribute called "Access mode", which can be defined for each node sharing the volume. There are two access modes; "Read and write possible" mode, and "Read only possible" mode.

For example, if a certain node will access the logical volume to create data backup, set the mode "Read only possible". That way, you can prevent writing to the volume by mistake.

Figure 1.19 Access Mode



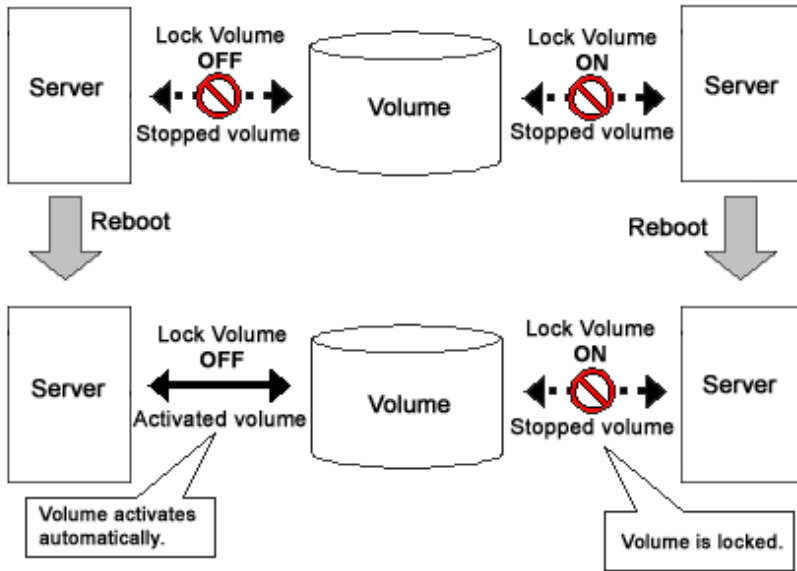
Lock Volume

When the node or cluster application is activated, logical volume starts automatically and becomes accessible. Likewise, when the cluster application is terminated, logical volume will also stop. This prevents a node with terminated application from accessing the logical volume.

However, rebooting a node could result in starting the volume.

In order to preclude the logical volume from starting in such a situation, the user can define the "Lock volume" attribute. When "Lock volume" is selected, volume will not be activated even when the node is rebooted, or cluster application activated.

Figure 1.20 Lock Volume



1.3.6 Realizing Large Capacity and I/O Load Balancing

In our current SAN environment, the demand for large capacity disk units and the amount I/O data processing is increasing daily.

This sub-section describes the functions that realize flexible disk configuration and I/O load balancing for efficient management of large volumetric data.

 **Note**

Logical partitioning, disk concatenation, and disk striping can be applied to local disks and shared disks. They cannot be applied to system disks.

1.3.6.1 Logical Partitioning

Logical partitioning divides a physical disk into logical devices based on its original method, and not based on the disk slice management provided by the VTOC.

In Solaris, physical disks can only be partitioned to a maximum of 7 disk slices (physical slices).

GDS allows users to use physical disks and objects equivalent to disks dividing them into a maximum of 1024 logical devices.

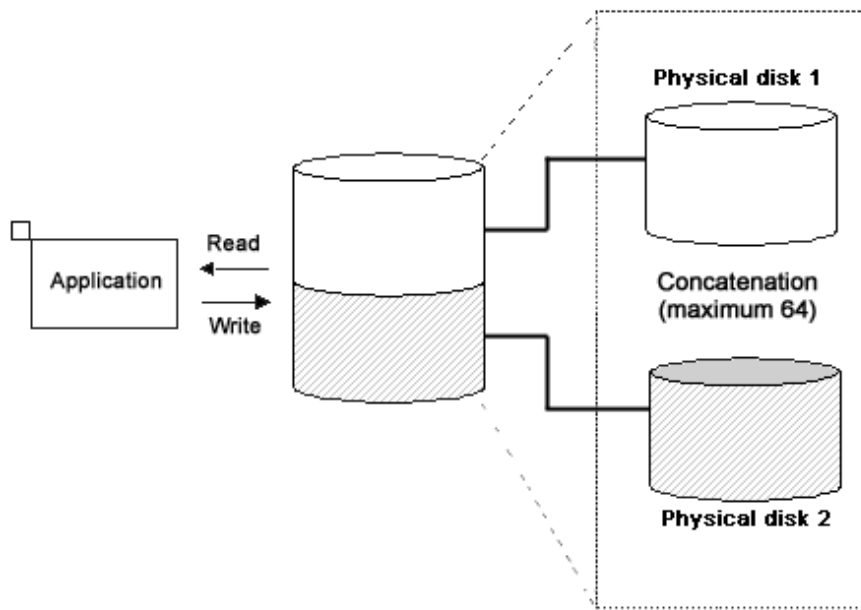
For details, see "[2.5 Logical Volume](#)" and "[D.4 sdxvolume - Volume operations.](#)"

1.3.6.2 Disk Concatenation

Disk concatenation combines multiple physical disks to form a single, large logical disk.

By using the disk concatenation function, users are able to configure a large-capacity logical disk device without being restricted to the limitation of one physical disk.

Figure 1.21 Disk Concatenation



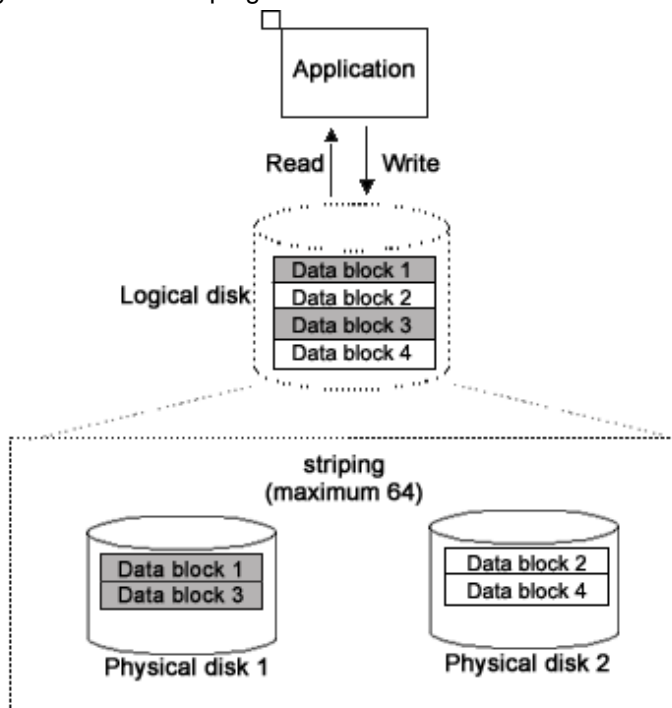
1.3.6.3 Disk Striping

Disk striping maps equal-sized data units on multiple disks so that the data is interleaved in sequence among two or more physical disks.

Disk striping function can balance the I/O load by distributing the divided I/O requests to multiple physical disks simultaneously.

You can also stripe concatenated disks.

Figure 1.22 Disk Striping




1.3.6.4 Combining Disk Striping with Mirroring

Concatenation and striping does not provide data redundancy. Since it involves more disks compared to a usual disk configuration, the risk of data loss caused by a disk failure is actually larger.

GDS delivers data redundancy as well as high-capacity storage and I/O load distribution by mirroring concatenated or striped disks. When using concatenation or striping, you are recommended to use mirroring as well.

Figure 1.23 Combining Striping and Mirroring




1.3.7 Online Volume Expansion

Overview

Volume expansion is a function that expands the capacity of volumes, retaining volume data. Volume expansion can be conducted without stopping applications using the volumes. This function is referred to as online volume expansion.

Volumes are expanded by disk area addition to areas after the last blocks of the volumes. Therefore, to expand a volume, there must be sufficient continuous free space following the last block of the volume.

Figure 1.24 Online Volume Expansion



Concatenation and Online Volume Expansion

Even if there is no sufficient continuous free space after the last block of the volume, by concatenating unused disk, online volume expansion will be available.




Use Conditions on "Concatenation and Online Volume Expansion"

This function can expand only volumes that meet the following conditions.

- Concatenation and mirroring are both applied.
In other words, concatenated disks have been mirrored.
- The number of concatenated disks can be one.
- The multiplicity of mirroring can be one.

For details, see "Concatenation and Online Volume Expansion" in "[A.2.15 Online Volume Expansion](#)."

Figure 1.25 Concatenation and Online Volume Expansion



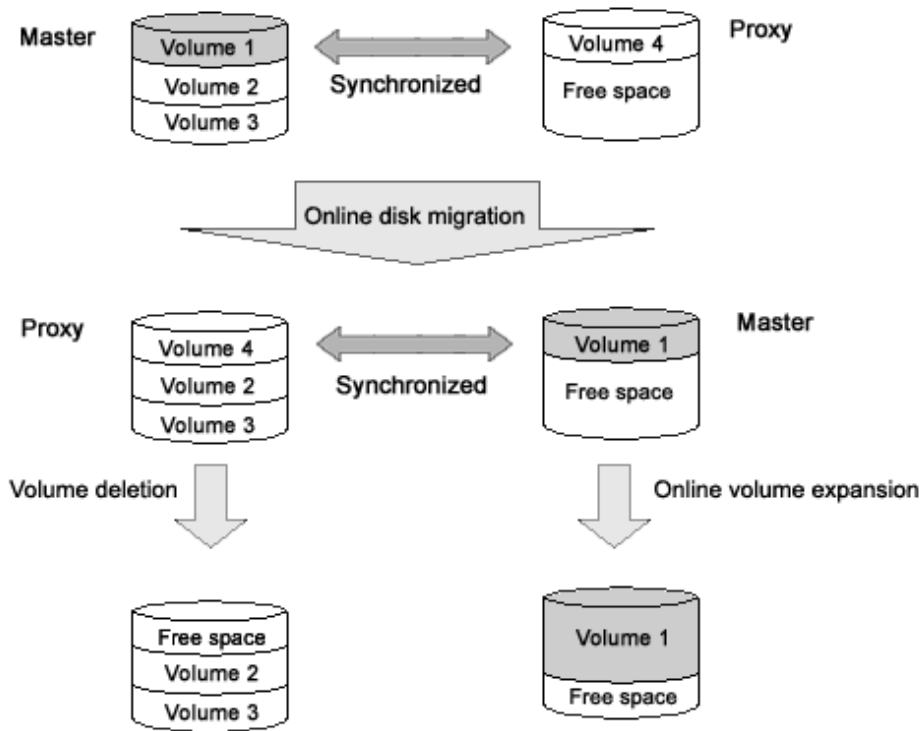
Online Disk Migration and Online Volume Expansion

Even if there is no sufficient continuous free space after the last block of the volume, by migration of the volume to a disk with sufficient free space, online volume expansion will be available. For the volume migration, use the GDS Snapshot online disk migration function.



For online disk migration, see "[1.5.5 Online Disk Migration](#)."

Figure 1.26 Online Disk Migration and Online Volume Expansion



Online System Volume Expansion

Online volume expansion is also available for system volumes. Even if there is not sufficient continuous free space after the last block of the volume to be expanded, by alternative boot environment creation with GDS Snapshot, online volume expansion will be available.



See

For the procedures for creating alternative boot environments and expanding system volumes without sufficient free space after the last blocks, see "D.20 Volume Expansion Using Commands."

Operating Instructions

Use the `sdxvolume -S` command. For details, see "D.4 sdxvolume - Volume operations."



Note

Points of Concern for Online Volume Expansion

See "A.2.15 Online Volume Expansion."

1.3.8 Snapshots by Slice Detachment


By temporarily detaching slices from mirror volumes, you can create snapshots (replications) of the volumes. In other words, by use of detached slices, volume data at the moment can be accessed.

Detached slices are logical devices separated from volumes, and other applications that use the detached slices can run simultaneously with service applications that use the volumes. For example, backup applications can be executed concurrently with service applications.

For volumes shared on multiple nodes, by separating nodes running service applications that use the volumes from ones running applications that use detached slices, a management mode in which process loads are not imposed mutually can also be established.

After applications that use detached slices end, to create snapshots again, it is necessary to reattach the detached slices to volumes. At this point, copying will be performed to synchronize data on the slices and the volumes. For this copying, GDS uses JRM (Just Resynchronization Mechanism) for slices to realize high-speed resynchronization by copying only the portions that were updated while slices were detached.

Figure 1.27 Snapshot by Slice detachment



 Note

Ensuring Consistency of Snapshot Data

If you create a snapshot while an application is accessing the volume, the data consistency may not be ensured, as the volume data will be incomplete.

To ensure the consistency of your snapshot data, you must stop the application that is accessing the volume in advance. After creating the snapshot, start the application again.

For example, when using the volume as a file system such as GFS or ufs, unmount the volume with the `umount(1M)` command before creating a snapshot, and mount it with the `mount(1M)` command afterwards. This way, you can ensure the consistency of the snapshot data.

To create a snapshot while running the application, the file system or database system you are using to manage the data must be able to ensure data integrity.

 **Note**

JRM for Slices

JRM for slices speeds up the resynchronization process when attaching a detached slice again to the volume. GDS records the changes made on the volume and slice in the memory while the slice is being detached. The resynchronization copy performed when the detached slice is reattached copies the updated portions only to realize high-speed resynchronization.

JRM for slices becomes effective when a slice is detached while the `jrm` attribute of the slices is on. However, if a system is stopped or if the slice is taken over by the `sdxslice -T` command while the slice is detached, just resynchronization is not conducted when the temporarily detached slice is attached again. Resynchronization is performed by copying the entire data, not only the updated portions.

Therefore, if you plan to shut down the system, or have a slice taken over, attaching the slice to the volume in advance is highly recommended.

 **Information**

What is JRM?

JRM stands for Just Resynchronization Mechanism, a feature that only allows duplication of the damaged part of data.

 **Information**

Three Types of JRM

There are three types of Just Resynchronization Mechanism (JRM): for volumes, for slices and for proxy. For details, see "[A.2.14 Just Resynchronization Mechanism \(JRM\)](#)."

1.4 GDS Snapshot Features

GDS Snapshot is software that provides additional functions to GDS.

GDS secures data from disk failures using the mirroring function and realizes continuous service. However, the mirroring function is not capable of protecting data against accidental erasure by the user or data crash due to an application malfunction. Therefore, data back up is mandatory. In a traditional system, the service must be stopped to back up data. GDS Snapshot provides a function for realizing the backup operation that does not influence the service and allows for continuous service.

The causes of cessation of services are not only problems such as disk failure and data damage. In a traditional system, the service must be stopped even for intentional maintenance work such as volume expansion and work applying patches to software. The function provided by GDS Snapshot also reduces system stop time and service stop time when performing maintenance work.

Once GDS Snapshot is implemented, proxy volumes and shadow volumes become available. These volumes are virtual devices that are bound up with GDS logical volumes. The user can make use of the GDS Snapshot function by operating and accessing those proxy volumes and shadow volumes.

The following sections discuss the GDS Snapshot functions provided by proxy volumes and shadow volumes.


1.5 Proxy Volume

Snapshots (replications at a certain moment) of volumes for service applications can be created on other volumes. The former volumes are referred to as the master volumes and the latter as the proxy (or alternate) volumes.

Using this proxy volume will resolve various issues. For example, stopping a system with a large amount of data for backup purposes is becoming more and more difficult.

Conventionally, data was backed up overnight when no one was on the system. However, as the amount of data grows, backup often cannot be completed by the following morning. Also, the widespread usage of Internet and service diversification has made it impossible to stop many systems.

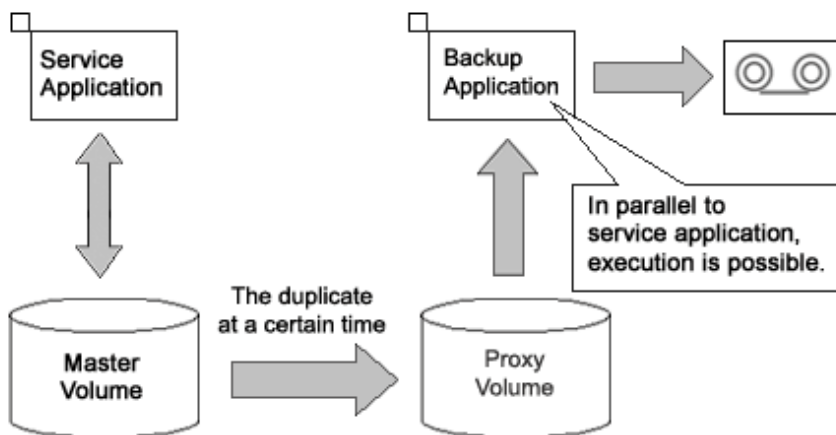
Figure 1.28 Problems Inherent to a 24-hour Operation System



Since the proxy volume can be accessed separately while the service application is accessing the master volume, users can run a backup application in parallel without worrying about the time.

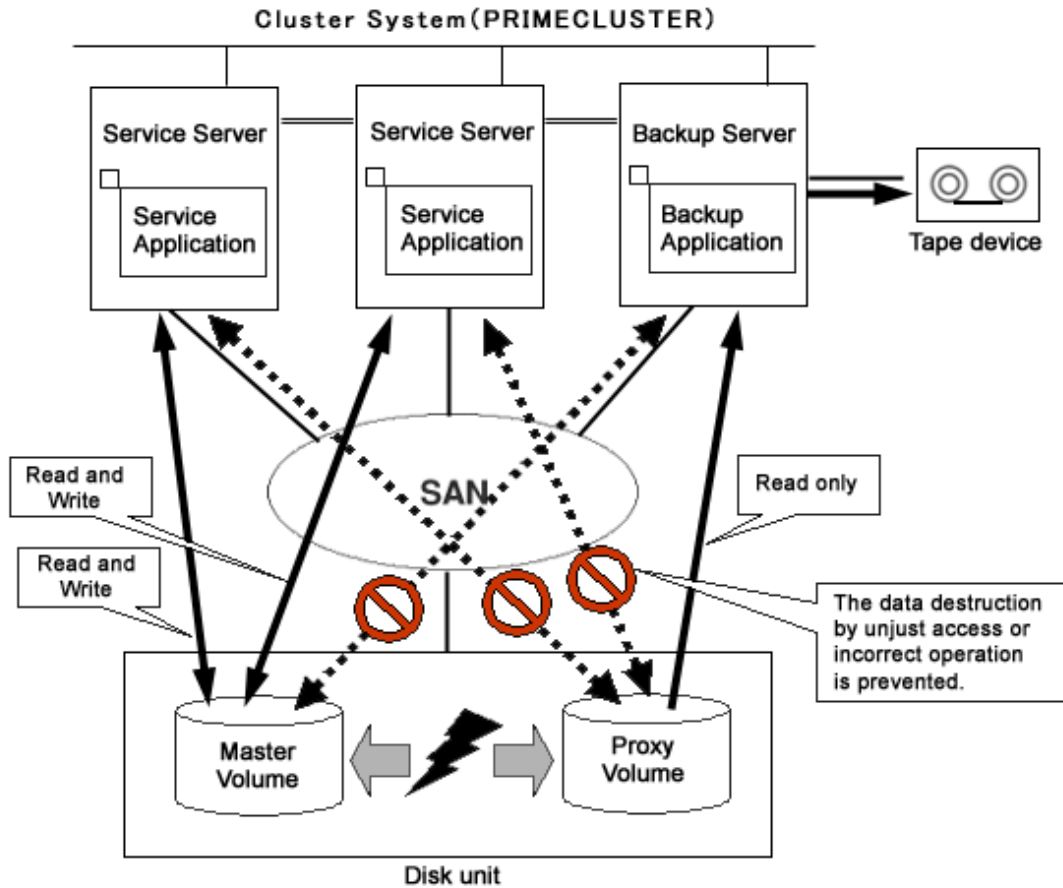
By using the proxy volume, users can conduct various tasks such as non-disruptive backup, data analysis, verification, and data protection from disaster, without affecting the main service application.

Figure 1.29 Backup Using the Proxy Volume



The user can utilize given logical volumes managed by GDS as master volumes or proxy volumes, including the root volume (system volume), local volumes in single server configuration, and shared volumes in cluster configuration. It also provides snapshot operation that is consistent with GDS's various volume management function, such as access control function.

Figure 1.30 Liaising with Volume Management Function



The snapshot functions using a proxy volume are explained here.

1.5.1 Snapshot by Synchronization

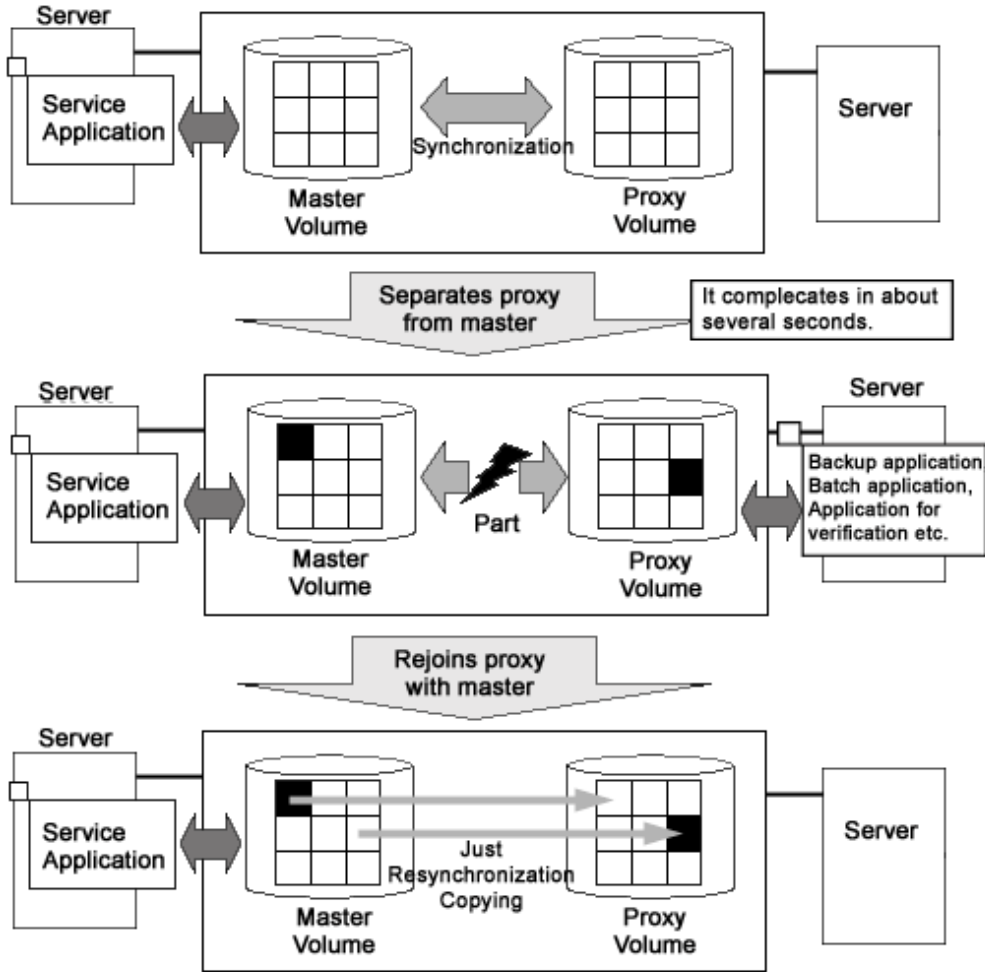
A master volume snapshot can be created by synchronizing data of the master volume and the proxy volume in advance and then separating the synchronized proxy volume from the master volume at any moment.

Snapshot by synchronization minimizes the influence on the service application and is suitable for routine scheduled backup.

As long as the two volumes are synchronized, snapshot can be created easily by separating the proxy volume from the master volume. Creation of snapshot will only take a few seconds even if the amount of data is large.

When the proxy volume is rejoined with the master volume for the next snapshot creation, high-speed resynchronization is performed by copying only updated portions on the master and proxy during the separation using JRM (Just Resynchronization Mechanism) for proxies.

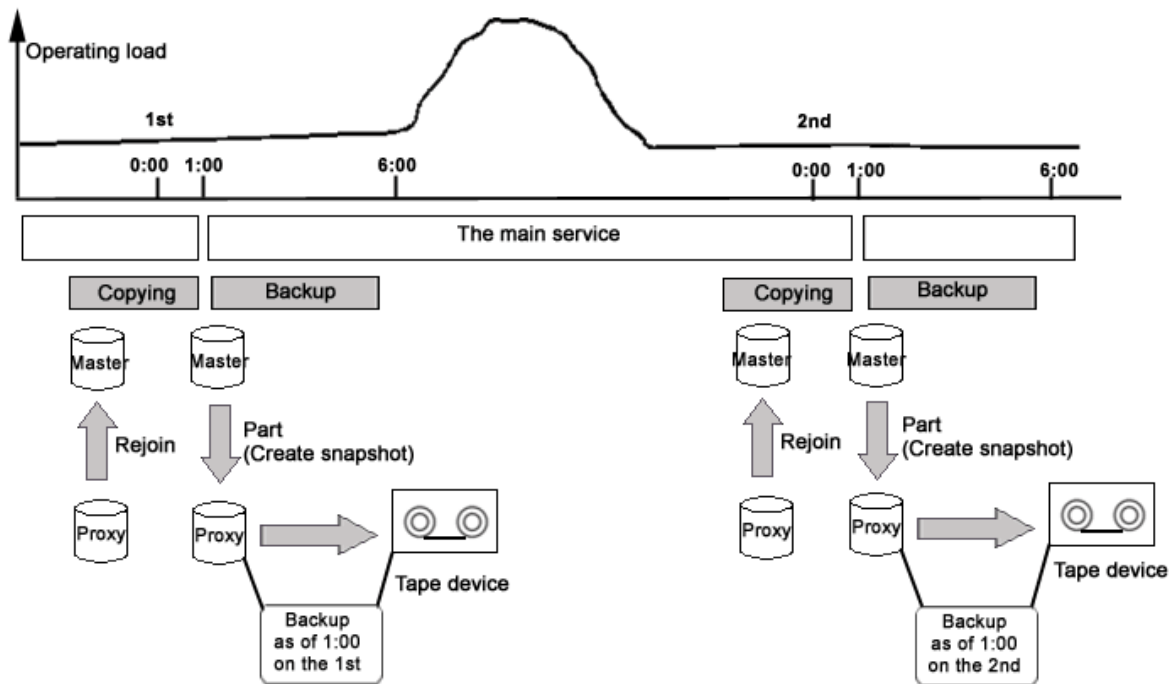
Figure 1.31 Snapshot by Synchronization



Therefore, both the snapshot creation and the resynchronization process before rejoining the proxy with master will have minimal influence on the service application.

For example, in a situation where loads imposed on disks are higher during the day than at night, if saving data to tape takes no more than 5 hours and resynchronization copying no more than an hour, the routine backup schedule is assumed as follows.

Figure 1.32 Example of a Daily Backup Schedule



 Note

JRM for Proxies

JRM for proxies speeds up the just resynchronization process when joining a parted proxy again to the master and when the master data is restored from the proxy. GDS records the changes made on the master and the proxy on the memory while the proxy is parted. The just resynchronization conducted when rejoining or restoring copies only updates the updated portions to realize high-speed synchronization.

JRM for proxies is enabled when the `pjrm` attribute of a proxy volume is set to on and the proxy volume is parted. However, if any node that is included in the scope of the class is stopped while the proxy is parted, just resynchronization is not put in operation. In other words, the entire data, not only the updated portions, is copied.

Therefore, if you plan to shut down the system, joining the proxy to the master in advance is highly recommended.

Reference to these matters is not necessary when you are using the copy function of a disk unit.

 Information

Three Types of JRM

There are three types of Just Resynchronization Mechanism (JRM): for volumes, for slices and for proxy. For details, see "[A.2.14 Just Resynchronization Mechanism \(JRM\)](#)."

 Note

Ensuring Consistency of snapshot data

If you create a snapshot while an application is accessing the volume, the data consistency may not be ensured, as the volume data will be incomplete.

To ensure the consistency of your snapshot data, you must stop the application that is accessing the volume in advance. After creating the snapshot, start the application again.

For example, when using the volume as a file system such as GFS or ufs, unmount the volume with the `umount(1M)` command before creating a snapshot, and mount it with the `mount(1M)` command afterwards. This way, you can ensure the consistency of the snapshot data.

To create a snapshot while running the application, the file system or database system you are using to manage the data must be able to ensure data consistency.

For details, see "[6.4 Online Backup and Instant Restore through Proxy Volume.](#)"



The Difference between a Mirrored Slice and Synchronized Proxy Volume

Although mirrored slices and synchronized master volume and proxy volume both have the same data, the ways they are used are different.

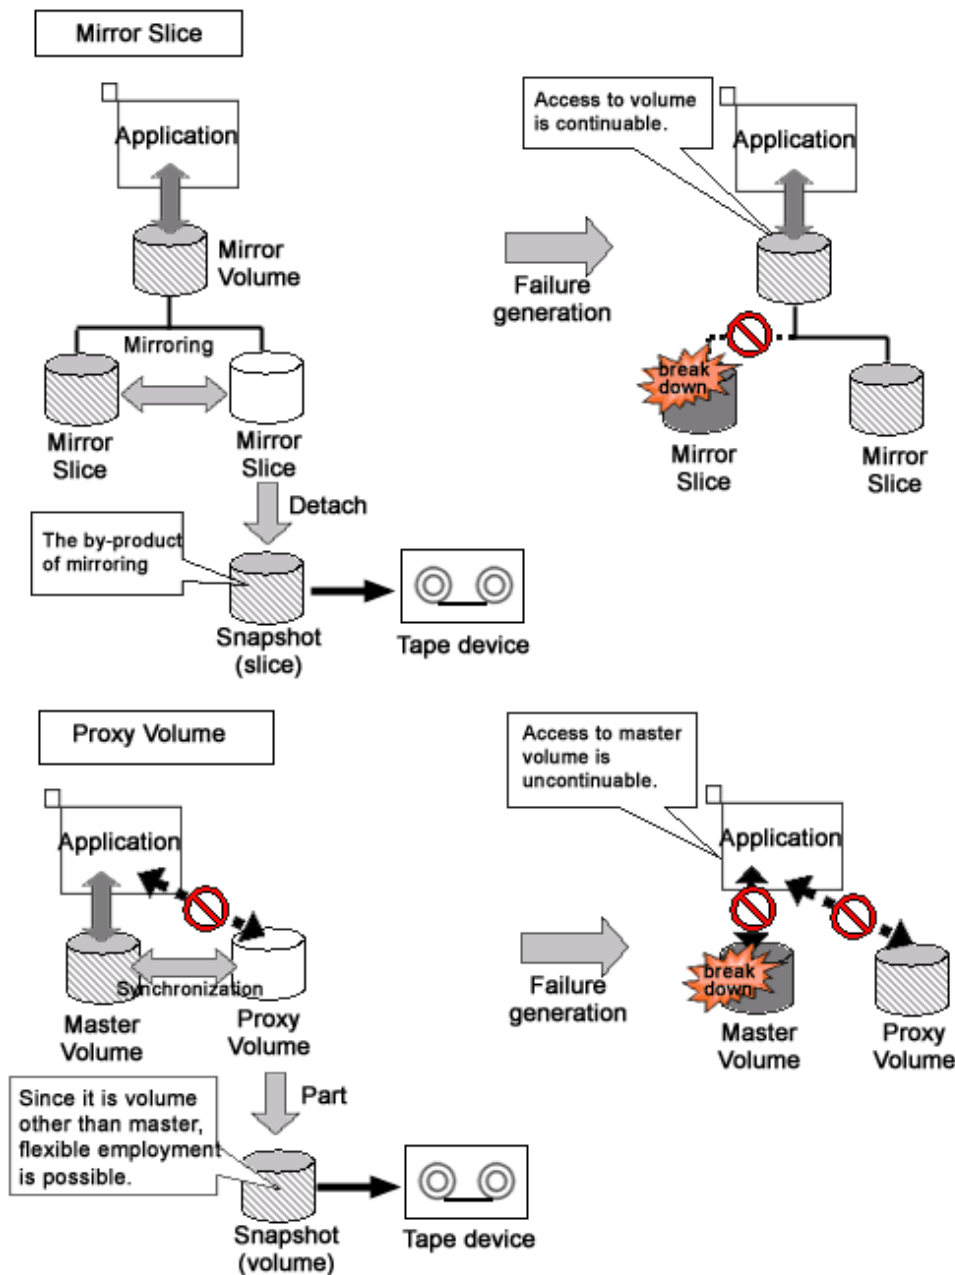
Mirrored slices are equals, and their purpose is to maintain data redundancy. So even if an abnormality is detected in one of the slices, the volume can be accessed as long as the other slice is normal.

However, even if the master volume and the proxy volume are synchronized, they are separate volumes and not equals. You may consider the master the primary volume, and the proxy the secondary volume. This means that even if the proxy volume is normal while all slices comprising the master volume are abnormal, you cannot continue accessing the master volume. The purpose of proxy volumes is to separate the master volume in order to perform a different task in parallel, and not to improve the data redundancy used by the service application.

In short, the snapshot by slice detachment is a by-product of mirroring, which is the primary objective, but the primary objective of the snapshot through use of proxy volumes is creating snapshots.

Therefore, GDS Snapshot provides more flexible disk configuration and supports various tasks.

Figure 1.33 Difference between a Mirrored Slice and Synchronized Proxy Volume

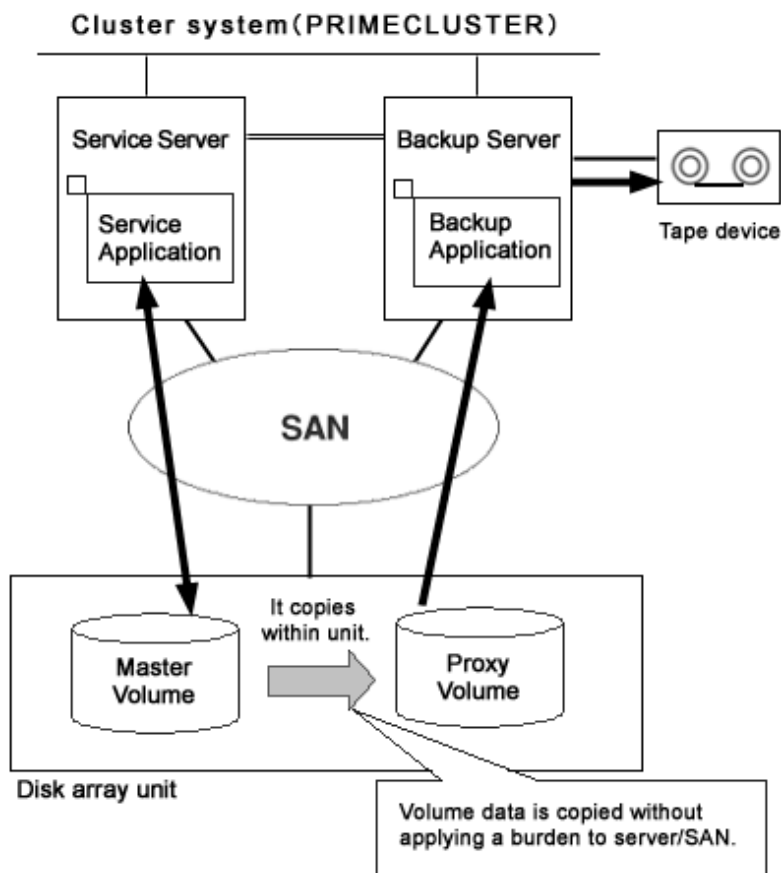


1.5.2 Snapshot Function with No Load to the Server/SAN

By cooperation with the Advanced Copy function of Fujitsu Storage Systems ETERNUS or TimeFinder and SRDF of EMC'S Symmetrix storage unit, processes such as copying for resynchronization, copying for maintaining synchronization, and recording the portion upgraded while proxies are being parted are performed within the disk units. Therefore, snapshots can be created without imposing loads on the server accessed by the service application, or the SAN (Storage Area Network).

As the operation is identical even when using GDS Snapshot with a disk unit lacking special copy function, there is no need for users to be aware of the differences between disk units.

Figure 1.34 Snapshot Function with No Load to the Server/SAN



1.5.3 Instant Snapshot by OPC

By cooperation with the OPC (One Point Copy) function provided by Fujitsu Storage Systems ETERNUS, snapshots can be created instantly at any given moment.

Unlike snapshot by synchronization, instant snapshot by OPC does not require synchronization with volumes in advance, so scheduling is unnecessary. However, as copying process will take place within the disk array after the snapshot has been created, I/O performance of the disk will be affected until the copying is complete.

Note

Instant Snapshot by OPC

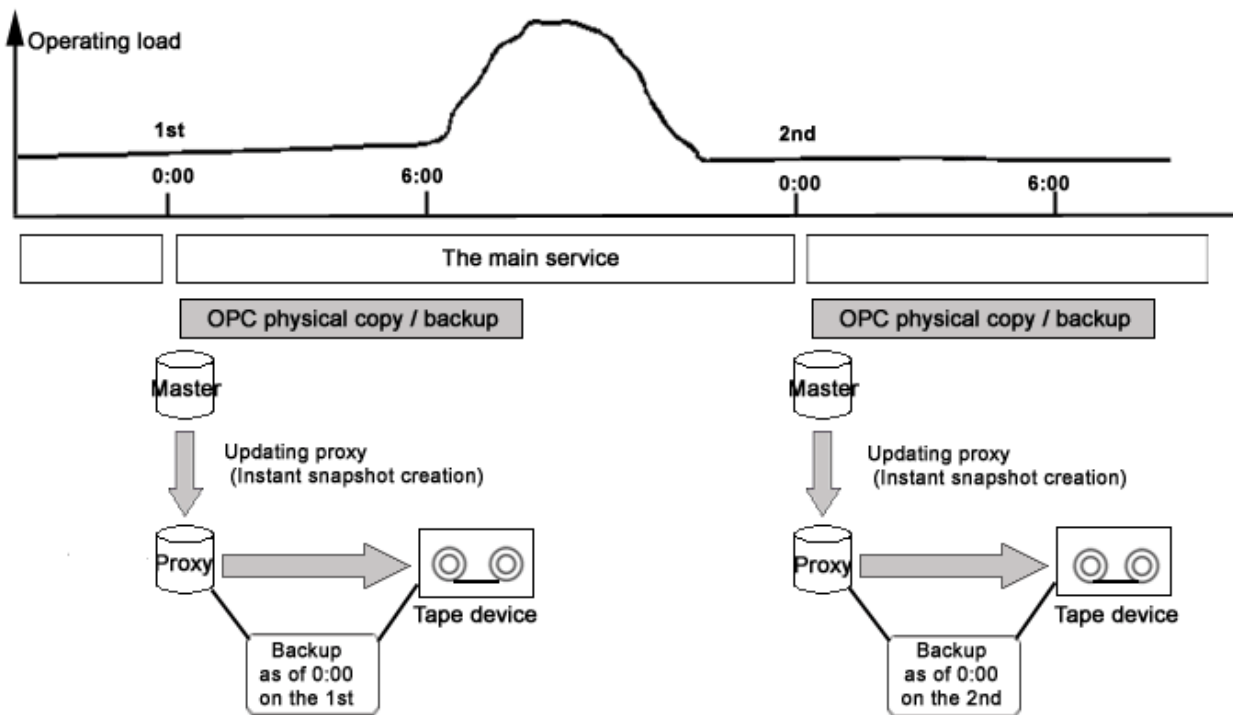
Instant snapshot by OPC is only available with the following storage systems supporting the function:

- ETERNUS DX60 / DX80 / DX90 (*1)
- ETERNUS DX80 S2 / DX90 S2 (*1)
- ETERNUS DX410 / DX440 (*1)
- ETERNUS DX410 S2 / DX440 S2 (*1)
- ETERNUS DX8100 / DX8400 / DX8700 (*1)
- ETERNUS8000
- ETERNUS6000
- ETERNUS4000 (Model 100 or higher) (*2)
- ETERNUS3000 (Model 100 or higher)

- ETERNUS2000 (Model 100 or higher) (*3)
- ETERNUS GR720
- ETERNUS GR730
- ETERNUS GR740

*1) GDS patch 914423-10 or later and GDS Snapshot patch 914457-03 or later are required.
 *2) GDS patch 914423-01 or later is required.
 *3) GDS patch 914423-05 or later is required.

Figure 1.35 Instant Snapshot by OPC



1.5.4 Instant Restore

For example, if a misoperation damages master volume data, you can still restore the master volume data using the proxy volume.

In this case, JRM (Just Resynchronization Mechanism) provides high-speed resynchronization by copying only portions that were updated while the master and proxy volumes were being separated. You can resume the service application immediately without waiting for the completion of the copying process. Synchronization copying process will take place in parallel with other tasks.


Users may use this function with any disk unit, as it is not dependent on a specific function of a disk unit.

Note

System Volume Restoration


To execute instant restore, the master volume and the proxy volume must be stopped temporarily. An active system volume running as a file system such as /, /usr, and /var cannot be the restore destination or source of instant restore because it cannot be stopped. For the method of restoring system volumes back from proxy volumes, see "[1.5.6 Creating an Alternative Boot Environment.](#)"

Figure 1.36 Instant Restore



For example, in a situation where loads imposed on disks are higher during the day than at night, if saving data to tape takes no more than 5 hours and resynchronization copying no more than an hour, the routine backup schedule that enables instant restore always from disks (proxy volumes) but not from tape is assumed as follows.

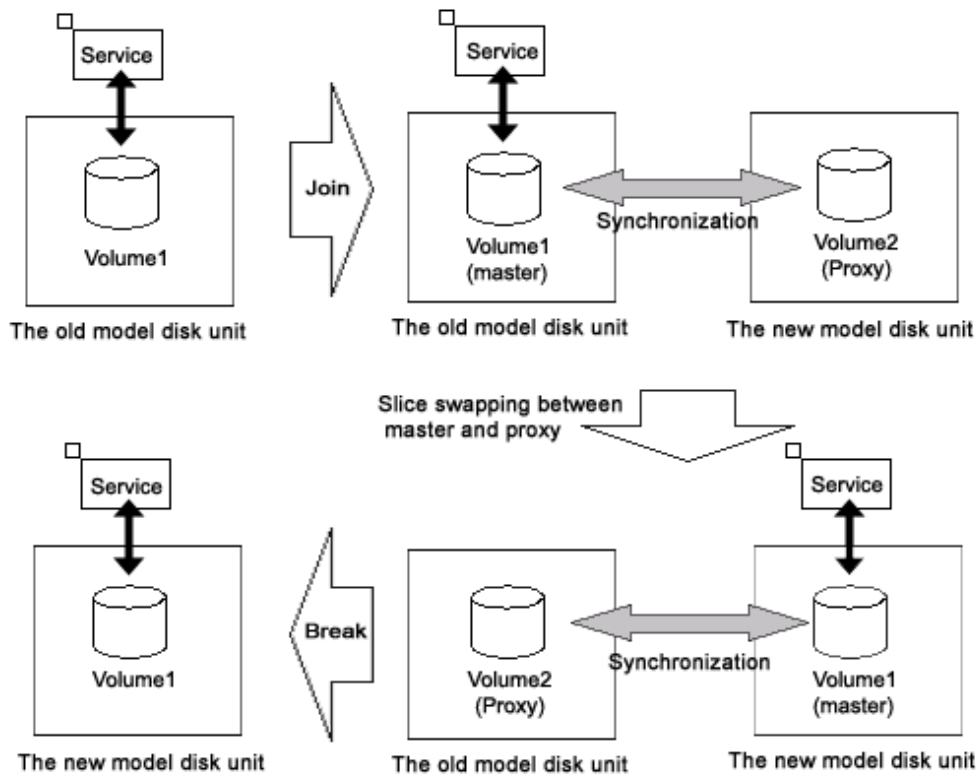
Figure 1.37 Example of a Backup Schedule Allowing Instant Restore at All Times



1.5.5 Online Disk Migration

When the master and proxy volumes are synchronized, migration from the physical disk used with the master volume to another is possible without affecting the service application. This is done by exchanging the slice comprising the master volume with the slice comprising the proxy volume and is useful in situations such as upgrading your disk units.

Figure 1.38 Online Disk Migration



1.5.6 Creating an Alternative Boot Environment

Using the snapshot function by proxy volumes, another boot environment (alternative boot environment) can be created without influencing the current boot environment during system operation. System stop time at the time of backup and restore work for system volumes (*/*, */usr*, */var*, swap area) can be reduced drastically by using the alternative boot environment. In addition, system and service stop time when modifying the system such as reconfiguring system volumes and applying patches can also be reduced.

In a SAN environment with a Fujitsu Storage Systems ETERNUS, it is possible to create an alternative boot environment using the Advanced Copy function without imposing loads on the server.

An alternative boot environment can be created with the GDS Snapshot command following the procedure as below.

- Creating snapshots of system volumes

Create snapshots of system volumes in the current boot environment using proxy volumes. Snapshots can be created during system operation.

- Configuring an alternative boot environment

Make booting available through snapshot volumes (proxy volumes) instead of system volumes in the current boot environment. This configuration changes the */etc/vfstab* files, the */etc/system* files and so on copied from the root volume (*/*) to the proxy volumes when snapshots were created. The current boot environment is not changed.

The environment can be switched to the alternative boot environment by simply specifying the name of a boot device name in the alternative boot environment and rebooting the system. The original boot environment can be restored by simply specifying the name of a boot device in the original boot environment and rebooting the system.

Using an alternative boot environment can reduce system and service stop time when performing maintenance work and reconfiguration work for the system as follows.

- Backing up the system

An alternative boot environment can be assumed as a backup of the system. Since an alternative boot environment can be created during system operation, stopping the system is no longer necessary at the time of backup work.

- Restoring when the system fails

When booting the system becomes unavailable due to a system disk failure or data crash, the system operation can be resumed by switching to an alternative boot environment. In addition, data on original system volumes can be restored from data in an alternative boot environment during system operation in the alternative boot environment, and the original boot environment can be switched back to by simply rebooting the system.

- Reconfiguring the system

The disk configuration of system volumes can be changed by creating an alternative boot environment and switching to the alternative boot environment. In addition, expanding the capacity and applying patches for volumes used in the alternative boot environment are available during system operation. These operations can be performed during system operation, and have no effect on the current boot environment. If booting an alternative boot environment fails or the performance of an alternative boot environment is improper, the current boot environment can be switched back to by simply rebooting the system.



See

For examples of using an alternative boot environment, see "[6.2 Backing Up and Restoring a System Disk through an Alternative Boot Environment](#)" and "[D.20 Volume Expansion Using Commands](#)."

1.6 Shadow Volume

Volumes in other domains (cluster systems or single servers) connected to a SAN can be recognized as volumes (shadow volumes) and accessed.

Proxy volumes and replications created with the copy function of a disk unit can also be recognized as shadow volumes and accessed in exchange for the volumes used in the service. Using shadow volumes on a server that does not belong to the domain running the service can realize the backup operation that does not impose loads on the domain running the service. Additionally, you can configure a backup server that consolidates backup of volumes from multiple domains in a SAN environment, and a disaster recovery system that conducts the service alternatively at a remote place in the unlikely event that disaster occurs.

This section describes the shadow volume feature.



Note

SAN Boot Environment

System volumes in a SAN boot environment are not recognized as shadow volumes.

1.6.1 Accessing Volumes from an External Server

In a SAN environment, disk areas used as volumes in other domains (cluster systems and single servers) can be recognized as volumes (shadow volumes) and accessed.

This feature allows a server that does not belong to the domain running the service to recognize and access volumes used in the service via the SAN. Not only the volumes used in the service but also the proxy volumes can also be recognized as shadow volumes by an external server. This capability enables another service, such as backup, restore, batch processing, and data analyses, to run concurrently with the service without imposing loads on the CPU, I/O, and the network of the domain running the service.


One server can recognize and access volumes in multiple domains connected to a SAN as shadow volumes. For example, you can configure a backup server that consolidates the backup and restoration of volumes in multiple domains.



See

For the operation patterns and the procedure for backup and restore by shadow volumes through an external server, see "[6.6 Backing Up and Restoring through an External Server](#)."

Figure 1.39 Backup Server



1.6.2 Using the Copy Function of a Disk Unit

Volumes can be copied with the copy function of a disk unit and the copy destination disk areas can be recognized as volumes (shadow volumes) and accessed.


To run another service concurrently with the service, the application volume must be replicated. The operation pattern that uses the copy destination disk areas created with the copy function of a disk unit is available as well as the operation pattern that uses proxy volumes, as replications of the volumes used in the service.



See

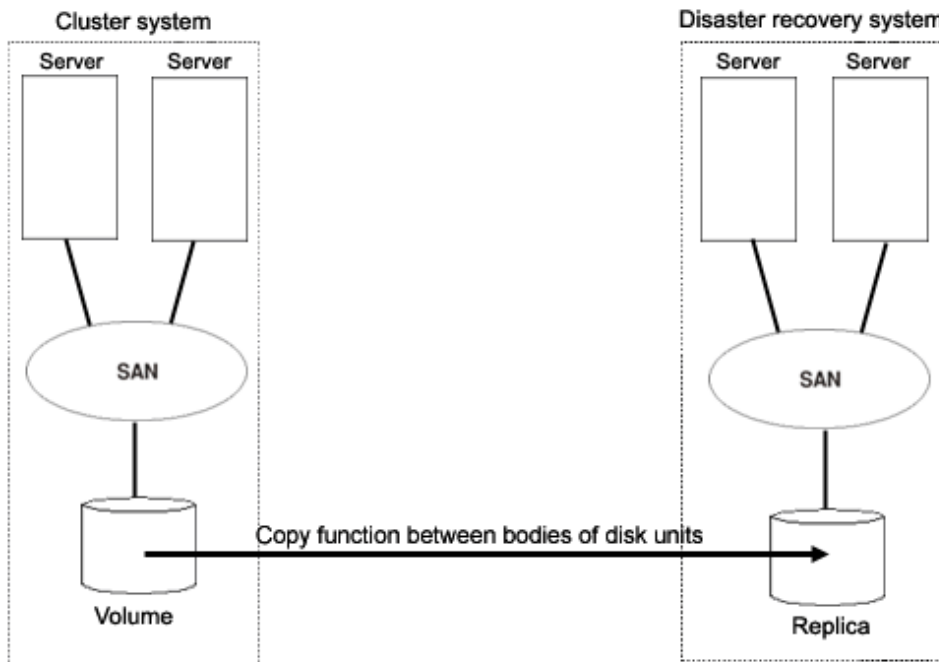
For the operations patterns and the procedure for backup and restore using the copy function of a disk unit and using shadow volumes, see "6.5 Backing Up and Restoring through Disk Unit's Copy Functions."

Figure 1.40 Backing Up a Replication Created with the Disk Unit Copy Function



You can configure a disaster recovery system at a place geographically distant using the copy function between disk bodies. If the domain running the service is affected by disaster, the disaster recovery system will conduct the service alternatively using the copy destination disk areas as shadow volumes until the domain recovers from disaster.

Figure 1.41 Disaster Recovery System



Chapter 2 Objects

GDS provides mirroring function and consistent manageability for disk units managed by GDS by virtualizing disk units as logical volumes. In this chapter, GDS object structure is explained so that you can understand the premise of logical volume management.

As GDS Snapshot, optional software of GDS, is implemented, proxy volumes and shadow volumes become available. By using proxy volumes and shadow volumes, the backup operation that can be conducted during service without affecting the running service can be realized.

This chapter systematically describes the virtual object structure that is the underlying basis of GDS and GDS Snapshot functionality to make it clear.

2.1 SDX Object

Virtual resources managed by GDS are called SDX objects, or objects.

There are five kinds of objects: classes, disks, groups, volumes and slices.

2.2 Disk Class

Disk class is the largest unit managed by GDS. Disk class may be referred to as "class".

In order to manage disk units (physical disks) with GDS, you must first register the physical disk with a certain class. A class is like a container that holds multiple disks for management.

By using the registered physical disks, GDS creates and manages objects such as disks, groups, volumes, and slices within a class.

Objects including classes can be named as the user wants. Respective names must be unique throughout the entire system.



Same Class Names

When multiple single nodes on which classes with the same name exist are changed over to a cluster system through installation of the cluster control facility, the duplicate class names come to exist in the cluster system. For details, see "[A.2.26 Changing Over from Single Nodes to a Cluster System.](#)"

Attributes

A class has the following attributes.

Name

This attribute identifies the class within a system.

Type

This attribute specifies the type of class. You can set to one of the following.

Root

Objects managed in this type of class are available only on the current node. Unlike with the local type class, system disks can be managed in this class.

Local

Objects managed in this type of class are available only on the current node.

Shared

Objects managed in this type of class are sharable with multiple nodes.

Scope

This attribute indicates a group of nodes on which objects in the class are available.

Hot Spare

This attribute indicates the hot spare operation mode. Either of the following values can be set.

on

Hot spare is enabled.

off

Hot spare is disabled. Automatic connections of spare disks are prevented.

Hot Spare Mode

This attribute indicates the spare disk selection mode of automatic connection for hot spare. One of the following values can be set.

External

If an I/O error occurs in a disk of a disk array unit, this method selects a spare disk that belongs to a different disk case from that of the failed disk.

If an I/O error occurs in a disk irrelevant to a disk array unit (such as an internal disk), it selects a spare disk that is connected to a different controller from that of the failed disk. When no applicable unconnected spare disk is found there, a spare disk that belongs to the same disk case or is connected to the same controller as that of the disk with the I/O error, is selected.

Internal

If an I/O error occurs in a disk of a disk array unit, this method selects a spare disk that belongs to the same disk case as that of the failed disk.

If an I/O error occurs in a disk irrelevant to a disk array unit (such as an internal disk), it selects a spare disk that is connected to the same controller as that of the failed disk. When no applicable unconnected spare disk is found there, spare disk automatic connection is not performed.

Operation

The following operations are available for classes.

Create

Registering a disk with the `sdxdisk -M` command and specifying a new class name will automatically create a new class.

For details on GDS Management View, see "[5.2.3.1 Class Configuration](#)."

Delete

Deleting the last registered disk from a class using the `sdxdisk -R` command will automatically delete the class. You can also delete the classes with the `sdxclass -R` command or the `sdxconfig Remove` command.

For details on GDS Management View, see "[5.5.4 Removing a Class](#)."

Status Display

Class status can be displayed with the `sdxinfo` command.

For details on GDS Management View, see "[5.3.1 Viewing Configurations/Statures and Monitoring Statures](#)."

Change Attributes

Class attribute values can be changed with the `sdxattr -C` command.

For details on GDS Management View, see "[5.4.1 Class Configuration](#)."

Recover

Closed classes can be recovered with the `sdxfix -C` command.
The operation from GDS Management View is unsupported.

Backup

The object configuration within a class can be backed up in configuration table format to a configuration file with the `sdxconfig Backup` command.
The operation from GDS Management View is unsupported.

Restore

The object configuration of a class can be restored according to the configuration table in the configuration file with the `sdxconfig Restore` command.
The operation from GDS Management View is unsupported.

Convert

A configuration table in a configuration file can be converted with the `sdxconfig Convert` command in order to restore the object configuration of the class in environment other than the original.
The operation from GDS Management View is unsupported.

Status

A class can have the following status.

Closed

All objects within a class are unavailable for reference.
For details, see "[F.1.4 Class Status Abnormality](#)."

Reference

There are the following points of concern for classes.

Rules

[A.1.1 Object Name](#)

[A.1.2 Number of Classes](#)

Important Points

[A.2.1 Physical Disk Driver](#)

[A.2.2 Managing System Disks](#)

[A.2.4 Booting from a CD-ROM Device \(or Network\)](#)

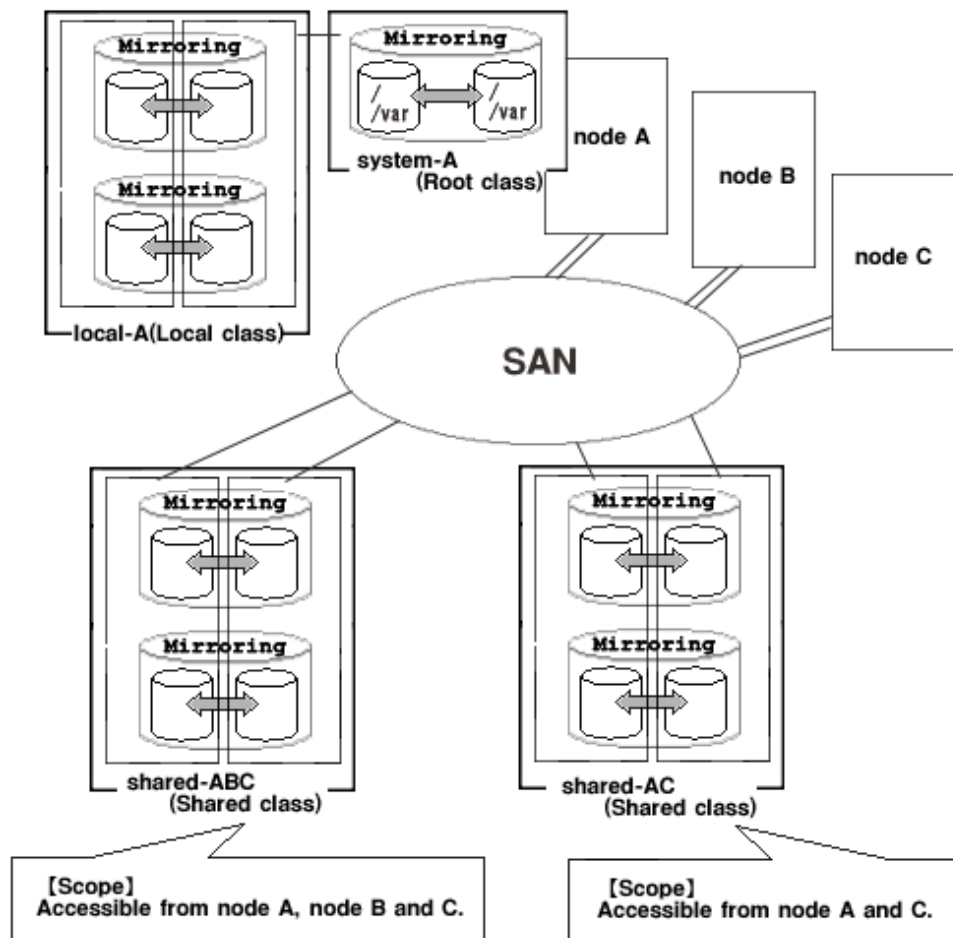
[A.2.9 System Disk Mirroring](#)

[A.2.25 Operation in Cluster System](#)

[A.2.26 Changing Over from Single Nodes to a Cluster System](#)

[A.2.29 Configuration Tables](#)

Figure 2.1 Disk Class



2.3 SDX Disk

Physical disks managed by GDS are registered with a certain class and are called SDX disks, or simply "disks."

Disks that no longer need to be managed with GDS can be deleted from the class and are returned to their original physical disk status.

Attributes

A disk has the following attributes.

Name

This attribute identifies the disk within a class.

Type

This attribute indicates the type of disk. Either of the following values will be set.

Mirror

The disk has been connected to a mirror group.

Stripe

The disk has been connected to a stripe group.

Concatenation

The disk has been connected to a concatenation group.

Switch

The disk is connected to a switch group.

Keep

Retains the disk format and data when registered with a class or connected to a group.

Single

A volume can be created without connecting to a group.

Spare

The disk will be used as a spare disk.

Undefined

The disk has been registered with the class without its usage being specified.

Operation

The following operations are available for disks.

Create

A disk can be created by registering a physical disk with a certain class, using the `sdxdisk -M` command.
For details on GDS Management View, see "[5.2.3.1 Class Configuration](#)."

Delete

A disk can be deleted from a class with the `sdxdisk -R` command.
For details on GDS Management View, see "[5.5.4 Removing a Class](#)."

Connect

A disk can be added to a certain group with the `sdxdisk -C` command.
For details on GDS Management View, see "[5.4.2 Group Configuration](#)."

Disconnect

A disk can be disconnected from the group with the `sdxdisk -D` command.
For details on GDS Management View, see "[5.4.2 Group Configuration](#)."

Swap

A disk is made ready for disk swap with the `sdxswap -O` command.
For details on GDS Management View, see "[5.3.4 Disk Swap](#)."

Recover

Swapped disks can be recovered with the `sdxswap -I` command.
For details on GDS Management View, see "[5.3.4 Disk Swap](#)."

Status Display

The disk status can be displayed with the `sdxinfo` command.
For details on GDS Management View, see "[5.3.1 Viewing Configurations/Statures and Monitoring Statures](#)."

Change Attributes

Disk attribute values can be changed with the `sdxattr -D` command.
For details on GDS Management View, see "[5.4.1 Class Configuration](#)."

Clear Errors

The I/O error status of SDX disks can be cleared with the `sdxfix -D` command.
The operation from GDS Management View is unsupported.

Status

A disk can have the following status.

ENABLE

The disk is in the normal status.

Unless you conduct a special operation, or a special event occurs, a disk is usually in the ENABLE status.

DISABLE

Data on the disk is inaccessible.

The configuration information of disks is checked on system startup, and if any disk is determined as a disk not managed with GDS, it is DISABLED. If this happens, all slices related to the disk are made the NOUSE status.

For details, see "[F.1.2 Disk Status Abnormality](#)."

SWAP

Data on the disk is inaccessible and the disk is ready for disk swap.

Successful completion of the `sdxswap -O` command will make the disk the SWAP status. If this happens, all slices related to the disk are made the NOUSE status.

Reference

There are the following points of concern for disks.

Rules

[A.1.3 Number of Disks](#)

[A.1.6 Number of Keep Disks](#)

Important Points

[A.2.3 Restraining Access to Physical Special File](#)

[A.2.5 Initializing Disk](#)

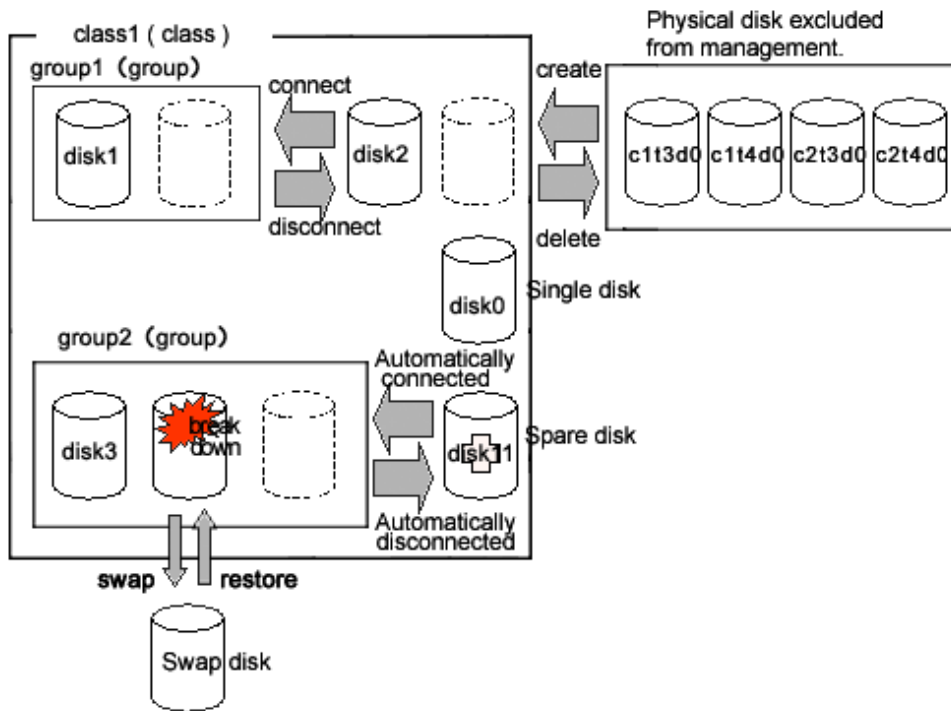
[A.2.6 Disk Size](#)

[A.2.8 Hot Spare](#)

[A.2.10 Keep Disk](#)

[A.2.11 GFS Partition on a Keep Disk](#)

Figure 2.2 SDX Disk



2.4 Disk Group

Disk group acts as container for disks that will be mirrored, striped, concatenated or switched. They may be simply referred to as "groups."

Among group attributes is a type attribute. Type attribute values may be "mirror", "stripe", "concat" or "switch". By connecting multiple disks to a group, the disks within that group can be mirrored, striped, concatenated or switched one another, depending on the group's type attribute value.

You can also create a hierarchy of groups by connecting a group to another group. A group belonging to another group is called a lower level group, and a group above that group is called a higher level group.

A group belonging to another group will retain its own function specified by type attribute, while also acting as a disk belonging to that group. For example, when more than one "stripe" type group is connected to a "mirror" type group, the connected "stripe" type groups will be mirrored to each other. So the disks belonging to the "stripe" type group will be striped and mirrored.

Attributes

A group has the following attribute.

Name

This attribute identifies the group within a class.

Type

This attribute specifies the type of group. You can set to one of the followings.

mirror

Disks and lower level groups belonging to the same group will be mirrored to each other. A maximum of 8 disks and lower level groups can be connected to a group collectively. In other words, a maximum of eight-way multiplex mirroring is supported.

stripe

Disks and lower level groups connected to the same group will each configure a stripe column, and will be striped. Since a maximum of 64 columns can be connected, you can stripe across a maximum of 64 columns.

concat

Disks connected to a concatenation group will be concatenated. Since a maximum of 64 disks can be connected, a maximum of 64 disks can be concatenated.

switch

The group consists of one active disk and one or less inactive disk. If an inactive disk is connected, the roles of the active and inactive disks can be switched.

Stripe Width

This is an attribute of "stripe" type group, which indicates the data unit size when data is striped.

Active Disk

This attribute indicates the active disk between disks that are connected to the switch type group.

Operation

The following operations are available for groups.

Create

Indicating a new (higher level) group when connecting a disk with the `sdxdisk -C` command, or connecting a group with `sdxgroup -C` command will automatically create a (higher level) group.

For details on GDS Management View, see "[5.2.3.3 Group Configuration.](#)"

Delete

Disconnecting the only remaining disk with the `sdxdisk -D` command, or disconnecting the only remaining lower level group with the `sdxgroup -D` command will automatically remove the (higher level) group. You can also delete a group using the `sdxgroup -R` command.

For details on GDS Management View, see "[5.5.3 Removing a Group.](#)"

Connect

You can add a group to another group with the `sdxgroup -C` command. A group that is connected to another group is called a lower level group, and a group to which the group is added is called a higher level group.

For details on GDS Management View, see "[5.4.2 Group Configuration.](#)"

Disconnect

The `sdxgroup -D` command disconnects the lower level group from the higher level group.

For details on GDS Management View, see "[5.4.2 Group Configuration.](#)"

Status Display

Group status can be displayed with the `sdxinfo` command.

For details on GDS Management View, see "[5.3.1 Viewing Configurations/Statures and Monitoring Statures.](#)"

Change Attributes

Group attribute values can be changed with the `sdxattr -G` command.

For details on GDS Management View, see "[5.4.2 Group Configuration.](#)"

Reference

There are the following points of concern for groups.

Rules

[A.1.4 Number of Groups](#)

A.1.7 Creating Group Hierarchy

Important point

A.2.27 Disk Switch

Guide

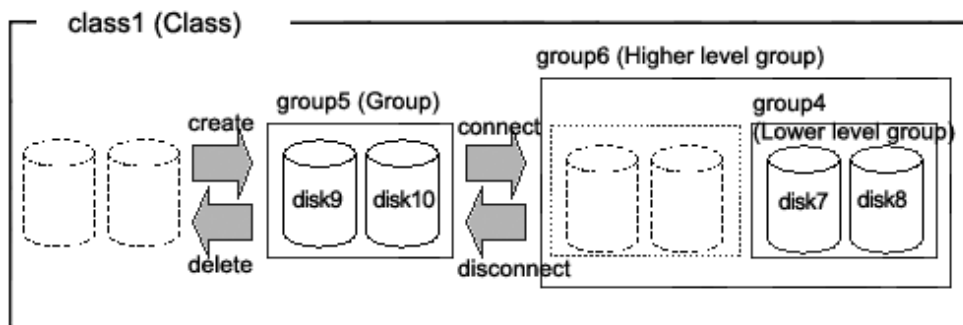
A.3.1 Guidelines for Mirroring

A.3.2 Guidelines for Striping

A.3.3 Guidelines for Concatenation

A.3.4 Guidelines for Combining Striping with Mirroring

Figure 2.3 Disk Group



2.5 Logical Volume

GDS provides mirroring function and unified manageability by virtualizing physical disks as logical volumes.

An application will access the logical volume instead of the physical disk.

Logical volumes are also called "volumes."

There are five kinds of volumes as follows:

- Single Volume Created in a Single Disk

Data will not be redundant. Single volumes are used when it does not necessary to be mirrored, but needs to be managed by GDS. By connecting the single disk to a mirror group, a single volume can be changed to a mirror volume while retaining its data.

- Mirror Volume Created in a Mirror Group

When multiple disks and lower groups are connected to the mirror group, data will be redundant through the mirroring function.

- Stripe Volume Created in a Stripe Group

The striping function enables I/O load sharing across multiple disks. Data will not be redundant.

- Volume Created in a Concatenation Group

Concatenation group allows users to create a large capacity volume spanning multiple disks. Its data will not be redundant.

- Switch Volume Created in a Switch Group

If an inactive disk is connected to a switch group in addition to the active disk, the active disk can be changed to the inactive disk with the disk switch function.

A volume created in the highest level group of a hierarchy will feature functions of its lower level groups as well. For example, a mirror volume created in a mirror group to which more than one stripe group are connected will feature both the I/O load distribution by the striping function and data redundancy by the mirroring function.

In Solaris, physical disks can only be partitioned to a maximum of 7 disk slices (physical slices). GDS allows users to create both a volume with a corresponding physical slice and a volume without a corresponding physical slice. So if you total the volumes with and without physical slices, single disks and groups can be partitioned to a maximum of 1024 volumes.

Attributes

A volume has the following attributes.

Name

This attribute identifies the volume within a class.

JRM

This attribute indicates the just resynchronization mechanism mode for volumes.

on

JRM for volumes is enabled.

off

JRM for volumes is disabled.

Lock Volume

This attribute sets the "Lock volume" mode. The value can be set to one of the followings.

on

The volume will be locked and prevented from activating.

off

The volume will not be locked.

Access Mode

This attribute sets the default access mode. If a volume is activated without specifying the access mode, the default setting will be applied. You can set to one of the followings.

rw

The default access mode is set to read and write.

ro

The default access mode is set to read only.

Physical Slice

This attribute indicates whether the volume has a physical slice or not. In other words, it indicates if the slice composing the volume is registered with the VTOC. The value can be set to one of the followings. However, note that physical slice attribute for volumes that are created in stripe group or concatenation group must be set to "off."

on

When the volume is a single volume, a slice in the single disk will be registered with the VTOC. When the volume is a mirror volume, and if there are disks directly connected to the mirror group, the slices on the disks will be registered with the VTOC. When the volume is a switch volume, slices on all of the disks connected to the switch group will be registered with VTOC. Mirror volumes created in mirror groups to which only lower level groups are connected have no physical slices even if this attribute is "on".

off

The volume has no physical slice. In other words, no slice in the volume is registered to the VTOC.

PJRM

This attribute indicates the just resynchronization mechanism mode for proxies. Either of the following values can be set.

on

JRM for proxies is enabled.

off

JRM for proxies is disabled.

Operation

The following operations are available for volumes.

Create

A volume can be created in a single disk or the highest level group with the `sdxvolume -M` command.

For details on GDS Management View, see "[5.2.3.4 Volume Configuration](#)."

Delete

Volume will be deleted with the `sdxvolume -R` command.

For details on GDS Management View, see "[5.5.2 Removing a Volume](#)."

Start

Volume will be started with the `sdxvolume -N` command.

For details on GDS Management View, see "[4.2.3 Operation](#)."

Stop

Volume will be stopped with the `sdxvolume -F` command.

For details on GDS Management View, see "[4.2.3 Operation](#)."

Expand

The volume size can be expanded with the `sdxvolume -S` command.

The operation from GDS Management View is unsupported.

Status Display

Volume status can be displayed with the `sdxinfo` command.

For details on GDS Management View, see "[5.3.1 Viewing Configurations/Statuses and Monitoring Statuses](#)."

Change Attributes

Volume attribute values can be changed with the `sdxattr -V` command.

For details on GDS Management View, see "[5.4.3 Volume Configuration](#)."

Copy control

Synchronization copying of volume can be controlled with the `sdxcopy` command.

For details on GDS Management View, see "[5.3.6 Copying Operation](#)."

Recover

Recovery attempt of an abnormal volume can be performed with the `sdxfix -V` command.

The operation from GDS Management View is unsupported.

Status

A volume can have the following status.

ACTIVE

Valid data is accessible.

After a volume is booted normally, it is given ACTIVE status. Here, there are one or more slices with ACTIVE or COPY (in the process of background copying) status within the volume.

STOP

Data is inaccessible, but the volume can be booted and made ACTIVE.

However, a proxy volume must first be parted from the master volume before activating. After the volume is stopped normally, it is given the STOP status. Here, there are one or more slices with STOP or COPY (in process of background copying) status within the volume.

INVALID

Data is invalid and inaccessible.

Here, the volume cannot be activated since there are no valid slices (ACTIVE or STOP) or slices in the COPY status (in process of background copying) within the volume.

For details, see "[F.1.3 Volume Status Abnormality](#)."

Reference

There are the following points of concern for volumes.

Rules

[A.1.5 Number of Volumes](#)

Important Point

[A.2.7 Volume Size](#)

[A.2.14 Just Resynchronization Mechanism \(JRM\)](#)

[A.2.24 Volume Access Mode](#)

Figure 2.4 Logical Volume


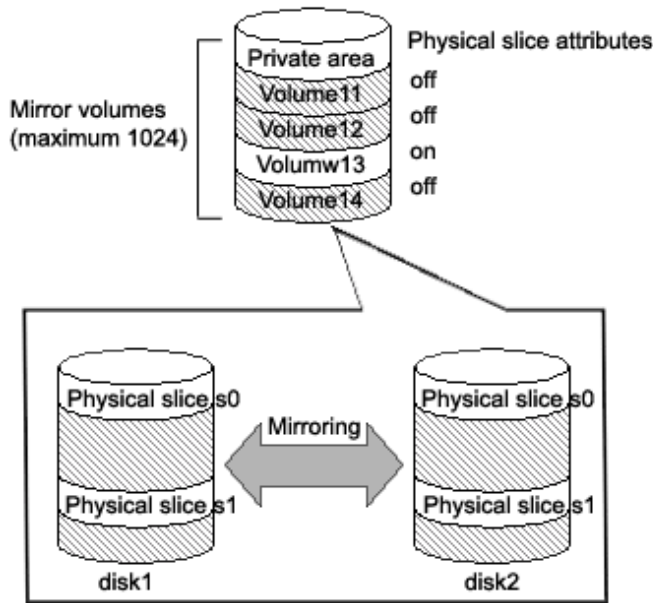


Figure 2.5 Logical Partitioning (Physical Slice Attribute)



2.6 Logical Slice

Logical slices are the smallest objects managed by GDS. A logical slice is also simply referred to as "slice."

Mirror volume is created by mirroring more than one logical slice. Each mirror volume will have one logical slice for every disk and lower level group belonging to the mirror group.

A single volume, stripe volume and a volume in concatenation group will each consist of one logical slice.

A switch volume consists of one logical slice, which is on the active disk.

Usually, an application cannot access the slices consisting a volume. But in the case of a mirror volume with its physical slice attribute set to "on," temporarily detaching the slice will allow the application to directly access the slice. A snapshot (replication) of the volume can be created with this function.

Attributes

A slice has the following attributes.

Name

This attribute identifies the slice within a class.

JRM

This attribute indicates the just resynchronization mechanism mode for slices.

on

JRM for slices is enabled.

off

JRM for slices is disabled.

Access Mode

This attribute sets the access mode. The value can be set to one of the followings.

rw

Sets access mode for read and write.

ro

Sets access mode for read only.

Operation

The following operations are available for slices.

Detach

A snapshot of a mirror volume can be created by detaching one of the slices consisting the mirror volume with the `sdxslice -M` command. For details on GDS Management View, see "[5.3.2.1 Backup \(by Slice Detachment\)](#)."

Attach

The temporarily detached slice can be attached to the volume again with the `sdxslice -R` command. For details on GDS Management View, see "[5.3.2.1 Backup \(by Slice Detachment\)](#)."

Activate

The detached slice can be activated with the `sdxslice -N` command. For details on GDS Management View, see "[5.3.2.1 Backup \(by Slice Detachment\)](#)."

Stop

The detached slice can be stopped with the `sdxslice -F` command. For details on GDS Management View, see "[5.3.2.1 Backup \(by Slice Detachment\)](#)."

Take over

Detached slice can be taken over from another node with the `sdxslice -T` command. The operation from GDS Management View is unsupported.

Status Display

Slice status can be displayed with the `sdxinfo` command. For details on GDS Management View, see "[5.3.1 Viewing Configurations/Statures and Monitoring Statures](#)."

Change Attributes

Attribute values of the detached slice can be changed with the `sdxattr -S` command. The operation from GDS Management View is unsupported.

Status

A slice can have the following status.

ACTIVE

Data is normal, and accessible.

Here, the volume is ACTIVE. The total number of slices that are in the ACTIVE, STOP, or COPY status (in process of background copying) within the mirror volume indicates the present multiplicity of the mirroring configuration (1 to 8).

STOP

Data is normal, but inaccessible.

Here, the volume is in the STOP status. The total number of slices that are in the STOP, ACTIVE, or COPY (in process of background copying) status within the mirror volume indicates the present multiplicity of the mirroring configuration (1 to 8).

INVALID

Data is invalid and inaccessible.

When an I/O error is detected in a slice during the mirroring process, the slice becomes INVALID. Promptly identify the problem and

restore the slice based on the disk driver message or the like. After restoring the slice, execute the synchronization copying. When it ends successfully, the slice will be valid (the slice will be made the ACTIVE or STOP status). If it fails, it will be INVALID again.

COPY

Synchronization copying is in process to ensure the data integrity.

Synchronization copying is executed from a slice in the valid status (ACTIVE or STOP) to a slice in the COPY status. When it ends successfully, the slice will be made valid (ACTIVE or STOP). If it fails, it will be INVALID.

TEMP

The slice is temporarily detached from the volume and can be accessed independently.

When the `sdxslice -M` command ends successfully, the slice is made the TEMP status. To restore the status (ACTIVE or STOP), execute the `sdxslice -R` command.

TEMP-STOP

The slice is temporarily detached from the mirror volume, but cannot be accessed separately.

To make the slice accessible, activate the slice executing the `sdxslice -N` command or the `sdxslice -T` command.

NOUSE

The slice is inaccessible for a special reason.

When the disk is not available for operation (DISABLE or SWAP), all slices related to the disk are made the NOUSE status.

Reference

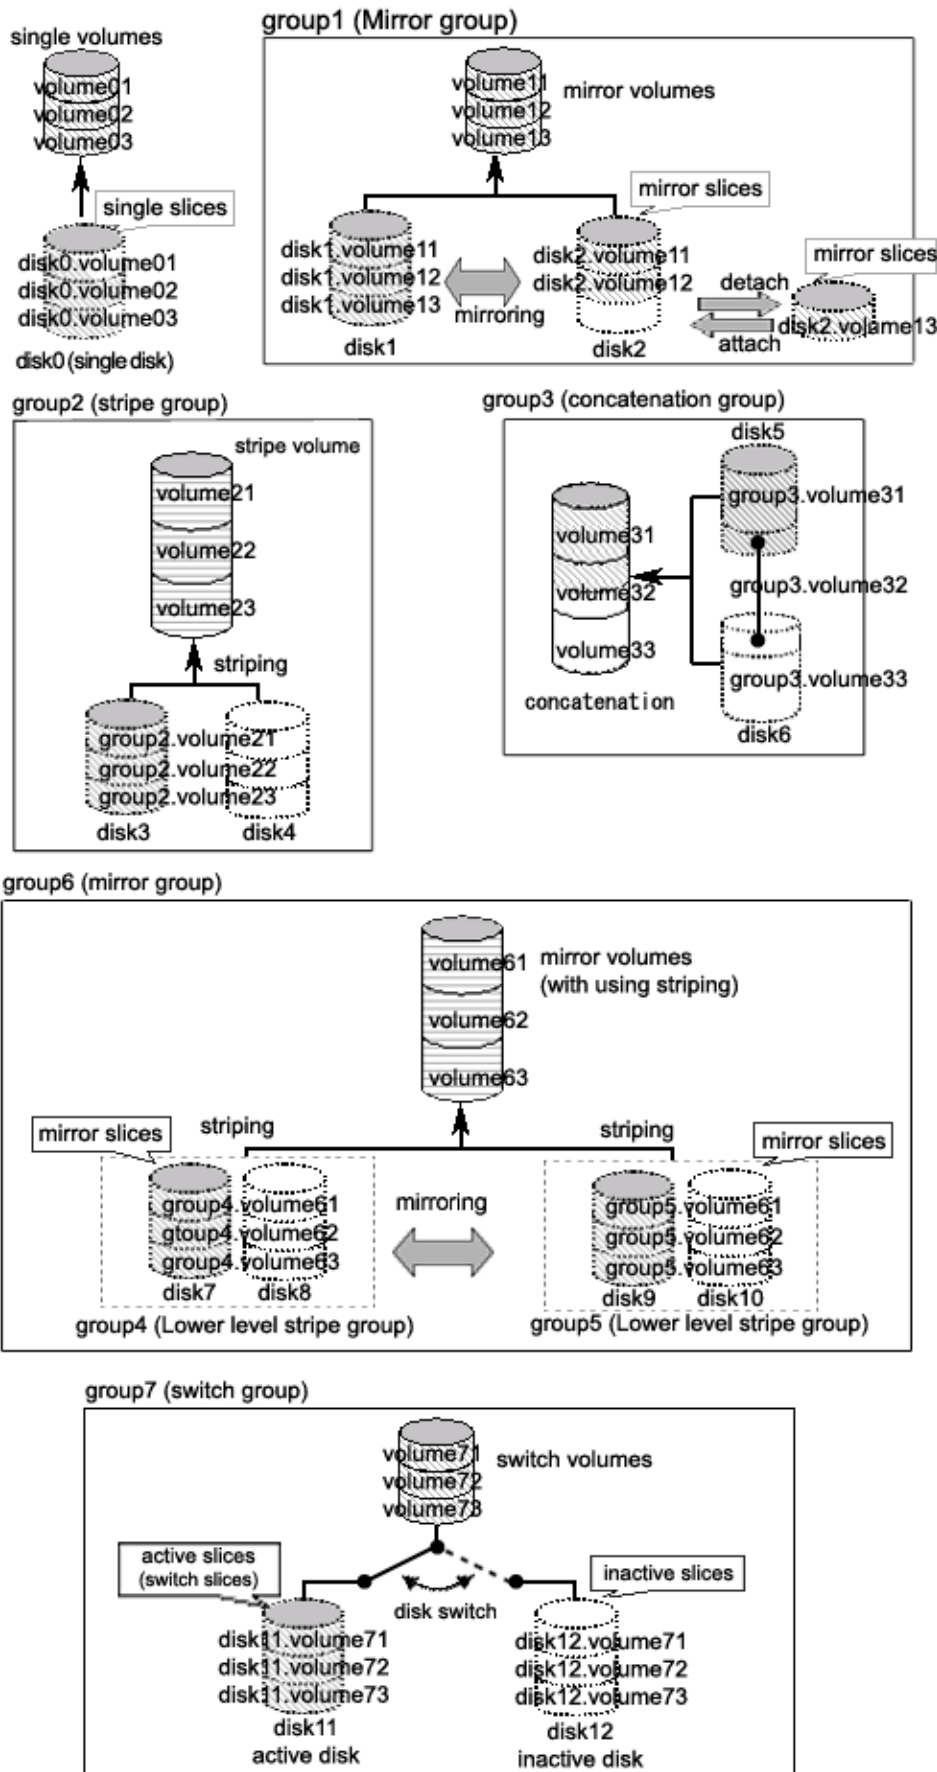
There are the following points of concern for slices.

Important Points

[A.2.12 Creating a Snapshot by Slice Detachment](#)

[A.2.14 Just Resynchronization Mechanism \(JRM\)](#)

Figure 2.6 Logical Slice



2.7 GDS Snapshot Objects

2.7.1 Proxy Object

An SDX object which is related to another SDX object (called master object) and plays the role of substitution of the master object is called a proxy object.

There are two kinds of proxy objects. The one is a proxy volume and the other is a proxy group.

A proxy volume can be joined and synchronized to a master volume, and they can be temporarily parted to become accessible as separate volumes. By joining and synchronizing the proxy volume to the master volume and then parting them, snapshot (copy at the moment) data of the master volume is instantly taken into the proxy volume.

A proxy group which is related to a master group has the same volume configuration as the master group. Each volume in the proxy group is a proxy volume of the corresponding master volume. By joining and synchronizing the proxy group to the master group and then parting them, snapshots of all the volumes in the master group are taken into the proxy group at a time.

Operation

The following operations are available for proxy objects.

Join

A pair of specified master and proxy volumes or master and proxy groups will be related and synchronized with the `sdxproxy Join` command.

For details on GDS Management View, see "[5.2.5.1 Join](#)."

Part

Master-proxy synchronization will be canceled and the proxy will be accessible as another device other than the master using the `sdxproxy Part` command. The master-proxy relationship will be maintained. The parted proxy will be a snapshot (a replication) of the master at that moment.

For details on GDS Management View, see "[5.3.2.2 Backup \(by Synchronization\)](#)."

Rejoin

The parted master and proxy will be resynchronized with the `sdxproxy Rejoin` command.

For details on GDS Management View, see "[5.3.2.2 Backup \(by Synchronization\)](#)."

Rejoin and Restore

The parted proxy is rejoined with the master, and master volume data will be restored using the proxy with the `sdxproxy RejoinRestore` command.

For details on GDS Management View, see "[5.3.3 Restore](#)."

Swap

The slices composing the master and those composing the proxy will be exchanged with the `sdxproxy Swap` command.

For details on GDS Management View, see "[5.3.5 Disk Migration](#)."

Relate

A pair of a master volume and a proxy volume or a master group and a proxy group can be related and parted, with the `sdxproxy Relate` command.

For details on GDS Management View, see "[5.2.5.2 Relate](#)."

Update

Data can be copied (overwritten) from a master to a parted proxy with the `sdxproxy Update` command. The updated proxy becomes a snapshot (a replica) of the master at the moment.

For details on GDS Management View, see "[5.3.2.3 Backup \(by OPC\)](#)."

Restore

Data from a parted proxy can be restored back to the master with the `sdxproxy Restore` command. The proxy data at the moment is copied (overwritten) to the master.

For details on GDS Management View, see "[5.3.3 Restore](#)."

Cancel Copy Session

The session of a copy function of a disk unit residing between the master and the parted proxy can be canceled with the `sdxproxy Cancel` command.

The operation from GDS Management View is unsupported.

Configure an Alternative Boot Environment

An environment can be set up with the `sdxproxy Root` command so that the system can be booted with the parted master or proxy.

The operation from GDS Management View is unsupported.

Break

The specified relationship between a pair of volumes or groups as the master and proxy will be cancelled and they will return to individual objects with the `sdxproxy Break` command.

For details on GDS Management View, see "[5.5.6 Breaking a Proxy](#)."

Status

The relationship between master volumes and proxy volumes will be one of the following statuses. These statuses can be viewed in the PROXY field for volumes displayed with the `sdxinfo` command.

Joined

Master and proxy are joined. In this status, the proxy cannot be accessed.

Parted

Proxy is parted from the master and can be accessed independently from the master (unless the volume is stopped explicitly).

Reference

There are the following points of concern for proxy objects.

Rules

[A.1.8 Proxy Configuration Preconditions](#)

[A.1.9 Number of Proxy Volumes](#)

[A.1.10 Proxy Volume Size](#)

[A.1.11 Proxy Group Size](#)

Important Points

[A.2.13 The Difference between a Mirror Slice and a Proxy Volume](#)

[A.2.14 Just Resynchronization Mechanism \(JRM\)](#)

[A.2.17 Object Operation When Using Proxy](#)

[A.2.18 Using the Advanced Copy Function in a Proxy Configuration](#)

[A.2.19 Instant Snapshot by OPC](#)

[A.2.21 Using EMC TimeFinder or EMC SRDF in a Proxy Configuration](#)

[A.2.22 Ensuring Consistency of Snapshot Data](#)

Figure 2.7 Operating Proxy Objects




Figure 2.8 Swapping Master Volume and Proxy Volume


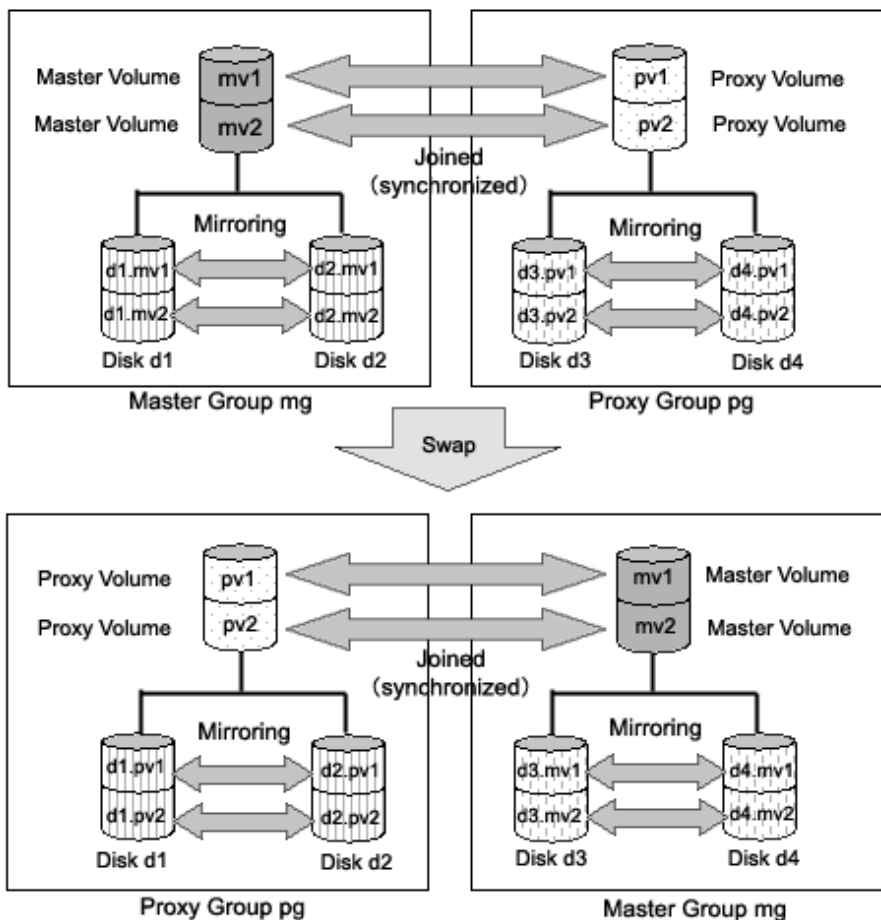


Figure 2.9 Swapping Master Group and Proxy Group



2.7.2 Shadow Object

There are five types of shadow object as follows: shadow classes, shadow disks, shadow groups, shadow volumes and shadow slices. These objects correspond to disk classes, SDX disks, disk groups, logical volumes and logical slices that are SDX objects respectively.

When SDX objects and shadow objects do not have to be classified particularly, they may be called "objects" collectively.

On the server where GDS Snapshot is installed, physical disks that are not registered with disk classes in the same domain and that store SDX disk management data can be managed as shadow objects and accessed. Physical disks conforming to either of the following conditions can be managed as shadow objects.

- Disks registered with a class (local class or shared class) of GDS or SynfinityDisk in another domain and managed as SDX disks
- Disks to which the private slice was copied from an SDX disk registered with a class of GDS or SynfinityDisk in another domain or in the same domain with the disk unit's copy function

Shadow objects have the following characteristics.

- When shadow objects, such as shadow disks and shadow volumes, are configured or broken up, the physical disk format and data remain unchanged. For this reason, shadow volumes can be created without affecting data stored on the physical disks, and the data can be read in or write out through the shadow volumes.
- The configuration information of shadow objects is not saved on the private slices but managed on memory. Shadow objects are cleared by server reboot, but they can be reconfigured. However, if those objects in the same configuration are not created after the server reboot, restoration is necessary.

For details, see ["Rebooting a Node"](#) in ["A.2.28 Shadow Volume."](#)

Shadow objects other than shadow disks can be named as desired. The object names must be unique in the entire system. For details on the shadow disk naming restrictions, see ["A.1.1 Object Name."](#)

2.7.2.1 Shadow Class

A shadow class is the shadow object corresponding to a disk class that is an SDX object.

When disk classes and shadow classes do not have to be classified particularly, they may be called "classes" collectively.

Physical disks conforming to one of the following conditions can be registered with shadow classes.

- Disks registered with a class (local class or shared class) of GDS or SynfinityDisk in another domain and managed as SDX disks
If this applies, with one shadow class, a physical disk registered with the class that has the same name in another domain can be registered. When multiple physical disks are registered with one class in another domain, part or all of those disks can be registered with a shadow class. It is also possible to register part of those disks to a shadow class and the other to another shadow class.
- Disks to which the private slice was copied from an SDX disk registered with a class (local class or shared class) of GDS or SynfinityDisk in another domain or the same domain with the disk unit's copy function
If this applies, with one shadow class, a copy destination physical disk of the SDX disk registered with the class that has the same name can be registered. When multiple physical disks are registered with one class, part or all of such copy destination physical disks can be registered with a shadow class. It is also possible to register part of such copy destination physical disks to a shadow class and the other to another shadow class.

In addition, the sizes of the private slices of physical disks to be registered with the same shadow class must be the same.

Attributes

A shadow class has the following attributes.

Name

This attribute identifies the shadow class in a system.

Type

This attribute indicates the type of shadow class. The following value is set.

Local

The object managed in a shadow class is available only on the current node.

Scope

This attribute indicates a group of nodes on which objects in the shadow class are available. The current node name is set.

Operation

The operations other than status display are not available from Management View. Use the command for each operation.

Create

A new shadow class will be automatically created by specifying a new class name when registering a disk with the `sdxshadowdisk -M` command.

Delete

The shadow class will be automatically deleted by deleting the last registered shadow disk from a shadow class with the `sdxshadowdisk -R` command.

Status Display

The shadow class status can be displayed with the `sdxinfo` command. A shadow class is indicated by 1 (one) in the SHADOW field for class information displayed with the `sdxinfo -e long` command.

For details on GDS Management View, see "[5.3.1 Viewing Configurations/Statures and Monitoring Statures.](#)"

Status

Unlike a disk class, there is no possibility that a shadow class is closed down.

Reference

There are the following points of concern for shadow classes.

Rules

[A.1.1 Object Name](#)

[A.1.2 Number of Classes](#)


Important Points

[A.2.1 Physical Disk Driver](#)

[A.2.25 Operation in Cluster System](#)

The disks that are registered with a class in a certain domain (domain alpha in the figure below) and accordingly are managed as SDX disks can be registered with a shadow class in another domain (domain beta in the figure below) connected to the same SAN.


Figure 2.10 Example of Common Configuration with a Shadow Class



The disk (c1t1d2 in the figure below) to which the private slice has been copied from an SDX disk (c1t1d1 in the figure below) with the disk unit's copy function can be registered with a shadow class on one of the following nodes.

- Node that is included in the scope of the class to which the SDX disk belongs (node1 or node 2 in the figure below)
- Node that belongs to the domain (domain alpha in the figure below) with which the SDX disk is registered, but that is not included in the scope of its class (node 3 in the figure below)
- Node that does not belong to the domain where the SDX disk is registered (domain alpha in the figure below), but that is connected to the same SAN (node 4 in the figure above)

Figure 2.11 Example of Configuration Using Disk Unit's Copy Function and a Shadow Class



Note

Physical Device Name

The same physical device name (such as c1t1d1) is not necessarily assigned to the identical physical device in domain alpha and domain beta.

2.7.2.2 Shadow Disk

A physical disk registered with a shadow class is called a shadow disk. A shadow disk is the shadow object corresponding to an SDX disk that is an SDX object.

When SDX disks and shadow disks do not have to be classified particularly, they may be called "disks" collectively.

Attributes

A shadow disk has the following attributes.

Name

This attribute identifies the shadow disk within a shadow class. There are the following restrictions.

- When it is registered with a class in another domain and managed as an SDX disk, the shadow disk must be named as same as the SDX disk in the said domain.
- When data has been copied on it from an SDX disk with the disk unit's copy function, the shadow disk must be named as same as the copy source SDX disk.

Type

This attribute indicates the type of shadow disk. Either of the following values will be set.

Mirror

The disk has been connected to a mirror group.

Stripe

The disk has been connected to a stripe group.

Concatenation

The disk has been connected to a concatenation group.

Single

A shadow volume can be created without connecting to a shadow group.

Undefined

The disk has been registered with the shadow class without its usage being specified.

Operation

The operation other than status display is not available from Management View. Use the command for each operation.

Create

A disk will be created by registering a physical disk with a certain shadow class using the `sdxshadowdisk -M` command.

Delete

The shadow disk can be deleted from a shadow class with the `sdxshadowdisk -R` command.

Connect

The shadow disk can be added to a certain shadow group with the `sdxshadowdisk -C` command.

Disconnect

The shadow disk will be disconnected from the shadow group with the `sdxshadowdisk -D` command.

Status Display

The shadow disk status can be displayed with the `sdxinfo` command.

For details on GDS Management View, see "[5.3.1 Viewing Configurations/Statuses and Monitoring Statuses.](#)"

Status

A shadow disk can have the following status.

ENABLE

The shadow disk is in the normal status.

Reference

There are the following points of concern for shadow disks.

Rules

[A.1.3 Number of Disks](#)

Important Points

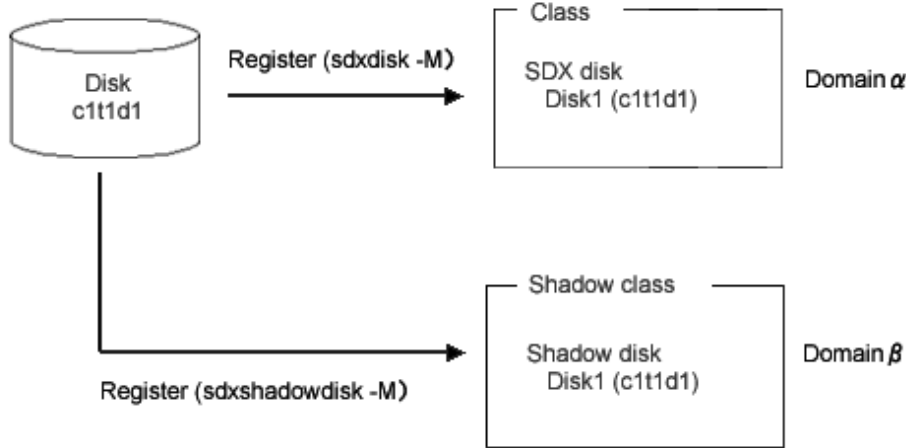
[A.2.3 Restraining Access to Physical Special File](#)

[A.2.5 Initializing Disk](#)

A.2.6 Disk Size

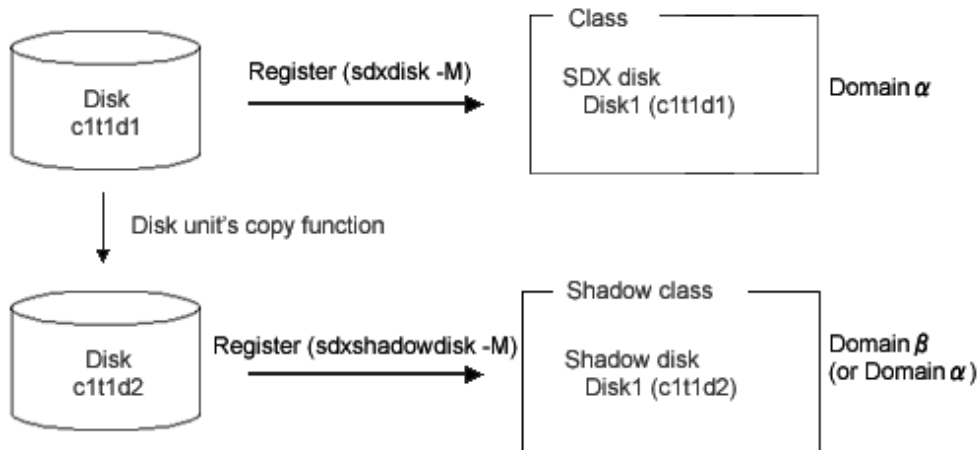
The physical disk that is registered with a class in a certain domain (domain alpha) with the `sdxdisk -M` command and accordingly is managed as an SDX disk can be registered with a shadow class in another domain (domain beta) with the `sdxshadowdisk -M` command and accordingly be managed as a shadow disk. The shadow disk must be named as same as the SDX disk.

Figure 2.12 An SDX Disk and a Shadow Disk



The physical disk (c1t1d2 in the figure below) to which the private slice has been copied with the disk unit's copy function from the physical disk that is registered with a class in a certain domain (domain alpha) with the `sdxdisk -M` command and accordingly is managed as an SDX disk (c1t1d1 in the figure below) can be registered with a shadow class in the same domain (domain alpha) or in another domain (domain beta) with the `sdxshadowdisk -M` command and accordingly be managed as a shadow disk. The shadow disk must be named as same as the SDX disk.

Figure 2.13 Disk Unit's Copy Function and a Shadow Disk



Note

Physical Device Name

The same physical device name (such as c1t1d1) is not necessarily assigned to the identical physical device in domain alpha and domain beta.

2.7.2.3 Shadow Group

A shadow group is the shadow object corresponding to a disk group that is an SDX object.

When disk groups and shadow groups do not have to be classified particularly, they may be called "groups" collectively.

To access data on a logical volume through a shadow volume, a shadow group must be created in the same configuration as that of the disk group to which the logical volume belongs. The following configurations must be identical.

- Group type. This can be viewed in the TYPE field displayed with the `sdxinfo -G -e` long command.
- Disks constituting the group and disks constituting the lower level group. The disk unit's copy function destination disk is also available. These can be viewed with the `sdxinfo -D` command and the `sdxinfo -G` command.
- Order of connecting disks and lower level groups when the group is the stripe type or the concatenation type. This can be viewed in the DISKS field displayed with the `sdxinfo -G` command.
- Stripe width when the group is the stripe type. This can be viewed in the WIDTH field displayed with the `sdxinfo -G -e` long command.

In addition, a shadow volume can be created by connecting only one of the following disks to a shadow group of the mirror type, and thus data on a single volume or on a slice temporarily detached can be accessed.

- Single disk. The disk unit's copy function destination disk is also available.
- Disk containing a slice temporarily detached (in the TEMP status). The disk unit's copy function destination disk is also available.

Attributes

A shadow group has the following attributes.

Name

This attribute identifies the shadow group within a shadow class.

Type

This attribute indicates the type of shadow group. One of the following values can be set.

mirror

Shadow disks and lower level shadow groups belonging to the shadow group will be mirrored to one another. A maximum of 8 disks and lower level shadow groups may be connected to a shadow group in total. In other words, a maximum of eight-way multiplex mirroring is supported.

stripe

Shadow disks and lower level shadow groups connected to the shadow group will configure stripe columns respectively and will be striped. Since a maximum of 64 shadow disks and lower level shadow groups can be connected to a shadow group in total, striping across a maximum of 64 columns is supported.

concat

Shadow disks connected to the shadow group are concatenated. Since a maximum of 64 shadow disks can be connected to a shadow group, a maximum of 64 shadow disks can be concatenated.

Stripe Width

This is an attribute of shadow groups of the "stripe" type, which indicates the data unit size when data is striped.

Operation

The operation other than status display from Management View is unsupported. Use the command for each operation.

Create

A shadow group (higher level) will automatically be created by specifying a new (higher level) shadow group name when connecting a shadow disk with the `sdxshadowdisk -C` command, or connecting a shadow group with the `sdxshadowgroup -C` command.

Delete

The shadow group (higher level) will be automatically removed by disconnecting the only remaining shadow disk with the `sdxshadowdisk -D` command, or by disconnecting the only remaining lower level shadow group with the `sdxshadowgroup -D` command. You can also delete a shadow group with the `sdxshadowgroup -R` command.

Connect

A shadow group will be added to another shadow group with the `sdxshadowgroup -C` command. A shadow group that is connected to another shadow group is particularly called a lower level shadow group, and a shadow group to which another shadow group is connected is particularly called a higher level shadow group.

Disconnect

The lower level shadow group will be disconnected from the higher level shadow group with the `sdxshadowgroup -D` command.

Status Display

The shadow group status can be displayed with the `sdxinfo` command.

For details on GDS Management View, see "[5.3.1 Viewing Configurations/Statures and Monitoring Statures.](#)"

Reference

There are the following points of concern for shadow groups.

Rules

[A.1.4 Number of Groups](#)

[A.1.7 Creating Group Hierarchy](#)

2.7.2.4 Shadow Volume

A volume created in a shadow group or a shadow disk of the single type is called a shadow volume. The users and applications will access data on the shadow volumes instead of the physical disks. A shadow volume is the shadow object corresponding to a logical volume that is an SDX object.

When logical volumes and shadow volumes do not have to be classified particularly, they may be called "volumes" collectively.

Create shadow volumes conforming to the following rules in order to use the shadow volumes to access data in the corresponding logical volumes.

- Equal to the corresponding logical volumes in size. For volume sizes, check the `BLOCKS` field displayed with the `sdxinfo -V` command.
- Having first block numbers consistent with the first block numbers of the corresponding logical volumes. Therefore, create shadow volumes within the same shadow group or shadow disk in ascending order in conformity to the order of the first block numbers of the corresponding logical volumes. For the first block numbers of volumes, check the `1STBLK` field displayed with the `sdxinfo -V` command.

Synchronization copying is not conducted when a shadow volume of the mirror type is created. When a shadow volume corresponding to a mirror volume is to be created, synchronization of the mirror volume must be ensured in advance with GDS or SynfinityDisk managing that mirror volume.

Shadow volumes and the corresponding logical volumes are managed independently. For example, the change on the slice status in one volume is not updated on the slice status in the other volume. For this reason, you must note certain operational particulars when using shadow volumes. For details, see "[A.2.28 Shadow Volume.](#)"

Attributes

A shadow volume has the following attributes.

Name

This attribute identifies the shadow volume within a shadow class.

JRM

This attribute identifies the just resynchronization mechanism mode. The following value is set.

off

JRM is inactive.

Lock Volume

This attribute identifies the "Lock volume" mode. The following value is set.

off

The volume will not be locked.

Access Mode

This attribute sets the default access mode. If a volume is activated without specifying the access mode, the default setting will be applied. The following value is set.

ro

The default is set to read only.

The access mode attribute value cannot be set to rw (read and write). To write data on a shadow volume, the shadow volume must first be stopped and reactivated in the read and write mode using the `sdxshadowvolume -N` command with the `-e mode=rw` option.

Physical Slice

This attribute is always set to off regardless whether the shadow volume has a physical slice, which means the shadow slice is registered with VTOC.

Operation

The operation other than status display from Management View is unsupported. Use the command for each operation.

Create

A shadow volume can be created in the highest level shadow group or a shadow disk of the single type with the `sdxshadowvolume -M` command.

Delete

The shadow volume can be deleted with the `sdxshadowvolume -R` command.

Start

The shadow volume can be started with the `sdxshadowvolume -N` command.

Stop

The shadow volume can be stopped with the `sdxshadowvolume -F` command.

Status Display

The volume status can be displayed with the `sdxinfo` command.

For details on GDS Management View, see "[5.3.1 Viewing Configurations/Statures and Monitoring Statures.](#)"

Status

A shadow volume can have the following status.

ACTIVE

Data is accessible.

After a shadow volume is started normally, it is given ACTIVE status. Here, there are one or more shadow slices in the ACTIVE status in the shadow volume.

STOP

Data is inaccessible, but the shadow volume can be activated and be ACTIVE.

After the shadow volume is stopped normally, it is given STOP status. Here, there are one or more shadow slices in STOP status in the shadow volume.

Reference

There are the following points of concern for shadow volumes.

Rules

[A.1.5 Number of Volumes](#)

Important Point

[A.2.7 Volume Size](#)


[A.2.24 Volume Access Mode](#)

[A.2.28 Shadow Volume](#)

Multiple physical disks virtualized as a logical volume in a certain domain (domain alpha in the figure below) can be virtualized as a shadow volume in another domain (domain beta in the figure below), and the shadow volume can be used in order to access the data on the logical volume in domain alpha from domain beta. The primary service (service A in the figure below) can be run with the logical volume in domain alpha, and another service (service B: for example, backup, restore, and batch processing) can be run with the shadow volume in domain beta. However, service A and service B cannot be run simultaneously. If they are run simultaneously, data consistency between disks is not ensured.

Figure 2.14 A Logical Volume and a Shadow Volume

* Service A and service B cannot be run simultaneously

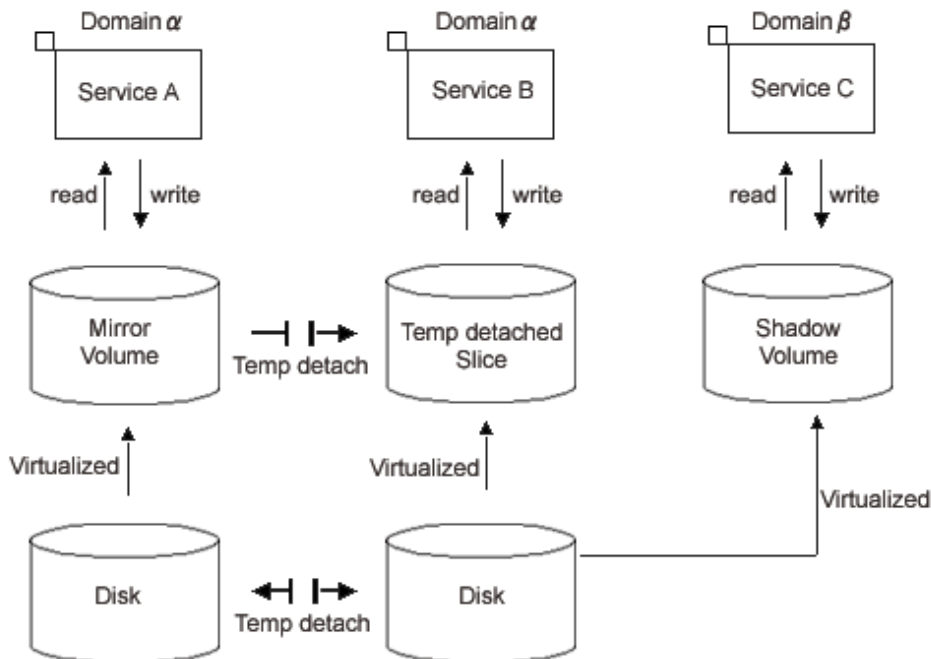


One physical disk temporarily detached from mirroring among multiple physical disks virtualized as a mirror volume in a certain domain (domain alpha in the figure below) can be virtualized as a shadow volume in another domain (domain beta in the figure below), and the

shadow volume can be used in order to access the data on the temporarily detached slice in domain alpha from domain beta. The primary service (service A in the figure below) can be run with the mirror volume in domain alpha from which one slice is temporarily detached, and another service (service C: for example, backup, restore, and batch processing) can be run with the shadow volume in domain beta simultaneously. However, when a different service (service B in the figure below) is run with the temporarily detached slice in domain alpha, service B and service C cannot be run simultaneously. If they are run simultaneously, data consistency between disks is not ensured.

Figure 2.15 Mirror Slice Temporary Detachment and a Shadow Volume

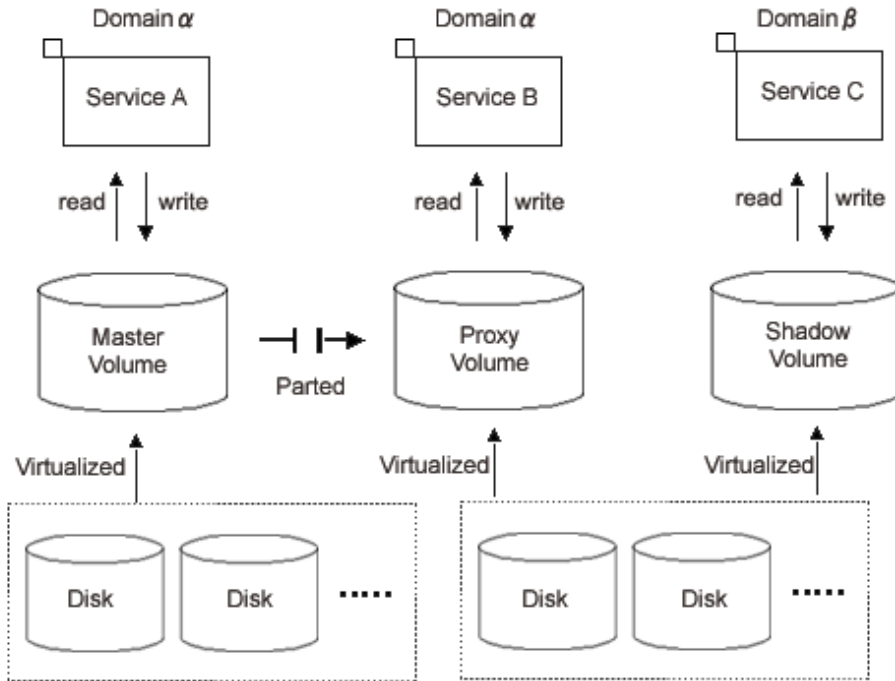
* Service B and service C cannot be run simultaneously



Multiple physical disks virtualized as a proxy volume in a certain domain (domain alpha in the figure below) can be virtualized as a shadow volume in another domain (domain beta in the figure below), and the shadow volume can be used in order to access the data on the proxy volume in domain alpha from domain beta. The primary service (service A in the figure below) can be run with the master volume in domain alpha from which the proxy is parted, and another service (service C: for example, backup, restore, and batch processing) can be run with the shadow volume in domain beta simultaneously. However, when a different service (service B in the figure below) is run with the proxy volume in domain alpha, service B and service C cannot be run simultaneously. If they are run simultaneously, data consistency between disks is not ensured.


Figure 2.16 A Proxy Volume and a Shadow Volume

* Service B and service C cannot be run simultaneously



Data on multiple physical disks virtualized as a logical volume in a certain domain (domain alpha in the figure below) can be copied to other physical disks with the disk unit's copy function, and the copy destination physical disks can be virtualized as a shadow volume in the same domain (domain alpha) or another domain (domain beta in the figure below). The primary service (service A in the figure below) can be run with the logical volume in domain alpha, and another service (service B: for example, backup, restore, and batch processing) can be run in the domain in which the shadow volume was created.

Figure 2.17 Disk Unit's Copy Function and a Shadow Volume



2.7.2.5 Shadow Slice

A shadow slice is a component of a shadow volume. A shadow volume of the mirror type consists of one or more mirrored shadow slices. A shadow volume of the single type, stripe type, or concatenation type consists of one shadow slice. A shadow slice is the shadow object corresponding to a logical slice that is an SDX object.

Independent access to a shadow slice detached from a shadow volume of the mirror type is unsupported.

When logical slices and shadow slices do not have to be classified particularly, they may be called "slices" collectively.

Attributes

A shadow slice has the following attributes.

Name

This attribute identifies the shadow slice within a shadow class.

Operation

The following operations are available for shadow slice.

Status Display

The slice status can be displayed with the `sdxinfo` command.

For details on GDS Management View, see "[5.3.1 Viewing Configurations/Statuses and Monitoring Statuses](#)."

Status

A shadow slice can have the following status.

ACTIVE

Data is accessible.

Here, the shadow volume is in the ACTIVE status. The total number of shadow slices in the ACTIVE or STOP status within a shadow volume of the mirror type indicates the present multiplicity of the mirroring configuration (1 to 8).

STOP

Data is inaccessible.

Here, the shadow volume is in the STOP status. The total number of shadow slices in the ACTIVE or STOP status within a shadow volume of the mirror type indicates the present multiplicity of the mirroring configuration (1 to 8).

INVALID

Data is inaccessible due to an I/O error.

When an I/O error is detected in a mirrored shadow slice, the shadow slice becomes INVALID.

For details, see "[F.1.1 Slice Status Abnormality](#)."

Chapter 3 Starting and Exiting GDS Management View

GDS Management View manages and monitors objects by using the Web browser.

This chapter explains how to start and exit the GDS Management View.

GDS Management View uses Web-Based Admin View / WWW Server for Admin View.

For details on Web-Based Admin View / WWW Server for Admin View, see "Web-Based Admin View Operation Guide."

3.1 Preparation for Starting GDS Management View

In order to start GDS Management View, the following must be completed.

- Decide the user group.
- Set up the client environment.
- Set up the Web environment.

3.1.1 Deciding the User Group

In order to use GDS Management View, you must create user groups that the operating system manages and register the user names on all nodes where GDS are installed.

3.1.1.1 User Group Types

There are two types of user groups.

wvroot

"wvroot" is a Web-Based Admin View administrator group and is created automatically when Web-Based Admin View is installed. This permits availability of all kinds of operation management products running on Web-Based Admin View, such as GDS management, environment setup, logging view, and all.

sdxroot

This is the administrator group for GDS management. This user type can perform GDS Management View operations.

3.1.1.2 Creating User Groups

After installing the software, only user group wvroot will be automatically created.

The other user group sdxroot, explained in "3.1.1.1 User Group Types" must be created as necessary.

User group can be created using the following command.

```
# groupadd sdxroot
```

3.1.1.3 Registering to a User Group

You can register a user name to a user group with a command.

Type an appropriate user group such as sdxroot, and execute the following command.

When registering an existing user to a user group

- When registering a group to "Primary Group"

```
# usermod -g Group name(you wish to register) User name
```

- When registering a group to "Secondary Group"

```
# usermod -G Group name(you wish to register) User name
```

When registering a new user to a user group

- When registering a group to "Primary Group"

```
# useradd -g Group name(you wish to register) User name
```

- When registering a group to "Secondary Group"

```
# useradd -G Group name(you wish to register) User name
```



Note

Registering a User to the wvroot Group

A user registered to the wvroot group will have the equivalent permission as a user assigned to the system administrator group. Only users responsible for the entire system should be registered to this group.

3.1.2 Setting up the Client Environment

For the operating environments, such as hardware, operating systems, and Web browsers of clients on which GDS Management View is used, see "Web-Based Admin View Operation Guide."

3.1.3 Setting up the Web Environment

In order to use GDS Management View, you must define the environment for the client and the Web browser.

For details on setting the Web browser, see "Web-Based Admin View Operation Guide."

3.2 Starting the GDS Management View

Start the GDS Management View after all the necessary settings are complete.

3.2.1 Starting Web-Based Admin View Operation Menu

Follow the procedures below to start Web-Based Admin View.

Procedures

1. Start the Web browser on your client.
2. Access the management server by specifying the following URL

When using Java Plug-in.

http://host name:port number/Plugin.cgi

host name

Specify the public LAN IP address for the primary or secondary management server, or host name.

port number

Specify "8081." When the port number has been changed, specify the new port number.

 See

For details on changing port numbers, see "Web-Based Admin View Operation Guide."

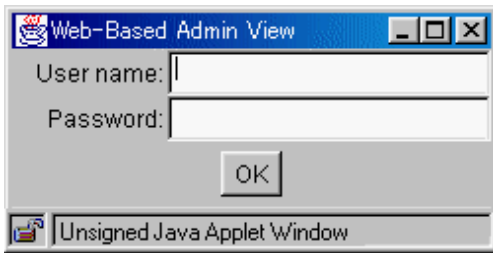
 Note

When the Web-Based Admin View does not start

If you specified the management server's host name for "host name", and the Web-Based Admin View does not start, specify the public LAN IP address instead.

3. After starting the Web-Based Admin View, the following user input screen appears.

Figure 3.1 User Name Input Screen



Type the user name and password for the management server, and click <OK>.

4. After completing authentication, the top menu of Web-Based Admin View appears.

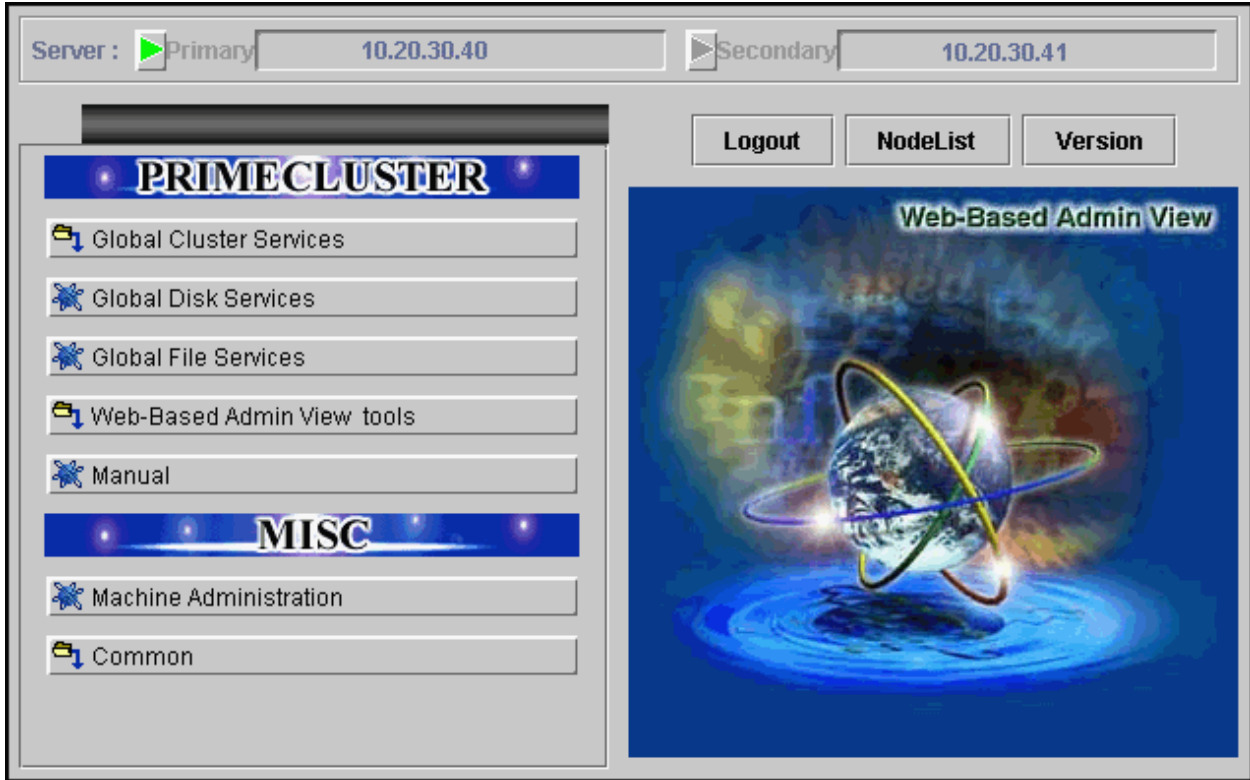
3.2.2 Web-Based Admin View Operation Menu Functions

The Web-Based Admin View screen supports the following facilities.

Menu	Outline
Global Disk Services	Start the GDS Management View.

See "Web-Based Admin View Operation Guide" about other operation menus.

Figure 3.2 Web-Based Admin View Operation Menu (Top Menu)



Note

The state of Web-Based Admin View Operation Menu

- The Web-Based Admin View menu varies depending on the products installed.
- If a dialog is displayed because of a Web-Based Admin View error, the picture on the right area of the screen turns red. Click the red picture, and then a hidden dialog is brought to the front. As any errors come into focus, be sure to make the picture viewable.

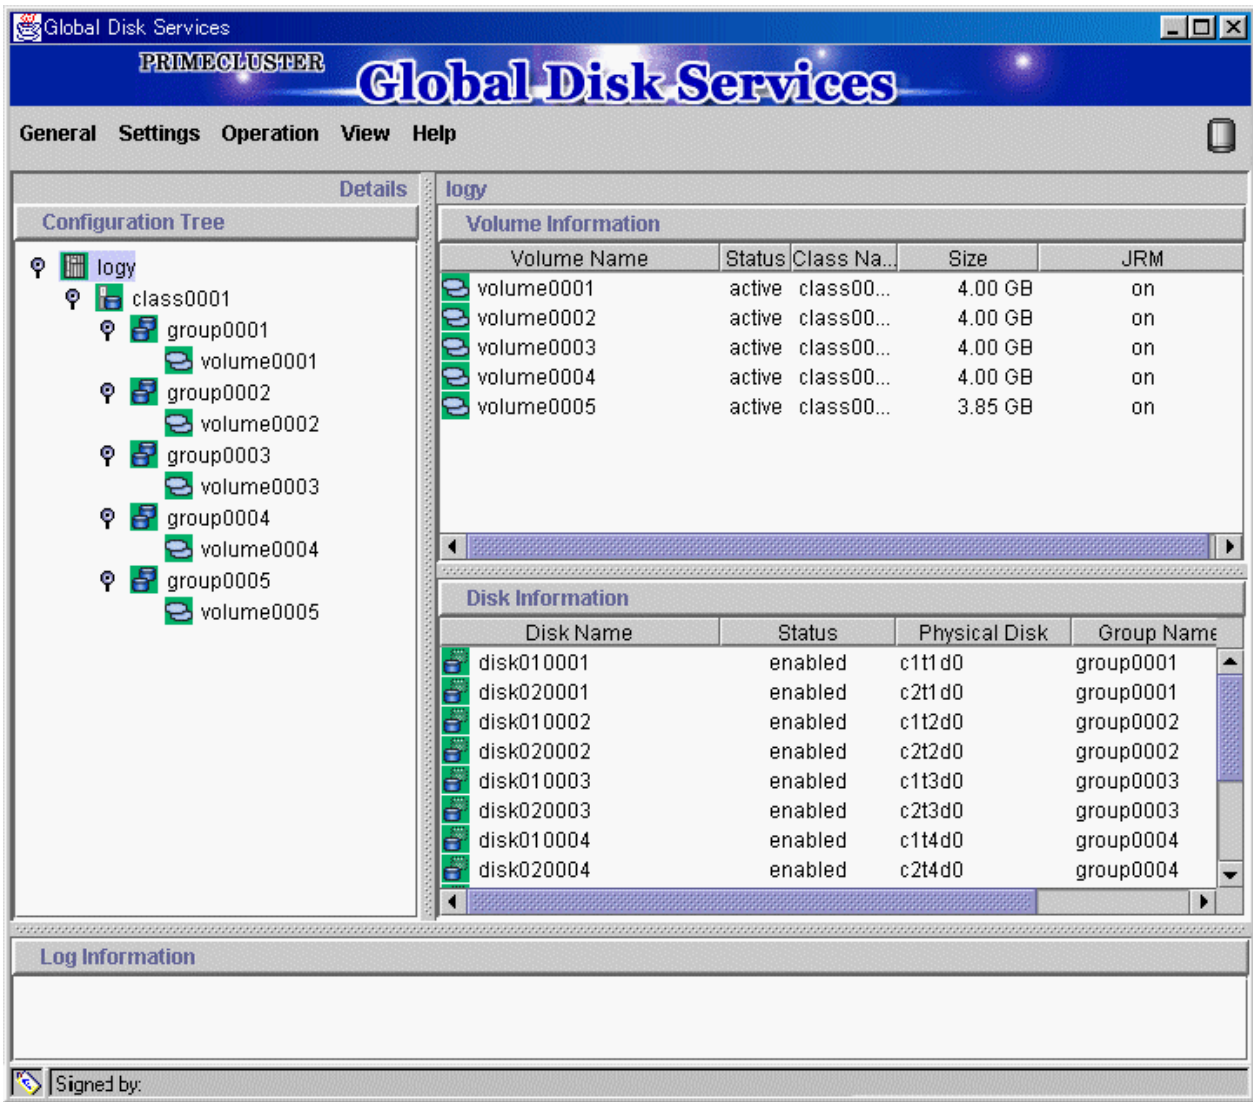
3.2.3 Starting GDS Management View

Click the GDS management icon on the Web-Based Admin View Operation menu to start the GDS Management screen (hereinafter main screen).

From the main screen, you can perform GDS object configuration, such as a class or a volume, status confirmation, and disk swap.

For details, see "[Chapter 5 Operation](#)"

Figure 3.3 GDS Management Main Screen



3.3 Exiting GDS Management View

How to exit the GDS Management View is described below.

On the [General] menu, click [Exit]. The following message appear

Figure 3.4 Exit Message



Click <Yes>, and Web-Based Admin View (top menu) appears.

Click <No>, and you will return to the main screen.

How to exit Web-Based Admin View is described below.

- Click <Log Out> from one of the following menus; top menu, GDS Management Operation menu, or Common Operation menu.

- Log In screen appears. Exit the browser, or click Back on the browser and exit GUI.



When the login screen remains displayed

After exiting the Web browser, the login screen may remain displayed but will disappear after a while.

3.4 Changing the Web-Based Admin View Settings

When changing one of the following settings after installing Web-Based Admin View, see "Web-Based Admin View Operation Guide" for details.

- Modifying the IP address of the public LAN
- Modifying the port number of the network service
- Changing the management server
- Modifying the operation of the secondary management server
- Modify the network environment in the management server

Chapter 4 Management View Screen Elements

This chapter explains the screen elements of GDS Management View.

Screen elements refer to the following items:

- screen configuration
- menu configuration and functions
- icon types and object status
- object information

4.1 Screen Configuration

Main Screen

Click <Global Disk Services> from Web-Based Admin View, and the screen below appears.

From the main screen, you can perform GDS object configuration, such as a class and a volume, status confirmation, and disk swap.

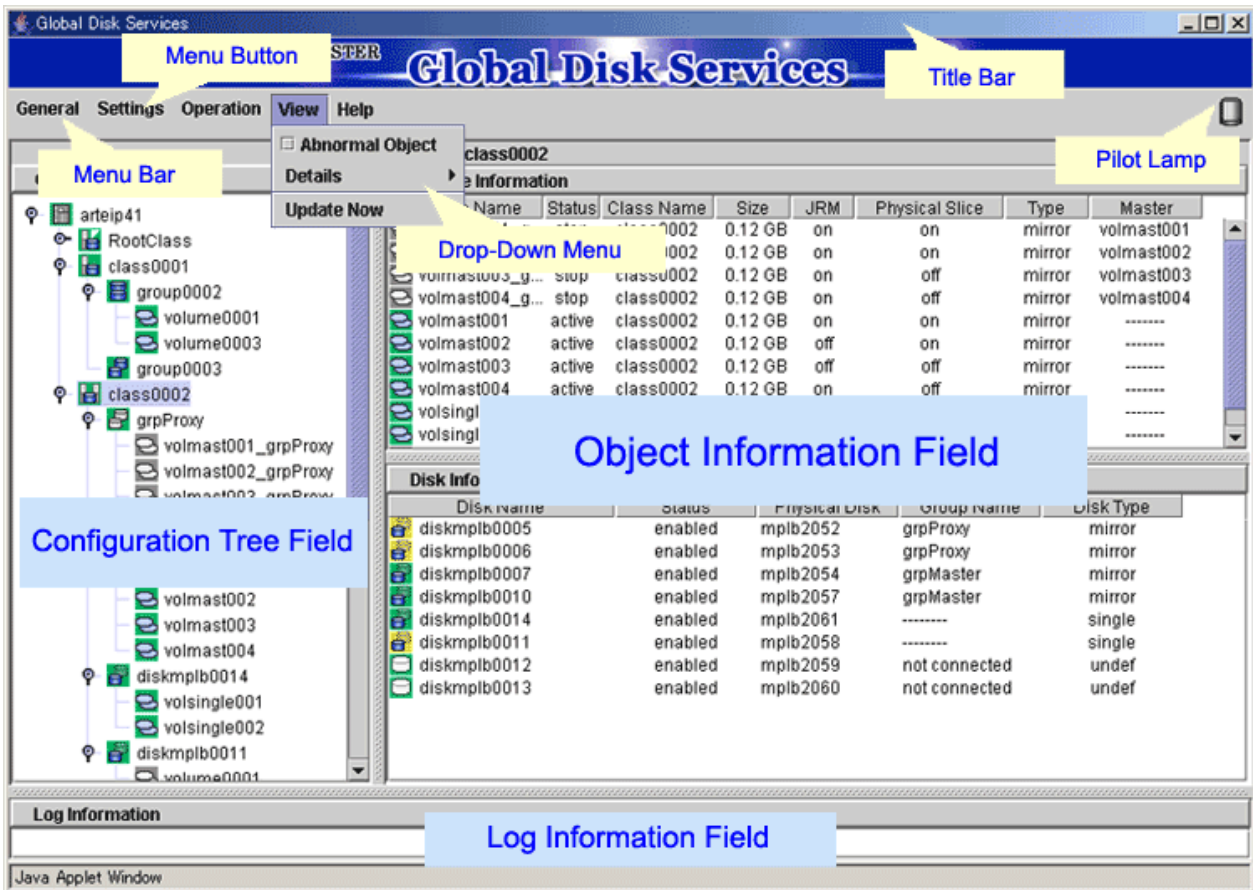
You can also configure and view the statuses of GDS Snapshot proxy objects and shadow objects and also operate those proxy objects.

The following operations are unsupported. Use commands for these operations.

- Object operations in classes that include switch groups
(Information of relevant objects is displayed in blue.)
- Operations of GDS Snapshot shadow objects
(Information of relevant objects is displayed in italic format.)

The screen configuration of the main screen is shown below.

Figure 4.1 GDS Management Screen (Main Screen)



Configuration Tree Field

Objects managed with GDS are displayed in a tree-structured directory system.

Each object has an icon depicting the type and the status of the object.

For details on the icon type and status, see "4.3 Icon Types and Object Status."

By selecting the node in the GDS configuration tree field, you can switch between the nodes you want to display or operate.

Object Information Field

Detailed information of objects is displayed in table format.

The displayed contents vary according to the menus selected in [View]:[Details] and the object types selected on the GDS configuration tree.

For details on displayed contents, see "4.4 Object Information."

Log Information Field

Displays error message output by the GDS daemon program.

Title Bar

Displays screen title (Global Disk Services).

Menu Bar

Displays the menu buttons.

Menu Button

Allows you to control the objects selected on screen.

There are [General], [Settings], [Operation], [View] and [Help].

Drop-Down Menu

When a menu button from the menu bar is selected, a drop-down menu will be displayed.

For details on drop-down menu, see "[4.2 Menu Configuration and Functions](#)."

Popup Menu






An on-screen menu that is displayed by a right-click on an object.

[Check Status] in the popup menu displays a description of the state of the object and the help to restore the object if it is faulty.

Pilot Lamp

Shows the status of the monitored objects.

The lamp will indicate the following status.


Pilot Lamp	Status	Meaning
 (Gray lit up)	Normal	Normal status.
 (Red blinking)	Abnormal	Critical abnormality (e.g. closed class, invalid volume)
 (Red lit up)	Abnormal	When red blinking warning lamp is single-clicked.
 (Yellow blinking)	Degradation	- Volume operating at degradation. (e.g. invalid slice, disabled disk) - Disk I/O error
 (Yellow lit up)	Degradation	When yellow blinking warning lamp is single-clicked.

GDS Configuration Settings screen

Select Configuration from [Settings] menu, and the "GDS Configuration Settings screen" shown below will appear.

Use the <Screen Switching Tab> to switch between "Class Configuration", "Group Configuration", and "Volume Configuration" settings screens.

Figure 4.2 GDS Configuration Settings Screen



For information on how to set each configuration, see "5.2.3 Operating from the Settings Menu."

4.2 Menu Configuration and Functions

Each menu button has a drop-down menu allowing you to operate the selected object on screen.

This section explains the menu configuration and functions.

The operations for shadow objects available with GDS Snapshot are not supported.

4.2.1 General



Change Monitoring Intervals

Sets the monitoring interval (min) of objects.

Figure 4.3 [General]: [Change Monitoring Intervals] Screen



Note


When Changing Monitoring Intervals

The default value for monitoring intervals is 5 minutes. When GDS Management is restarted, the monitoring interval is initialized to 5 minutes. To change the monitoring interval, modify the monitoring interval value each time GDS Management is restarted.

Exit

Exits GDS Management.

Figure 4.4 [General]: [Exit] Screen



4.2.2 Settings

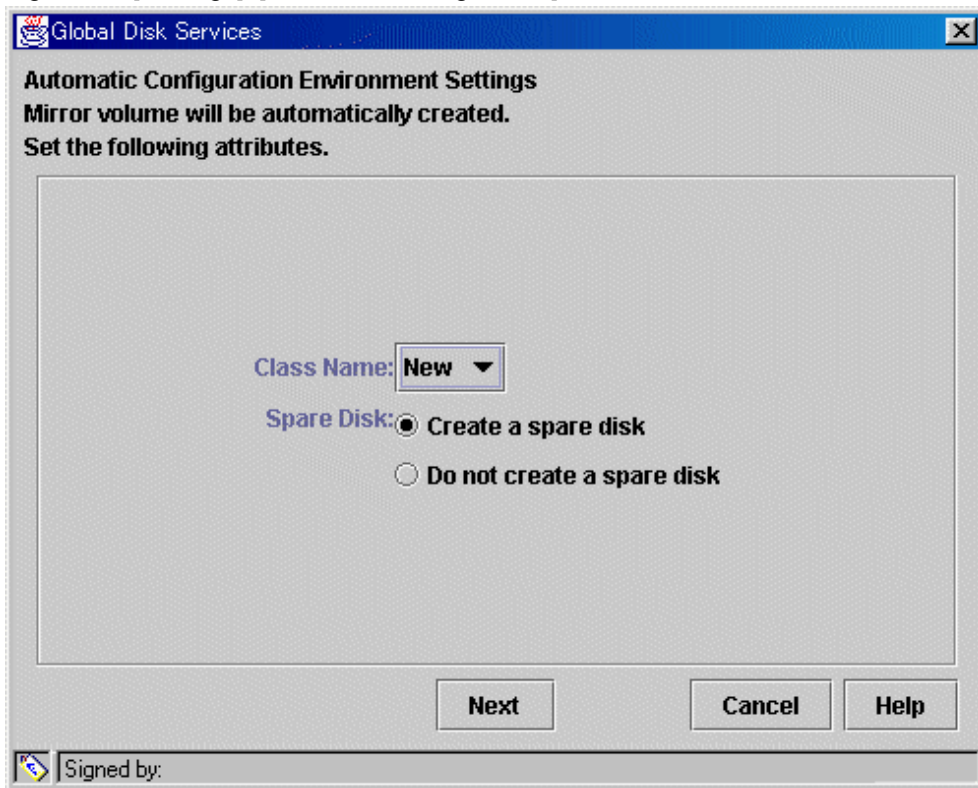


Automatic Configuration

Allows you to simply specify mirroring settings for a local disk by guiding you through the process with wizards.

For details on operation, see "[5.2.2 Automatic Configuration Wizard.](#)"

Figure 4.5 [Settings]: [Automatic Configuration] Screen

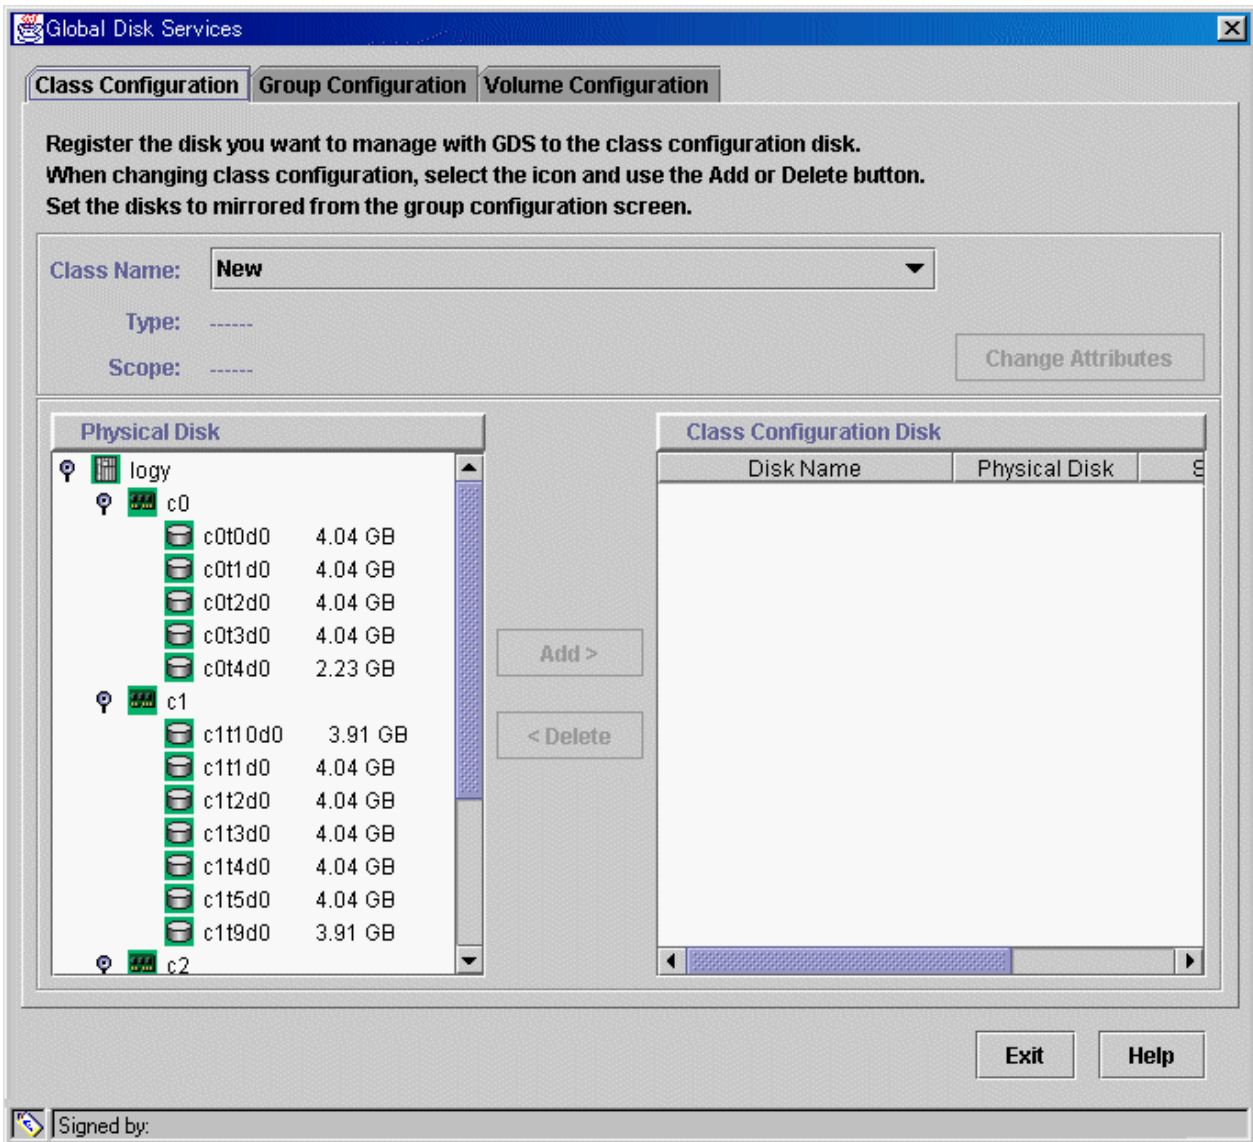


Class Configuration

Sets the class configuration.

For details, see "[5.2.3 Operating from the Settings Menu.](#)"

Figure 4.6 [Settings]: [Class Configuration] Screen

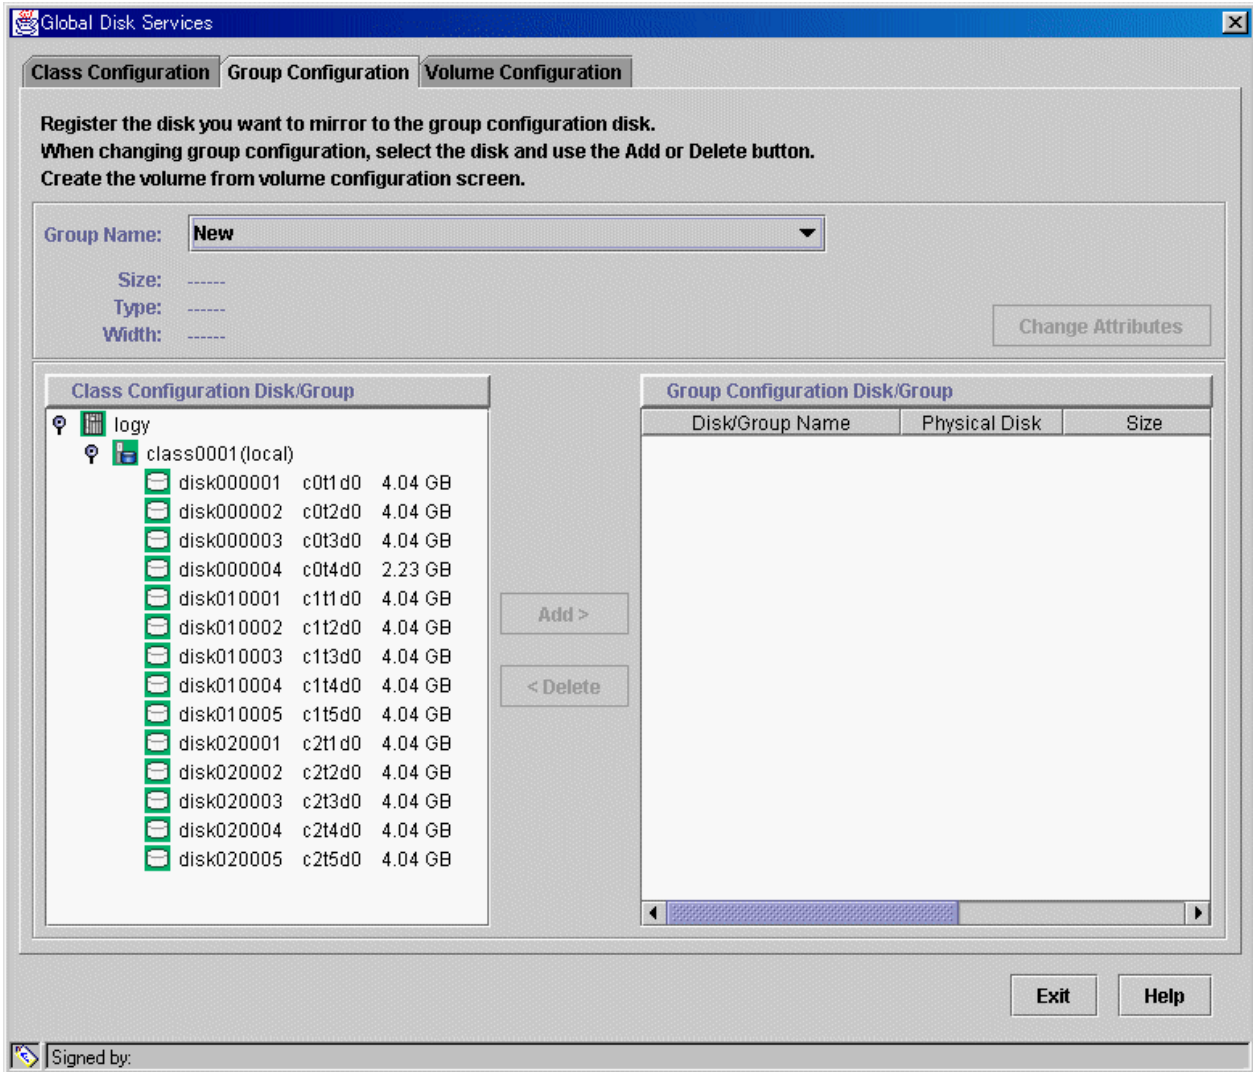


Group Configuration

Sets the group configuration.

For details, see "[5.2.3 Operating from the Settings Menu.](#)"

Figure 4.7 [Settings]: [Group Configuration] Screen




Volume Configuration

Sets the volume configuration.

For details, see "5.2.3 Operating from the Settings Menu."

Figure 4.8 [Settings]: [Volume Configuration] Screen

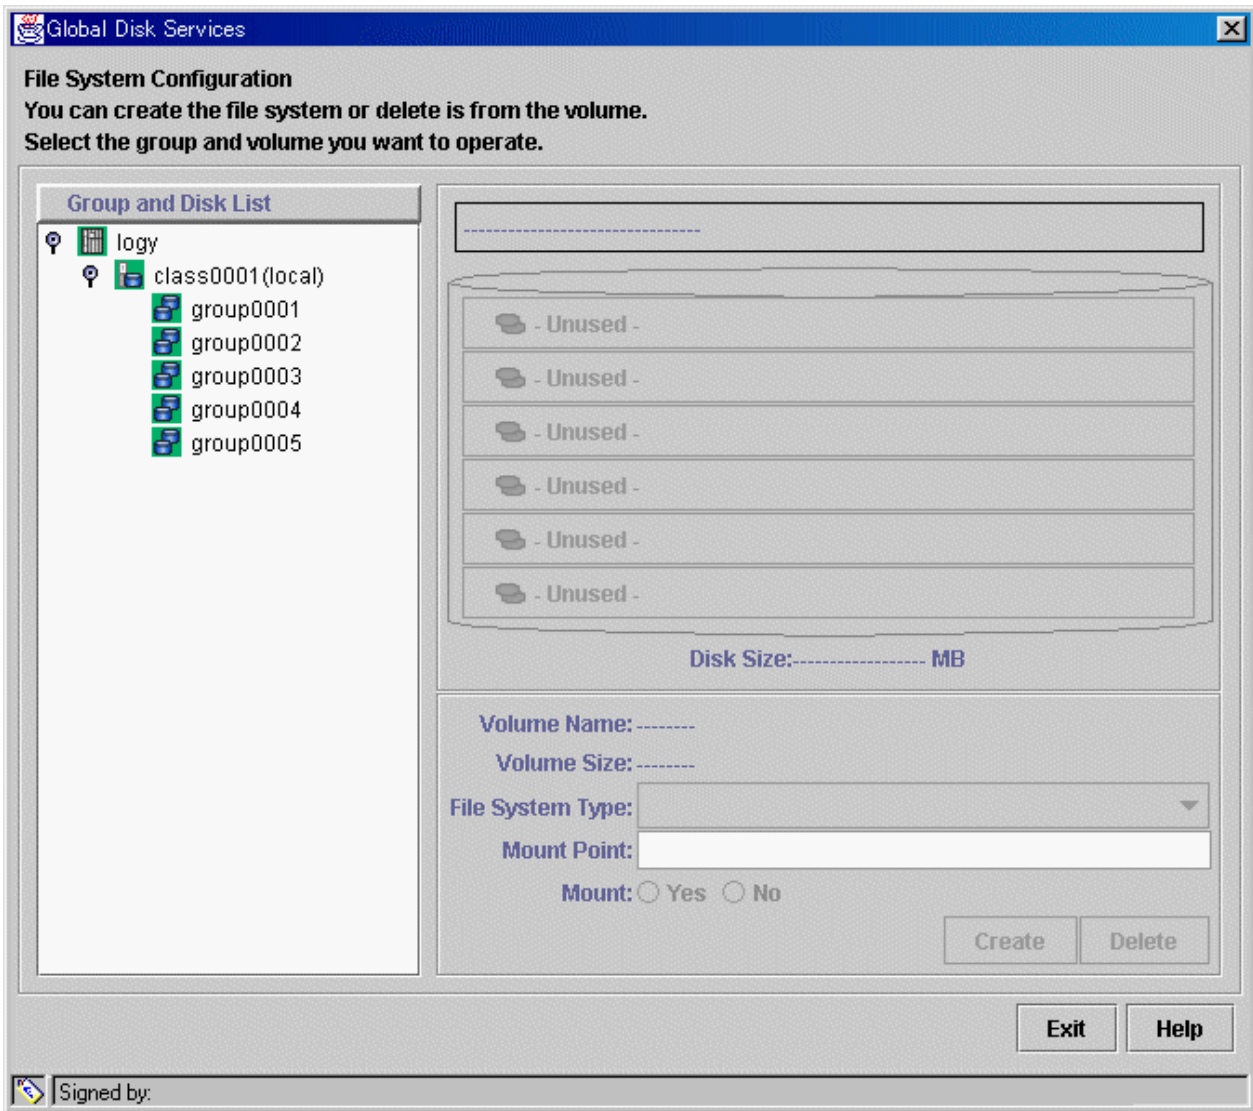


File System Configuration

Sets the file system configuration.

For details, see "[5.2.4 File System Configuration](#)."

Figure 4.9 [Settings]: [File System Configuration] Screen




System Disk Settings

Mirrors the system disk.

For details, see "[5.2.1 System Disk Settings](#)."

Figure 4.10 [Settings]: [System Disk Settings] Screen




Unmirror System Disk

Unmirrors the system disk.

For details, see "[5.5.5 Unmirroring the System Disk.](#)"

Figure 4.11 [Settings]: [Unmirror System Disk] Screen



4.2.3 Operation



Swap Physical Disk

Places the physical disk off-line for swapping physical disks.

For details, see "[5.3.4 Disk Swap](#)."

Restore Physical Disk

Places the swapped physical disk on-line for restoration after swapping physical disks.

For details, see "[5.3.4 Disk Swap](#)."

Detach Slice

Detaches one of the slices from Mirror volume to prepare for backup. The detached slice will become accessible as a separate logical device.

For details, see "[5.3.2.1 Backup \(by Slice Detachment\)](#)."

Attach Slice

Slice detached by [Detach Slice] is attached to Mirror volume again.

For details, see "[5.3.2.1 Backup \(by Slice Detachment\)](#)."

Stop/Activate Slice

Stop Slice

In order to protect data of the slice which has been detached to prepare for backup, a slice with "temp" status will temporarily become inaccessible.

Activate Slice

The detached slice which is now inaccessible ("temp-stop" status) as a result of [Stop Slice] operation or switching of nodes will be reactivated and become accessible.

For details, see "[5.3.2.1 Backup \(by Slice Detachment\)](#)."

Start Copying

A slice with "invalid" or "copy-stop" status as a result of [Cancel Copying] operation will be attached to a mirror volume, and synchronization copying will be performed.

For details, see "[5.3.6 Copying Operation](#)."

Cancel Copying

Execution of copying will be stopped to avoid effects caused by accessing the disk in the process of synchronization copying.

For details, see "[5.3.6 Copying Operation](#)."

Start Volume

Starts the stopped volume.

Stop Volume

Stops the volume.

Proxy Operation

Operates proxy objects. This menu is available only if GDS Snapshot is installed.

Join

Relates proxy objects to master objects and synchronizes them as preparation for snapshot creation by synchronization, or online disk migration.

For details, see "[5.2.5.1 Join](#)."

Part

Parts joined proxies from masters temporarily to make them accessible as logical devices that are different from the masters for snapshot creation by synchronization. The parted proxies can be used as snapshots (replicas) of the masters at the moment.

For details, see "[5.3.2.2 Backup \(by Synchronization\)](#)."

Rejoin

Joins parted proxies to masters again and synchronizes them as preparation for snapshot re-creation by synchronization.

For details, see "[5.3.2.2 Backup \(by Synchronization\)](#)."

Relate

Relates proxy objects to master objects and parts them as preparation for snapshot creation by OPC.

For details, see "[5.2.5.2 Relate](#)."

Update

Copies (overwrites) data from masters to parted proxies for snapshot creation by OPC. The updated proxies can be used as snapshots (replicas) of the masters at the moment.

For details, see "[5.3.2.3 Backup \(by OPC\)](#)."

Restore

Copies (overwrites) data from parted proxies to masters for restoration of damaged master data. The masters are recovered by the proxy data at the moment.

For details, see "[5.3.3 Restore](#)."

Swap

Swaps slices of synchronized masters and proxies for online disk migration.

For details, see "[5.3.5 Disk Migration](#)."

Break

Breaks the relationships between masters and proxies and makes them unrelated again.

For details, see "[5.5.6 Breaking a Proxy](#)."

Change Attributes

Changes the attributes of the selected object.

For details, see "[5.4 Changes](#)."

Update Physical Disk Information

You can update the disk information without rebooting the system.

This feature is useful in the situations given below.

- When physical disk size is not displayed properly.
- When you turn on the disk array or the disk unit after booting the system.
- When the disk has become inaccessible while operating the system.
- When the configuration of devices such as mphpd or mplb has been changed.
- When automatic resource registration was performed.

Check Status

Displays a description of the state of an object and the help to restore the object if it is faulty.

4.2.4 View



Abnormal Object

Only displays objects with abnormalities; not all objects.

Details

Changes displayed contents in the Object Information Field.

By default, [SDX Object] is selected.

SDX Object

Displays information of volumes, disks and slices.

Proxy Object

Displays information of proxy volumes, proxy groups and slices.

For details on the Objects Information Field, see "[4.4 Object Information.](#)"

Update Now

Usually, GDS Management screen updates the information on object status at an interval specified by [Change Monitoring Intervals] in the [General] menu.

If you select [Update Now], object status will be updated immediately regardless to the interval specified by [Change Monitoring Intervals].

In order to recognize the disk again, select [Update Physical Disk Information] from [Operation] menu.

4.2.5 Help



Help

Displays Help information.

4.3 Icon Types and Object Status

GDS Management uses icons to show the status of the objects.






Information


The SDX objects that belong to GDS classes and the shadow objects that belong to GDS Snapshot shadow classes are distinguished by fonts. Information related to shadow objects is displayed in italics.

The status and the icons of objects are shown below.




1. Nodes

Icon	Status	Meaning
 (Green)	Normal	-
 (Yellow)	Abnormal	Abnormality within the node
 (Red)	Abnormal	Abnormality of the node




2. Adapter

Icon	Status	Meaning
 (Green)	-	-



3. Classes (local)

Icon	Status	Meaning
 (Green)	Normal	-
 (Yellow)	Abnormal	Abnormality within the (local) class
 (Red)	Closed	Closed class within the class




4. Classes (shared)

Icon	Status	Meaning
 (Green)	Normal	-
 (Yellow)	Abnormal	Abnormality within the (shared) class
 (Red)	Closed	Closed class within the class




5. Classes (root)

Icon	Status	Meaning
 (Green)	Normal	-
 (Yellow)	Abnormal	Abnormality within the (root) class




6. Groups (mirror)

Icon	Status	Meaning
 (Green)	Normal	-
 (Yellow)	Abnormal	Abnormality within the (mirror) group
 (Red)	Abnormal	Closure of the class to which the (mirror) group belongs




7. Groups (stripe)

Icon	Status	Meaning
 (Green)	Normal	-
 (Yellow)	Abnormal	Abnormality within the (stripe) group
 (Red)	Abnormal	Closure of the class to which the (stripe) group belongs




8. Groups (concat)

Icon	Status	Meaning
 (Green)	Normal	-
 (Yellow)	Abnormal	Abnormality within the (concat) group
 (Red)	Abnormal	Closure of the class to which the (concat) group belongs







9. Groups (switch)

Icon	Status	Meaning
 (Green)	Normal	-
 (Yellow)	Abnormal	Abnormality within the (switch) group
 (Red)	Abnormal	Closure of the class to which the (switch) group belongs






10. Physical disks

Icon	Status	Meaning
 (Green)	Normal	-
 (Yellow)	Abnormal	Object status of some volumes or slices not active
 (Red)	Abnormal	Abnormality within the physical disk






11. Disks connected to a group


Icon	Status	Meaning
 (Green)	enabled	Operation enabled
 (Yellow)	enabled	Some volumes or slices with status other than active
 (Red)	enabled	I/O error on the disk
 (Red)	disabled	Operation disabled
 (Red)	close	Closure of the class to which the disk belongs
 (Light Brown)	swap	Disk swap possible

12. Single disks






Icon	Status	Meaning
 (Green)	enabled	Operation enabled
 (Yellow)	enabled	Some volumes or slices with status other than active
 (Red)	enabled	I/O error on the disk
 (Red)	disabled	Operation disabled
 (Red)	close	Closure of the class to which the single disk belongs

13. Spare disks







Icon	Status	Meaning
 (Green)	enabled	Operation enabled
 (Yellow)	enabled	Some volumes or slices with status other than active
 (Red)	enabled	I/O error on the disk
 (Red)	disabled	Operation disabled
 (Red)	close	Closure of the class to which the spare disk belongs

Icon	Status	Meaning
 (Light Brown)	swap	Disk swap possible










14. Unused disks

Icon	Status	Meaning
 (Green)	enabled	Operation enabled
 (Red)	enabled	I/O error on the disk
 (Red)	disabled	Operation disabled
 (Red)	close	Closure of the class to which the unused disk belongs
 (Light Brown)	swap	Disk swap possible



15. Volumes





Icon	Status	Meaning
 (Green)	active	Active volume
 (Yellow)	copy	Copying
 (Yellow)	warning	Object status of some slices not active (except slices in copying process)
 (Black)	stop	Volume stopped
 (Red)	invalid	Abnormality of volume
 (Red)	close	Closure of the class to which the volume belongs

16. Slices




Icon	Status	Meaning
 (Green)	active	Active slice
 (Black)	stop	Slice stopped
 (Red)	invalid	Abnormality of slice (Invalid data)
 (Red)	close	Closure of the class to which the slice belongs
 (Blue)	copy	Copying
 (Blue)	copy-stop	Synchronization copying temporarily interrupted
 (Light Blue)	temp	Slice temporarily excluded (accessible)
 (Light Blue)	temp-stop	Slice temporarily excluded (inaccessible)
 (Light Brown)	nouse	Operation disabled

17. Proxy volumes

Icon	Status	Meaning
 (Green)	active	Active proxy volume
 (Yellow)	copy	Copying

Icon	Status	Meaning
 (Yellow)	warning	Object status of some slices not active (except slices in copying process)
 (Black)	stop	Proxy volume stopped
 (Red)	invalid	Abnormality of proxy volume
 (Red)	close	Closure of the class to which the proxy volume belongs

18. Proxy groups

Icon	Status	Meaning
 (Green)	Normal	-
 (Yellow)	Abnormal	Abnormality within the group
 (Red)	Abnormal	Closure of the class to which the group belongs



Note

Abnormality Detected with GDS Management

What GDS Management displays as abnormal is limited to objects detected by GDS.

Therefore, even if the disk unit may have a hardware abnormality, it will be displayed as normal until it is accessed, and then abnormality is detected.

You should pinpoint the hardware error on a disk unit based on, for example, log messages for the disk driver output in the `/var/adm/messages` file. For details, see "[F.1.11 Disk Unit Error.](#)"

4.4 Object Information

This section describes information displayed in the Main Screen's Object Information Field.

The displayed contents in the Object Information Field vary according to the menus selected in [View]:[Details] and the object types selected on the GDS configuration tree.

By default, [View]:[Details]:[SDX Object] is selected.

- [View]:[Details]:[SDX Object] shows:

GDS Configuration Tree Object	Object Information Field (Upper)	Object Information Field (Lower)
Volume	Slice information	Disk information
Non-volume	Volume information	Disk information

- [View]:[Details]:[Proxy Object] shows:

GDS Configuration Tree Object	Object Information Field (Upper)	Object Information Field (Lower)
Volume	Proxy volume information	Slice information
Non-volume	Proxy group information	Proxy volume information

The Object Information Field displays the following information.

Field Name	Displayed Information
Volume Information	Volume name, status, assigned class name, size, JRM mode, physical slice attribute, type, master volume name
Disk Information	Disk name, status, physical disk name, assigned group name, type
Slice Information	Slice name, status, copy status, master-proxy GDS Snapshot copy type
Proxy Volume Information	Proxy volume name, master volume name, status, snapshot creation time (last parted or updated time), JRM (for rejoin and restore) mode
Proxy Group Information	Proxy group name, master group name

Chapter 5 Operation

This chapter describes the GDS operations from GDS Management View.

From GDS Management View, you can configure, operate, reconfigure, delete, view the configuration of, and monitor the status of GDS objects (SDX objects). If GDS Snapshot is installed, you can also operate proxy objects and view their configurations and statuses. For shadow objects, configuration viewing and status monitoring only are possible.

For details on operations supported in GDS Management View, see "[5.1.10 Operations from GDS Management View](#)."

5.1 Operation Outline

This section explains the GDS settings and operation management.



Note

In PRIMECLUSTER systems

Before defining the configuration of GDS objects such as classes and volumes, perform the following procedures.

1. Automatic Resource Registration

Register disk unit resources with the PRIMECLUSTER resource database. The resource registration is required even when only the system disk and local disks are in the management scope but shared disk units are not. For details on resource registration, see the "PRIMECLUSTER Cluster Foundation Configuration and Administration Guide."

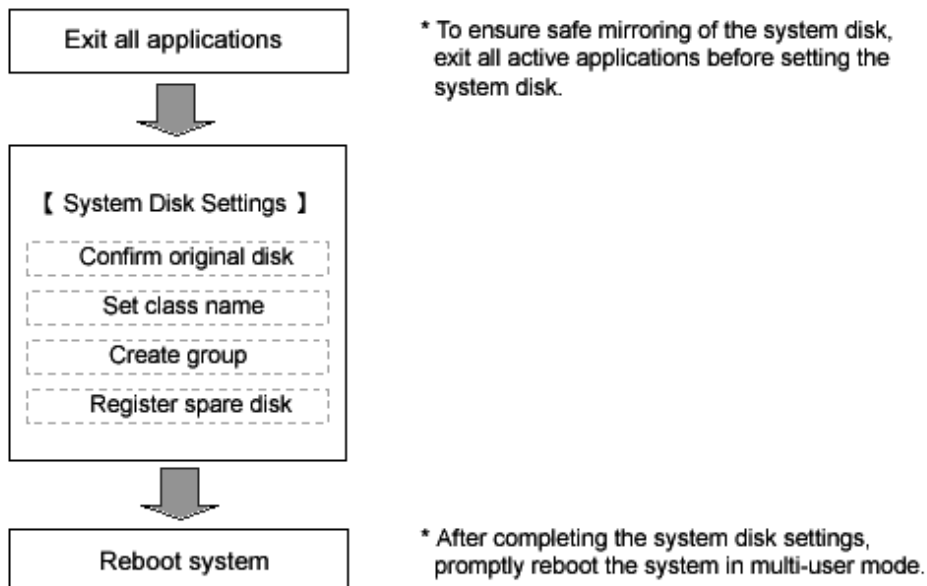
2. Physical Disk Information Update

From the [Operation] menu in GDS Management View, select [Update Physical Disk Information].

5.1.1 System Disk Settings

The operation outline of the setting of the system disk mirroring is shown below.

Figure 5.1 System Disk Settings Operation





See

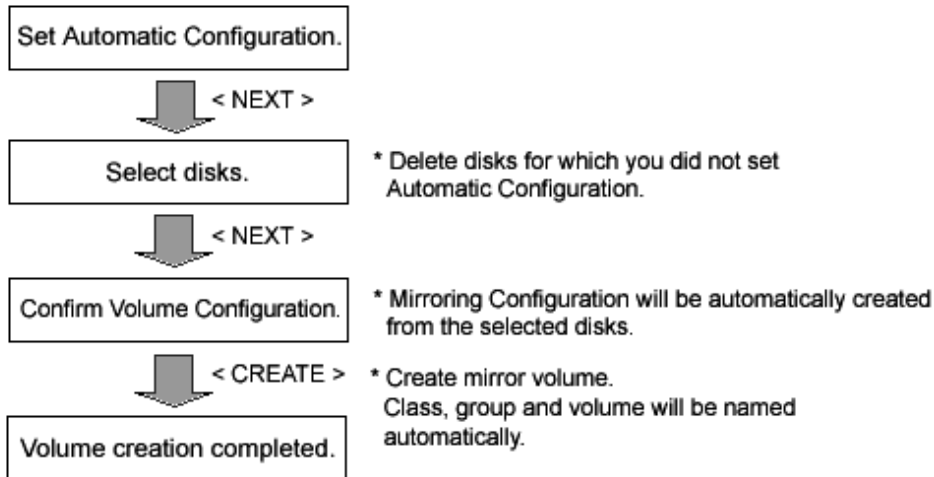
For details on the operation methods, see "5.2.1 System Disk Settings."

5.1.2 Automatic Configuration

The Automatic Configuration Wizard allows you to automatically set the basics by simply clicking <Next>.

The operation outline of the Automatic Configuration Wizard is shown below.

Figure 5.2 Automatic Configuration Wizard Operation



See

For details on the operation methods, see "5.2.2 Automatic Configuration Wizard."

5.1.3 Configuration Settings

If you are not using [System Disk Settings] or the Automatic Configuration Wizard, you must set the configurations using the relevant functions for classes, groups, and volumes.

The setting procedures differ depending on the type of volume you are creating.

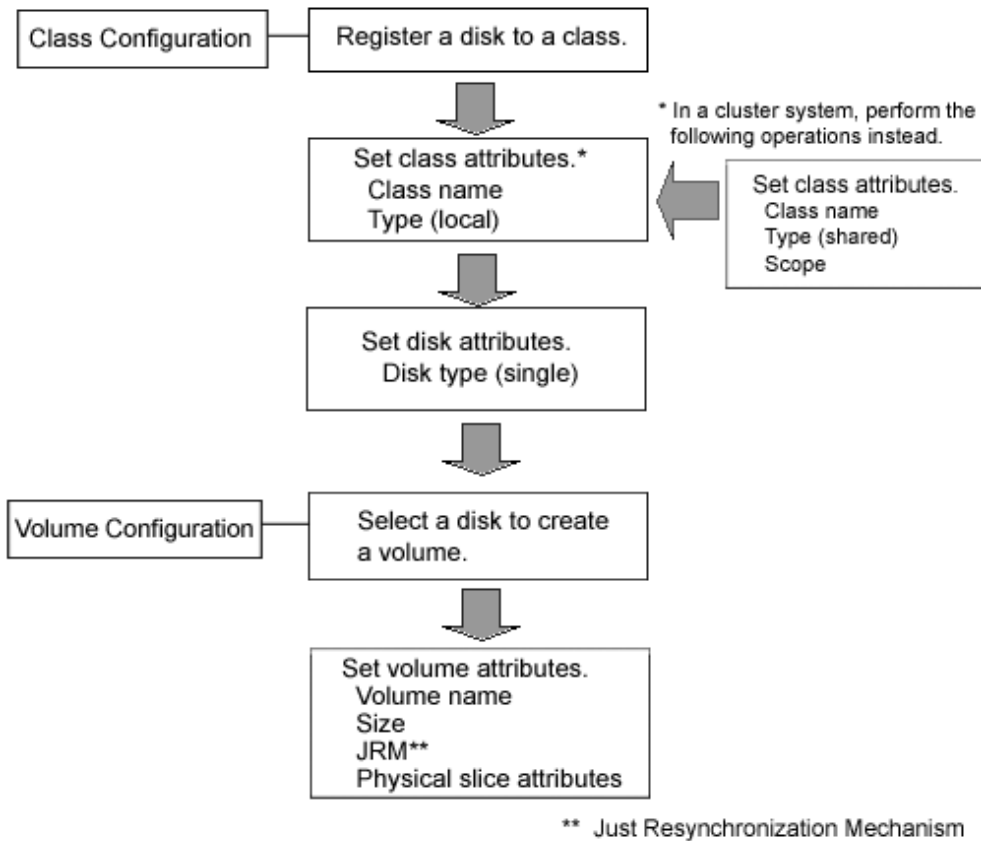
- Single volume configuration settings
- Other volume (mirror volume, stripe volume, volume in a concatenation group) configuration settings

Switch volume creation from GDS Management View is unsupported. For creating those volumes, use commands.

5.1.3.1 Single Volume Configuration Settings

The operation outline of single volume configuration settings is shown below.

Figure 5.3 Single Volume Configuration Settings Operation



 See

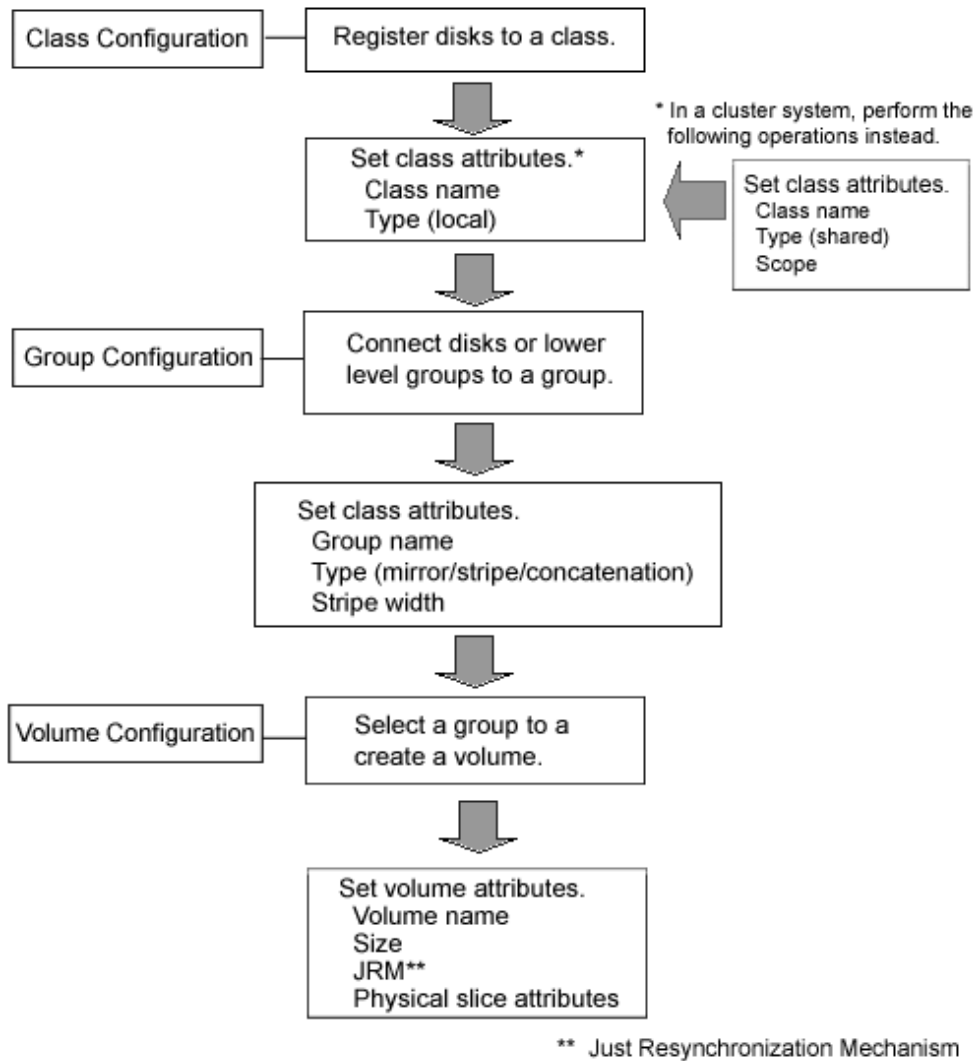
For details on the operation methods, see "5.2.3 Operating from the Settings Menu."

5.1.3.2 Other Volume Configuration Settings

The operation outline of the configuration settings for volumes other than single volume (mirror volume, stripe volume, volume in a concatenation group) is shown below.

Switch volume creation from GDS Management View is unsupported. For creating those volumes, use commands.

Figure 5.4 Volume Configuration Settings Operation



See

For details on the operation methods, see "5.2.3 Operating from the Settings Menu."

5.1.4 Backup

Volumes can be backed up while applications using them are running.

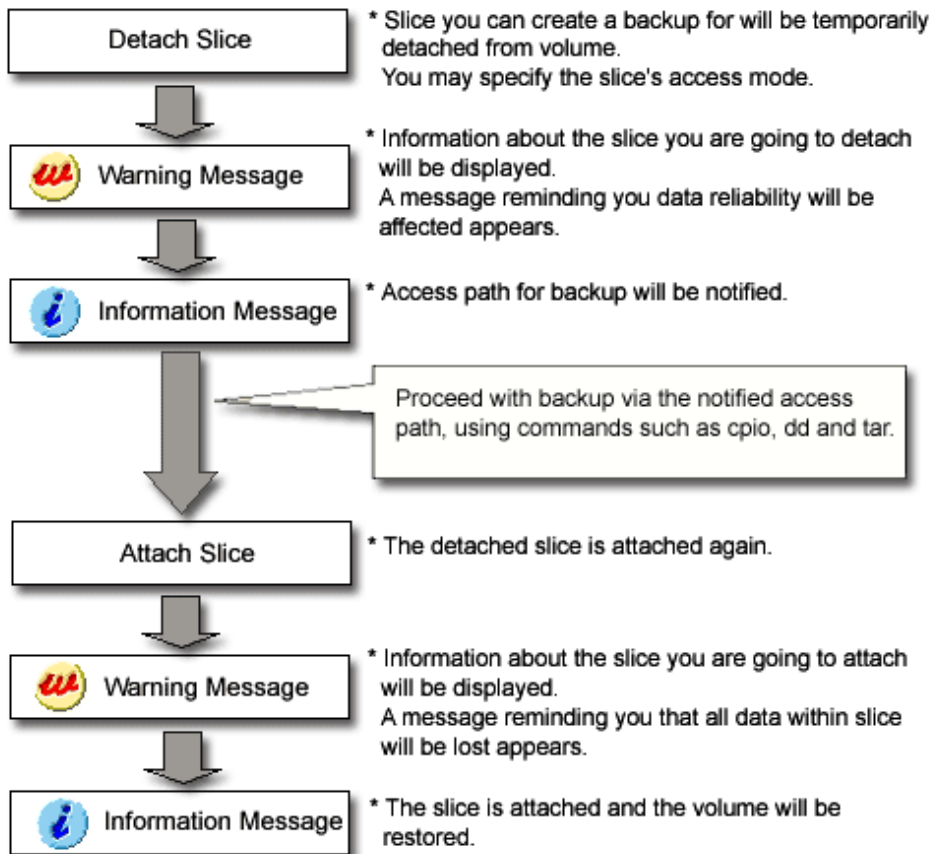
The following three backup methods are available.

- Slice detachment
- Synchronization (GDS Snapshot required)
- OPC (GDS Snapshot required)

5.1.4.1 Backup (by Slice Detachment)

The following figure shows the procedures for backing up mirror volumes through snapshot by slice detachment.

Figure 5.5 Backup Operation (by Slice Detachment)



Note

Use Conditions on Snapshot by Slice Detachment

Slices can be detached only from mirror volumes with physical slices. In other words, if disks are not connected directly to mirror groups, creation of snapshots by slice detachment is impossible.

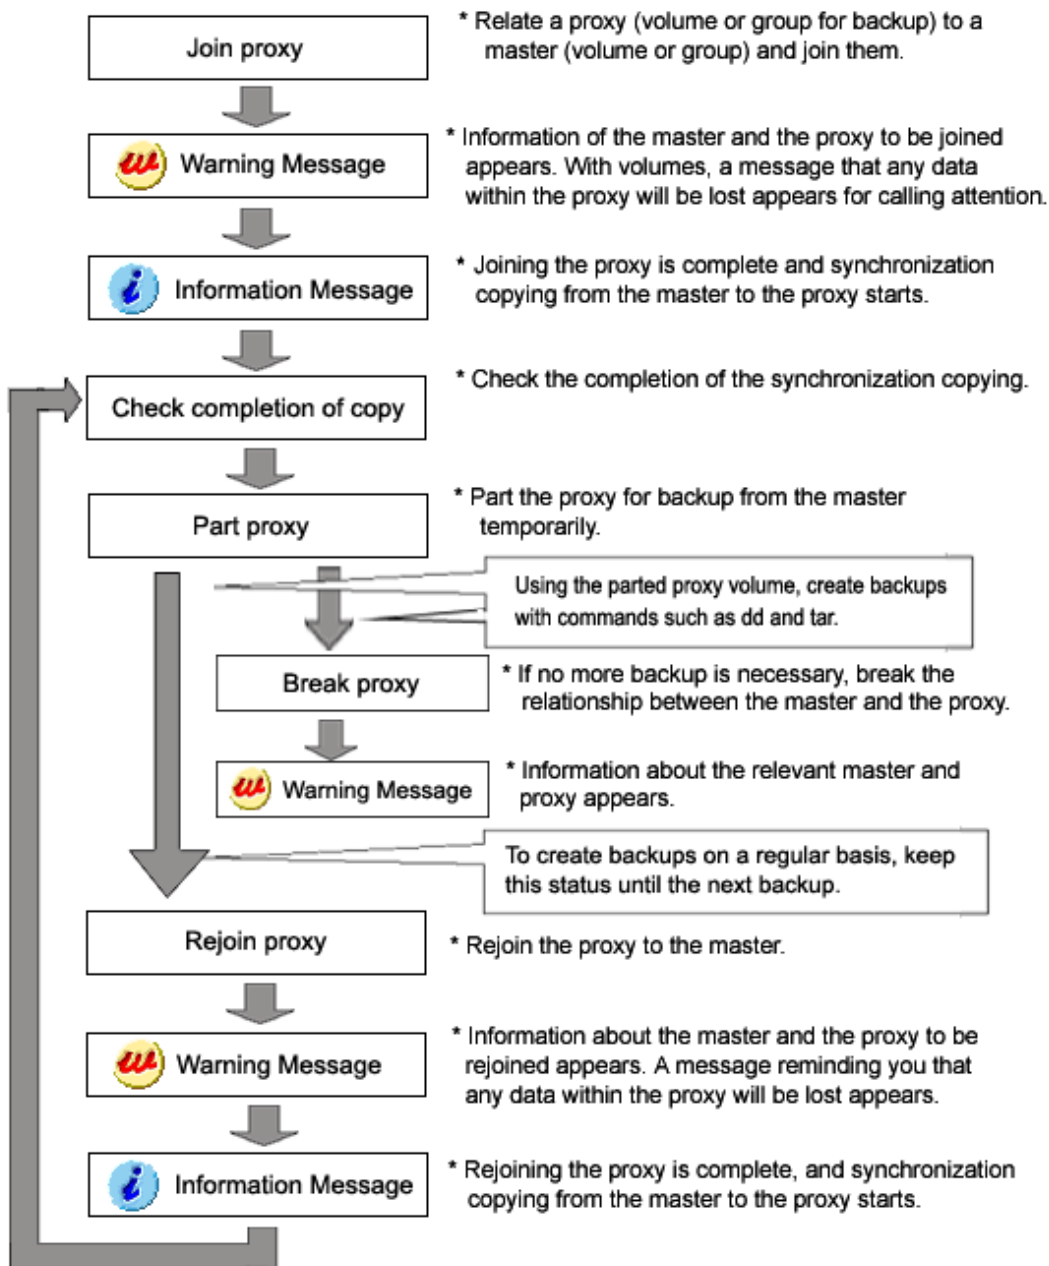
See

- For snapshot by slice detachment, see "[1.3.8 Snapshots by Slice Detachment](#)."
- For details on the operation methods, see "[5.3.2.1 Backup \(by Slice Detachment\)](#)."

5.1.4.2 Backup (by Synchronization)

The following figure shows the procedures for backing up volumes through use of snapshots (proxy volumes) of GDS Snapshot created by synchronization.

Figure 5.6 Backup Operation (by Synchronization)



Note

Snapshot by Synchronization Limitation

Snapshots by synchronization can be created only if the master and proxy type is mirror or single. For details, see "[A.1.8 Proxy Configuration Preconditions](#)."

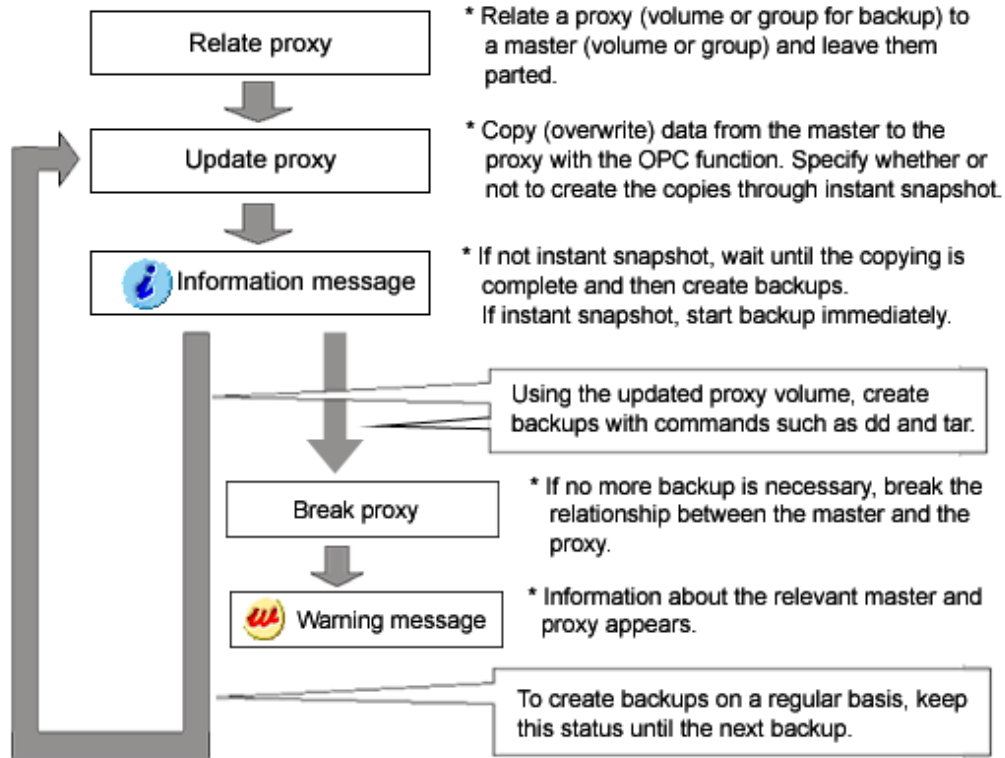
See

- For snapshot by synchronization, see "[1.5.1 Snapshot by Synchronization](#)."
- For details on the operation methods, see "[5.3.2.2 Backup \(by Synchronization\)](#)."

5.1.4.3 Backup (by OPC)

The following figure shows the procedures for backing up volumes through use of snapshots (proxy volumes) of GDS Snapshot created by OPC.

Figure 5.7 Backup Operation (by OPC)



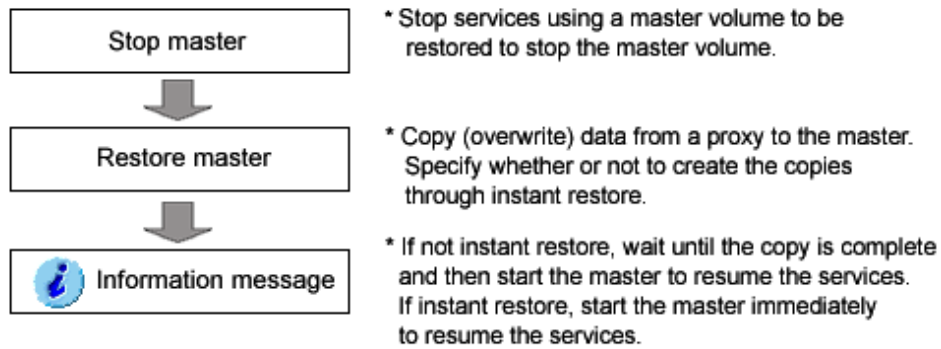
See

- For snapshot by OPC, see "1.5.3 Instant Snapshot by OPC."
- For details on the operation methods, see "5.3.2.3 Backup (by OPC)."

5.1.5 Restore

The following figure shows the procedures for restoring volumes through use of GDS Snapshot proxy volumes.

Figure 5.8 Restore Operation



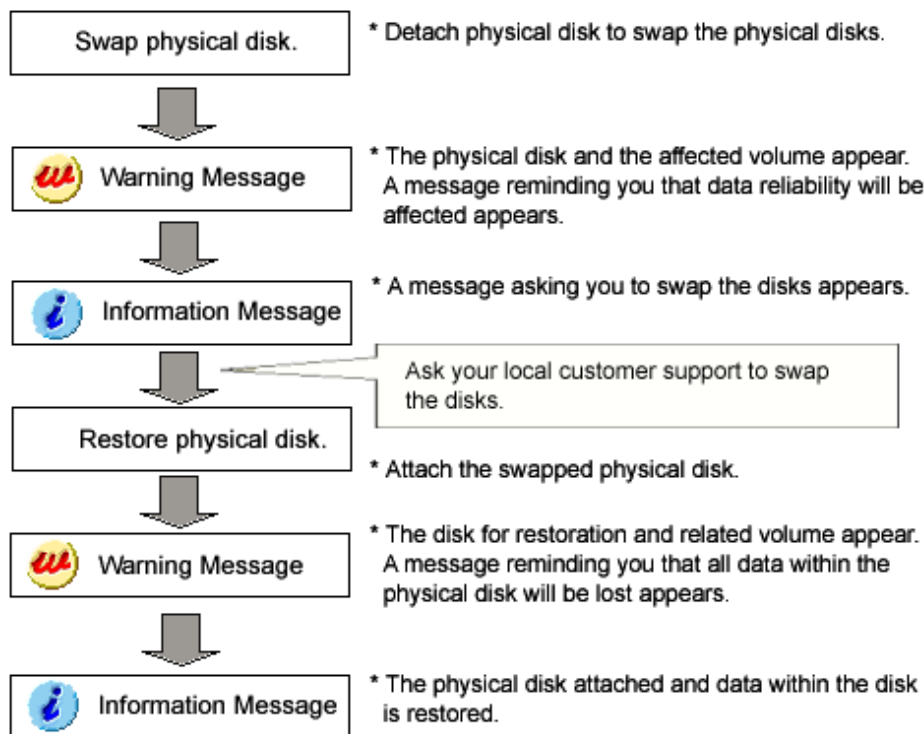
See

- For restore with proxy volumes, "1.5.4 Instant Restore."
- For details on the operation methods, see "5.3.3 Restore."

5.1.6 Disk Swap

The following operations are required for swapping disks in the event of a disk error or for the purpose of preventive maintenance.

Figure 5.9 Disk Swap and Restoration Operation



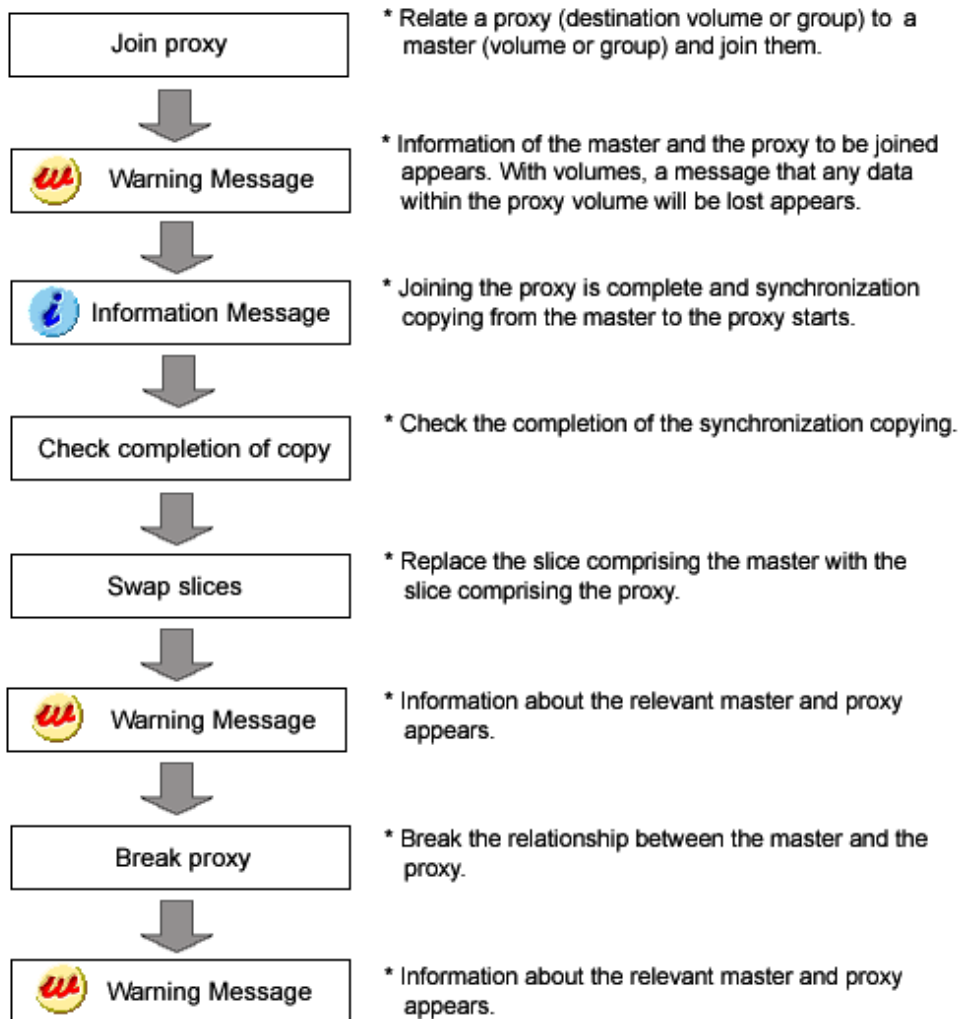
See

For details on the operation methods, see "5.3.4 Disk Swap."

5.1.7 Disk Migration

The following figure shows the procedures for migrating disks through use of GDS Snapshot proxy volumes.

Figure 5.10 Disk Migration Operation



Note

Disk Migration Precondition

Disk migration is supported only if the master and proxy type is mirror or single.

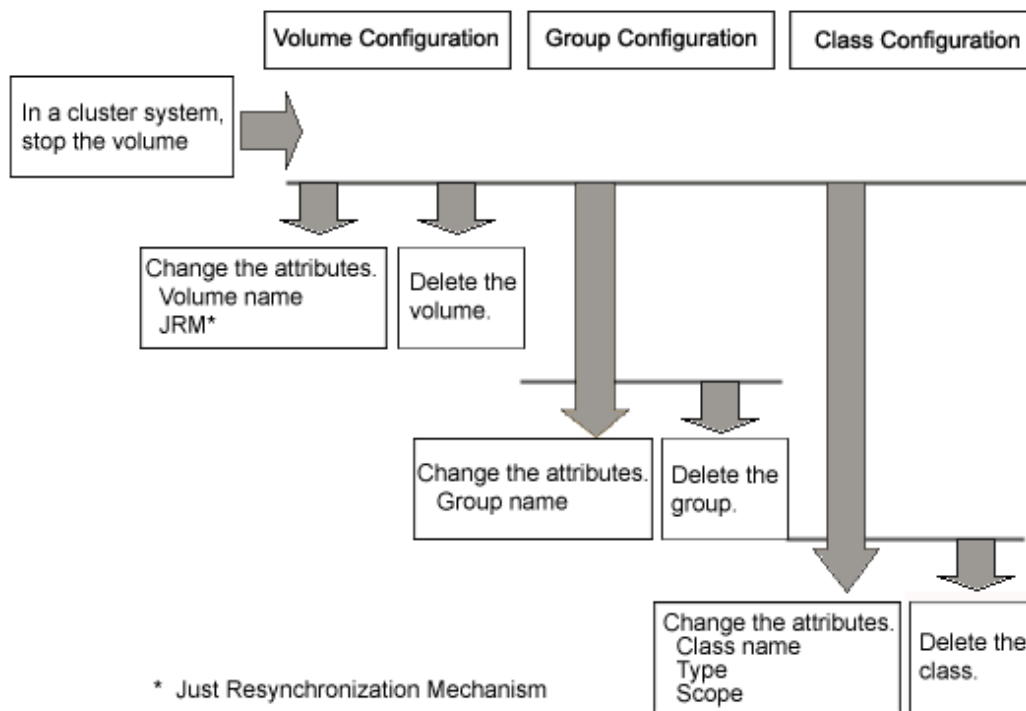
See

- For disk migration with proxy volumes, see "[1.5.5 Online Disk Migration](#)."
- For details on the operation methods, see "[5.3.5 Disk Migration](#)."

5.1.8 Configuration Change

The procedure for changing or removing the configuration settings is shown below.

Figure 5.11 Configuration Changing Operation



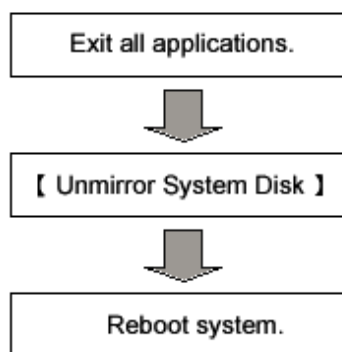
 See

For details on the operation methods, see "5.4 Changes" or "5.5 Removals."

5.1.9 Unmirroring the System Disk

The operation outline of unmirroring the system disk is shown below.

Figure 5.12 Unmirroring System Disk Operation



* To ensure safe unmirroring of the system disk, exit all active applications before unmirroring the system disk.
After unmirroring the system disk is complete, promptly reboot the system in multi-user mode.



See

For details on the operation methods, see "5.5.5 Unmirroring the System Disk."

5.1.10 Operations from GDS Management View

GDS Management View supports the following operations.

Use commands for operations that are unsupported in GDS Management View.

Class operation

Operation	Command Operation	GDS Management View Operation
Create	# sdxdisk -M	For details, see "5.2.3.1 Class Configuration" in "5.2.3 Operating from the Settings Menu."
Remove	# sdxdisk -R or # sdxclass -R	For details, see "5.5.4 Removing a Class" in "5.5 Removals"
Change attribute (Class name)	# sdxattr -C -a name= <i>classname</i>	For details, see "5.2.3.1 Class Configuration" in "5.4 Changes."
Change attribute (Type)	# sdxattr -C -a type={local shared}	For details, see "5.4.1 Class Configuration" in "5.4 Changes."
Change attribute (Scope)	# sdxattr -C -a scope= <i>node: node...</i>	For details, see "5.4.1 Class Configuration" in "5.4 Changes."
Change attribute (Hot spare)	# sdxattr -C -a hs={on off}	Unsupported
Change attribute (Hot spare mode)	# sdxattr -C -a hsmode={exbox bybox}	Unsupported
Restore	# sdxfix -V	Unsupported

Group operation

Operation	Command Operation	GDS Management View Operation
Connect	# sdxdisk -C or # sdxgroup -C	For details, see "5.2.3.3 Group Configuration" in "5.2.3 Operating from the Settings Menu."
Remove	# sdxdisk -D or # sdxgroup -D or # sdxgroup -R	For details, see "5.5.3 Removing a Group" in "5.5 Removals."
Change attribute (Group name)	# sdxattr -G -a name= <i>groupname</i>	For details, see "5.4.2 Group Configuration" in "5.4 Changes."
Change attribute (Active disk)	# sdxattr -G -a actdisk= <i>disk</i>	Unsupported

Volume operation

Operation	Command Operation	GDS Management View Operation
Create	# sdxvolume -M	For details, see "5.2.3.4 Volume Configuration" in "5.2.3 Operating from the Settings Menu."

Operation	Command Operation	GDS Management View Operation
Remove	# sdxvolume -R	For details, see " 5.5.2 Removing a Volume " in " 5.5 Removals ."
Start	# sdxvolume -N	For details, see "Start Volume" in " 4.2.3 Operation ."
Stop	# sdxvolume -F	For details, see "Stop Volume" in " 4.2.3 Operation ."
Expand	# sdxvolume -S	Unsupported
Change attribute (Volume name)	# sdxattr -V -a name= <i>volumename</i>	For details, see " 5.4.3 Volume Configuration " in " 5.4 Changes ."
Change attribute (JRM)	# sdxattr -V -a jrm={on off}	For details, see " 5.4.3 Volume Configuration " in " 5.4 Changes ."
Change attribute (Physical slice)	# sdxattr -V -a pslice={on off}	For details, see " 5.4.3 Volume Configuration " in " 5.4 Changes ."
Change attribute (Lock volume)	# sdxattr -V -a lock={on off}	Unsupported
Change attribute (Access mode)	# sdxattr -V ... -a mode={rw ro}	Unsupported
Change attribute (JRM for proxy)	# sdxattr -V -a pjrm= off	Unsupported
Start copying	# sdxcopy -B	For details, see " 5.3.6 Copying Operation " in " 5.3 Operation in Use ."
Cancel copying	# sdxcopy -C	For details, see " 5.3.6 Copying Operation " in " 5.3 Operation in Use ."
Restore	# sdxfix -V	Unsupported

Slice operation

Operation	Command Operation	GDS Management View Operation
Detach	# sdxslice -M	For details, see " 5.3.2.1 Backup (by Slice Detachment) " in " 5.3 Operation in Use ."
Attach	# sdxslice -R	For details, see " 5.3.2.1 Backup (by Slice Detachment) " in " 5.3 Operation in Use ."
Activate	# sdxslice -N	For details, see " 5.3.2.1 Backup (by Slice Detachment) " in " 5.3 Operation in Use ."
Stop	# sdxslice -F	For details, see " 5.3.2.1 Backup (by Slice Detachment) " in " 5.3 Operation in Use ."
Take over	# sdxslice -T	Unsupported
Change attribute	# sdxattr -S	Unsupported

Disk operation

Operation	Command Operation	GDS Management View Operation
Swap	# sdxswap -O	For details, see " 5.3.4 Disk Swap " in " 5.3 Operation in Use ."
Restore	# sdxswap -I	For details, see " 5.3.4 Disk Swap " in " 5.3 Operation in Use ."

Operation	Command Operation	GDS Management View Operation
Change attribute	# sdxattr -D	For details, see " 5.4.1 Class Configuration " in " 5.4 Changes ."
Clear errors	# sdxfix -D	Unsupported

Proxy operation

Operation	Command Operation	GDS Management View Operation
Join	# sdxproxy Join	For details, see " 5.2.5.1 Join " in " 5.2 Settings ."
Part	# sdxproxy Part	For details, see " 5.3.2.2 Backup (by Synchronization) " in " 5.3 Operation in Use ."
Rejoin	# sdxproxy Rejoin	For details, see " 5.3.2.2 Backup (by Synchronization) " in " 5.3 Operation in Use ."
Rejoin and restore	# sdxproxy RejoinRestore	For details, see " 5.3.3 Restore " in " 5.3 Operation in Use ."
Swap	# sdxproxy Swap	For details, see " 5.3.5 Disk Migration " in " 5.3 Operation in Use ."
Relate	# sdxpsoxy Relate	For details, see " 5.2.5.2 Relate " in " 5.2 Settings ."
Update	# sdxpsoxy Update	For details, see " 5.3.2.3 Backup (by OPC) " in " 5.3 Operation in Use ."
Restore	# sdxpsoxy Restore	For details, see " 5.3.3 Restore " in " 5.3 Operation in Use ."
Cancel hard copy session	# sdxpsoxy Cancel	Unsupported
Create alternative boot environment	# sdxproxy Root	Unsupported
Break	# sdxproxy Break	For details, see " 5.5.6 Breaking a Proxy " in " 5.5 Removals ."

Shadow operation

Operation	Command Operation	GDS Management View Operation
Operate shadow disk	# sdxshadowdisk	Unsupported
Operate shadow group	# sdxshadowgroup	Unsupported
Operate shadow volume	# sdxshadowvolume	Unsupported

Operation to display the configuration information

Operation	Command Operation	GDS Management View Operation
View	# sdxinfo	For details, see " 5.3.1 Viewing Configurations/Statuses and Monitoring Statuses " in " 5.3 Operation in Use ."

Other operations

Operation	Command Operation	GDS Management View Operation
System disk settings	# sdxdisk -M # sdxdisk -C # sdxroot -M <reboot>	For details, see " 5.2.1 System Disk Settings " in " 5.2 Settings ."
Unmirror system disk	# sdxdisk -D # sdxroot -R <reboot> # sdxvolume -F # sdxvolume -R # sdxdisk -D # sdxdisk -R	For details, see " 5.5.5 Unmirroring the System Disk " in " 5.5 Removals ."
Automatic configuration	# sdxdisk -M # sdxdisk -C # sdxvolume -M	For details, see " 5.2.2 Automatic Configuration Wizard " in " 5.2 Settings ."
Configuration parameter operations	# sdxparam	Unsupported
Object configuration operations	# sdxconfig	Unsupported

5.2 Settings

This section explains the GDS setting procedure.

5.2.1 System Disk Settings

System Disk Settings allows you to mirror system disks.

System disks cannot be mirrored under the following conditions.

Change the disk configuration before performing the system disk settings.

- The system disk's slice numbers are all in use, which means that the number of slices on the system disk has reached the upper limit (seven).
- The system disk does not have sufficient free space or a swap area.
- The system disk has a 2 GB or larger swap device in 32-bit Solaris.

For details on necessary free space (or a swap area), refer to "[A.2.6 Disk Size](#)."



Note

Disk that will be mirrored by system disk settings

For system disk settings from the GDS Management View, disks with /, /var, /usr, and a swap area are recognized as system disks.

Disks that are recognized as system disks will automatically be the original disks for mirroring. Among /, /var, /usr, and a swap area, disks with only swap areas can be excluded from mirroring, but disks with /, /var, and /usr are always to be mirrored.

Also, physical disks that are not recognized as system disks cannot be mirrored using the [System Disk Settings].

Note

For Safe System Disk Mirroring

To safely mirror the system disk, exit all active applications before proceeding to system disk settings.

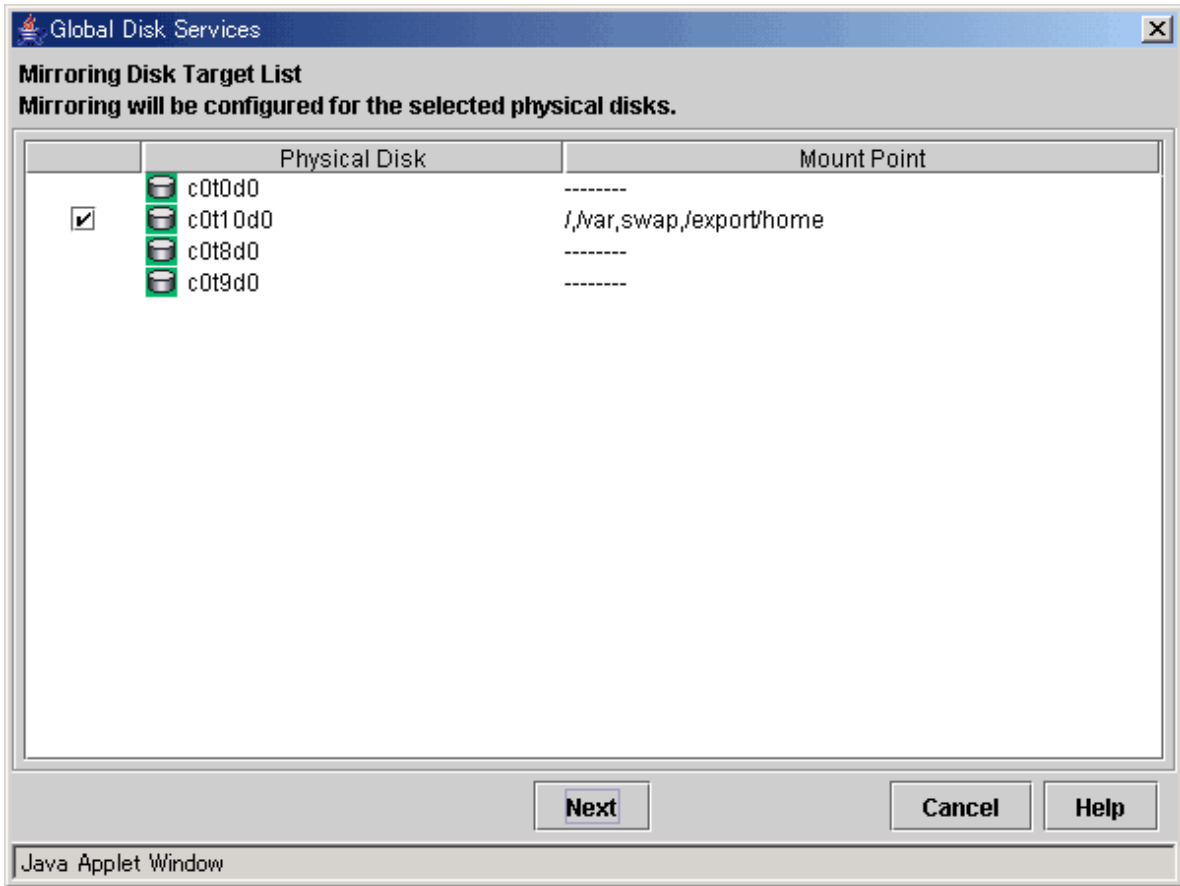
During the mirroring process, there may be a considerable degradation of application response time. After completing the system disk settings, promptly reboot the system.

How to mirror a system disk

In the [Settings] menu, select [System Disk Settings].

1. Confirming original disks

Figure 5.13 Mirroring Disk Target List



In the [Physical Disk] field, system disks that will be used as original disks will be displayed with a check mark.

For disks with mount information in /etc/vfstab, the [Mount Point] field displays their mount points.

Original disks for mirroring with check marks in the [Physical Disk] field are registered as keep disks with the root class and mirrored.

Uncheck system disks not to be registered with the root class.

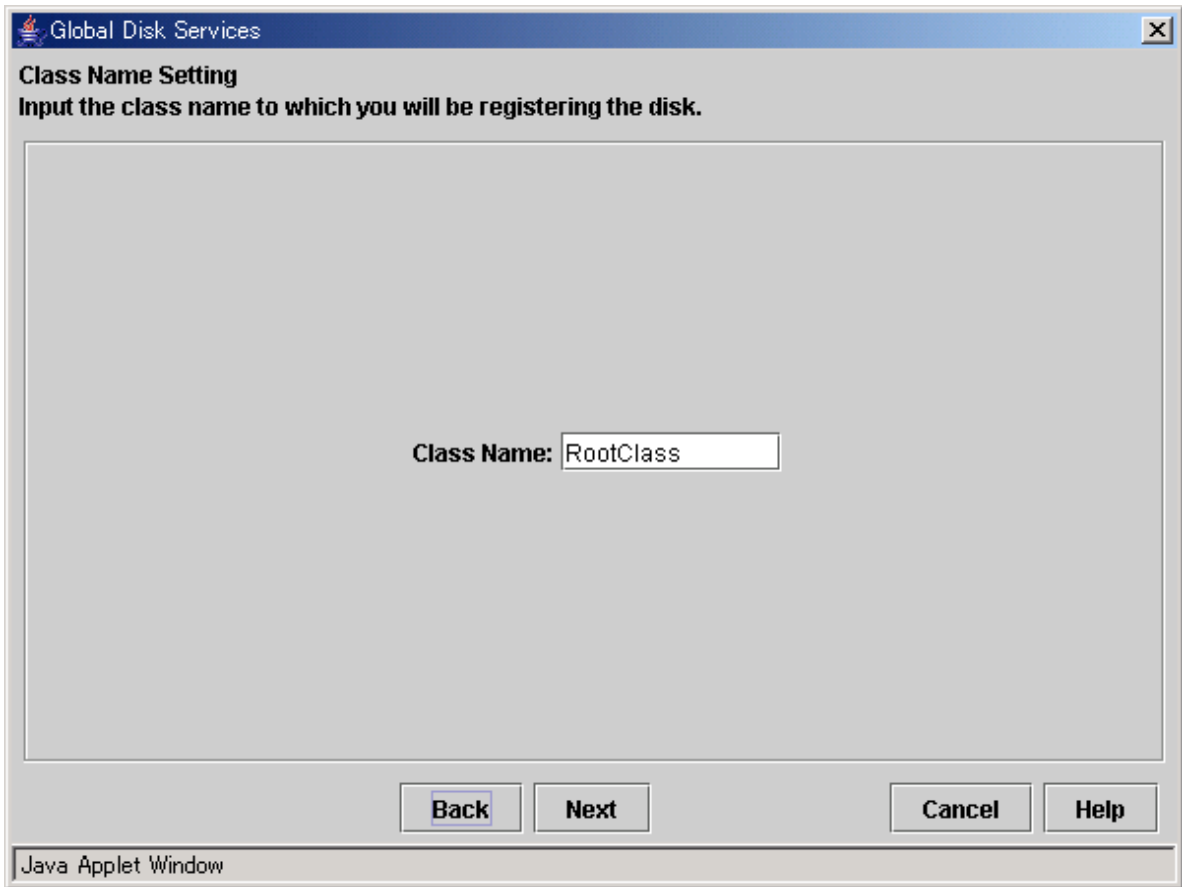
- It is possible to uncheck system disks for which only swap among /, /usr, /var, and swap is displayed in the [Mount Point] field.
- It is impossible to uncheck system disks for which /, /var, and /usr are displayed in the [Mount Point] field.

If you do not want to proceed with mirroring, click <Cancel>.

If you want to proceed with mirroring, click <Next>.

2. Creating root class

Figure 5.14 Class Name Setting



Type the class name of the root class.

If the root class is present, the class name cannot be set.

Information

Inputting the Class Name

The class name will be used for the device path name.

/dev/sfdsk/class_name/rdisk/volume_name

/dev/sfdsk/class_name/dsk/volume_name

You must be careful when inputting the class name, as once the volume is created, it cannot be changed.

Note

When Setting the System Disk in Cluster System

When setting the system disk in cluster systems, the class name of the root class should be different for each node.

See

For information on assigning a class name, see "A.1.1 Object Name."


Click <Next> to continue.

3. Creating group

Create the group by selecting the mirror disks.

When multiple system disks that will be used as original disks exist, perform group creation for every original disk.

Figure 5.15 Group Disk Selection: rootGroup



Group name, size and mount point will be displayed.

In the [Group Name], an automatically created group name appears as a default.

Change the [Group Name] if necessary.

In the [Group Configuration Disk] field, the original disk you selected will be displayed.

You cannot remove the original disk from the [Group Configuration Disk] field.

From the [Physical Disk List] field, select the mirror disk (i.e. the disk you want to mirror to) and click <Add>. The disk will be added to the [Group Configuration Disk] field.

You can select more than one physical disk at a time.

For mirroring disks, add one or more mirror disks. If no mirror disk is added, the original disk is registered with the root class and managed with GDS but will not be mirrored.

Double-click the [Disk Name] field in the [Group Configuration Disk] field, and change the name.

After adding all disks, click <Next> and create the next group.

Once you finish creating all groups, proceed to register the spare disks.




Note

Physical disks that can be registered as the Group Configuration Disk

GDS Management's system disk settings does not allow you to add a physical disk that is smaller than the original disk. Users should always add a physical disk that is larger than the original disk.

4. Registering spare disk

Figure 5.16 Spare Disk Selection



To register a spare disk, select the physical disk you want to set as a spare disk from the [Physical Disk List], and click <Add>.

Existing spare disks in the root class cannot be removed from the [Spare Disk] field.

After finishing the registration, click <Next>.

When you do not need to register a spare disk, no setting is necessary. Just click <Next>.

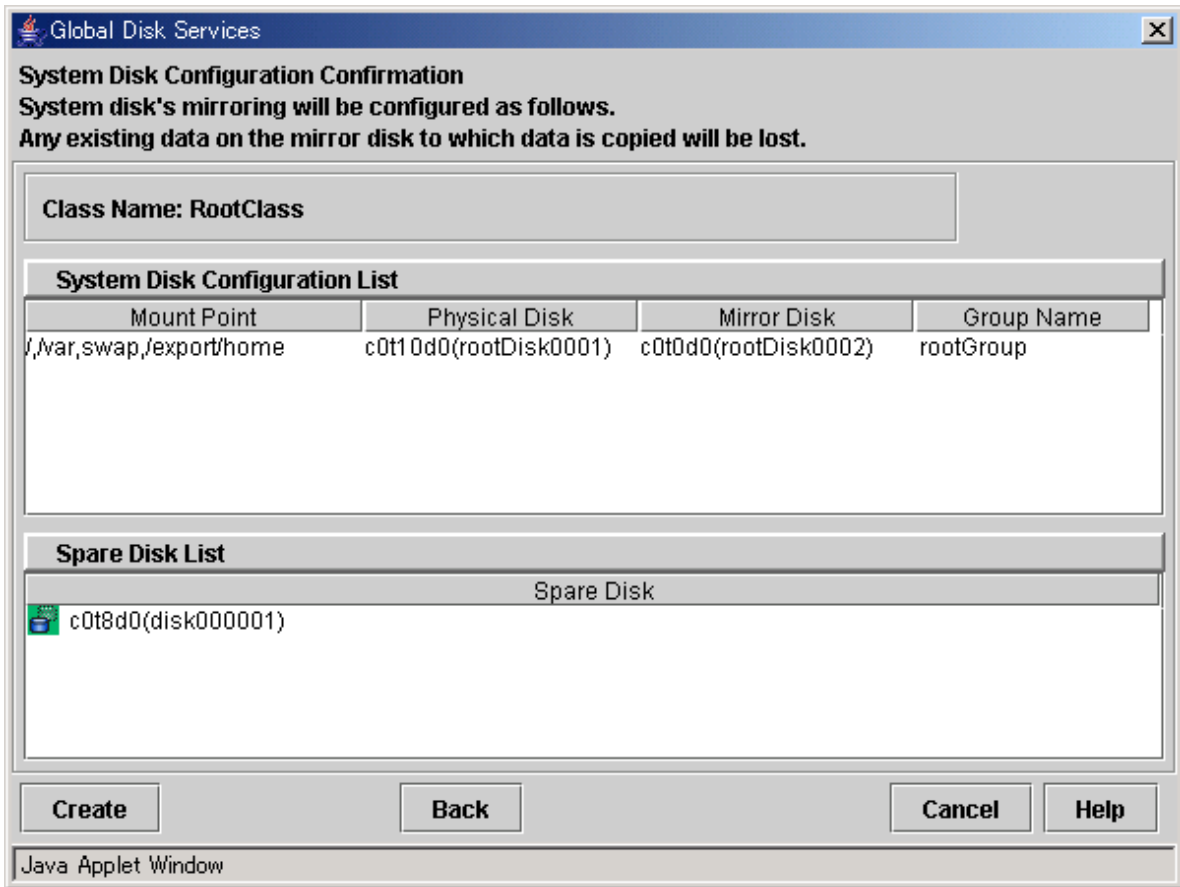
Point

Spare Disk Size

The hot spare function will not operate when there is not sufficient space on a spare disk to copy the configuration of volume in mirror group. Define the largest disk within the class as the spare disk.

5. Confirming system disk configuration

Figure 5.17 System Disk Configuration Confirmation



Confirm the system disk configuration.

In the [Physical Disk] field, the original disk will be displayed, and in the [Mirror Disk] field, the mirror disk will be displayed.

When mount information of a slice contained in physical disk is set for /etc/vfstab, mount point will be displayed in the [Mount Point] field.

Click <Create> to continue.

Information

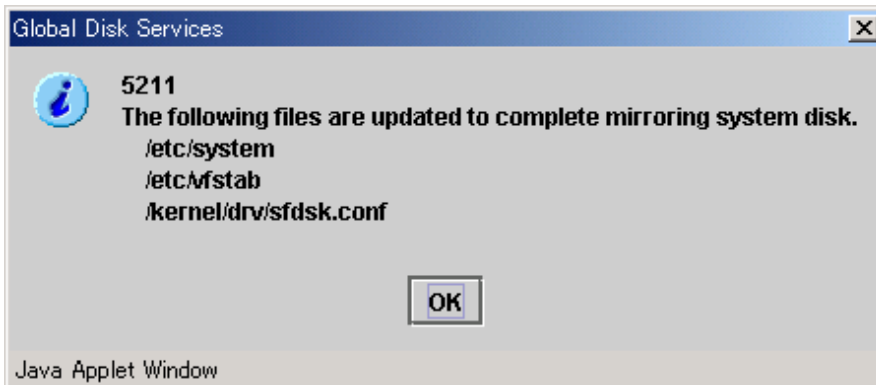
Automatically Generated Volume Names

The following volume names are automatically generated when setting the system disk.

- When mount information is set for /etc/vfstab, the name will be mount point + "Volume." (e.g. usrVolume) However, root partition will be rootVolume.
- When mount information is not set for /etc/vfstab, the name will be "Volume" + number. (e.g. Volume0001)

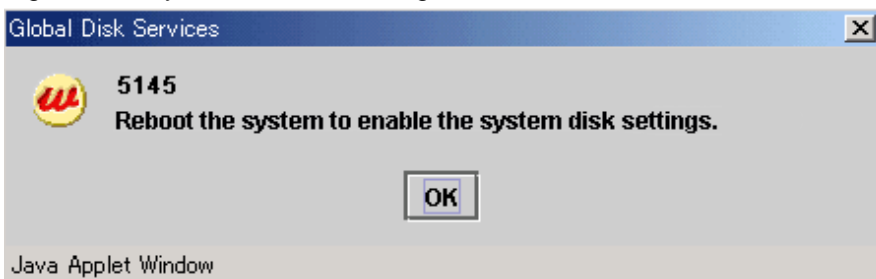
6. Completing system disk configuration

Figure 5.18 Setting System Disk Mirroring Complete



Confirm that the system disk configuration is complete, and click <OK>.

Figure 5.19 System Reboot Message



These system disk settings will be in effect after the system is rebooted. Click <OK> and reboot the system immediately.

Note

JRM and System Disk Settings

JRM (Just Resynchronization Mechanism) for a volume created by system disk settings will be set as "Yes."

To disable the JRM feature, select the volume in the Main screen and change the attribute by clicking [Change Attributes] in the [Operation] menu.

Note

GFS Partition in System Disk

Configuring a system disk causes the device special files of file systems within the system disk to be renamed. For this reason, when the system disk contains a partition used by the sfxfs file system, you must make the following modifications for the relevant partition prior to rebooting the system after completing the system disk configuration.

- Changing the partition information settings

For details on the method of changing the settings, refer to "PRIMECLUSTER Global File Services Configuration and Administration Guide" or the manual page of the sfxadm(IM) command.

- Modifying the /etc/vfstab file

You must change the device special file name from that for the physical slice to that for the logical volume.

In addition, a GFS Local File System cannot be used as /(root), /usr, /var, or /opt. Also a partition on the system disk cannot be used in a GFS Shared File System. For details refer to "PRIMECLUSTER Global File Services Configuration and Administration Guide."

5.2.2 Automatic Configuration Wizard

In order to configure mirroring for disks other than the system disk, use the following methods.

- Using Automatic Configuration Wizard
- Using the [Settings] menu

For details on operation using the [Settings] menu, see "[5.2.3 Operating from the Settings Menu.](#)"

Automatic Configuration Wizard allows you to create a volume by simply clicking <Next>.

In order to ensure an environment which allows an application to continuously access data at the time of I/O adapter, controller or cable failure, Automatic Configuration Wizard mirrors disks connected to different I/O adapters.

When there is an add-on expansion mechanism within the expansion unit (add-on file unit or add-on disk cabinet), expansion units are mirrored.

In order to ensure a safe mirroring environment, Automatic Configuration Wizard automatically creates one spare disk for every ten disks. For automatic configuration, make sure you have a minimum of three disks.



Note

Supported Disks

Disks that can be configured for mirroring with the Automatic Configuration Wizard are those disks whose physical disk names are in "cXtXdX" format. If all target disks are connected to one I/O adaptor, the Automatic Configuration Wizard cannot be used. There, perform configuration from the relevant configuration screen.



Note

Mirroring Expansion Units with Add-On Expansion Mechanism

When you are mirroring expansion units with add-on expansion mechanism, connect the two cables connected to the same expansion unit to the same I/O adapter on the processing unit.



Note

Shared Disk Mirroring

Shared disks in cluster system cannot be mirrored using the automatic configuration feature. Perform configuration from the relevant configuration screen.




Note

System Disk Mirroring

System disk cannot be configured automatically. Instead, perform mirroring of the system disk as in "[5.2.1 System Disk Settings.](#)"

To start Automatic Configuration Wizard, select [Automatic Configuration] in the [Settings] menu.

Figure 5.20 Automatic Configuration Environment Settings



Selecting the Configuration Environment for Mirroring

Set the configuration environment in order to create a volume using the Automatic Configuration Wizard.

a. Selecting the [Class Name]

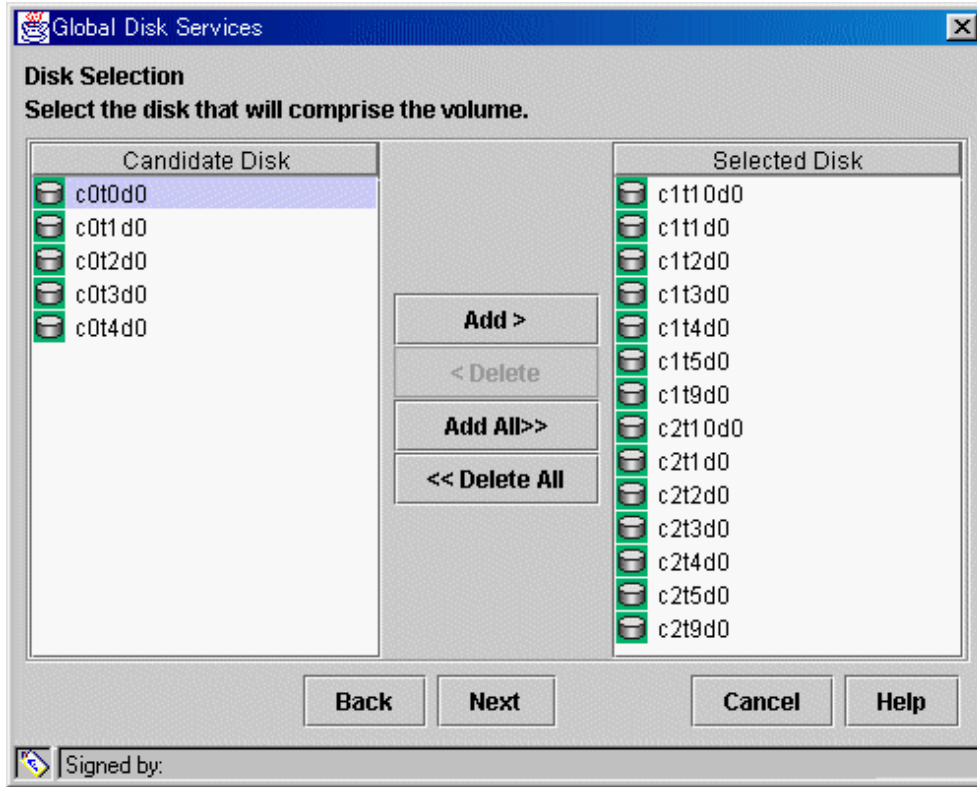
If you clicked <New> in the [Class Name] list, a new class will be created for the volume. If you are creating a volume within an existing class, select the appropriate class from the list. If no classes have been created, the only option available on the list will be <New>.

b. Selecting the [Spare Disk]

When creating a spare disk from the selected disk, select [Create a spare disk]. If a spare disk is unnecessary, select [Do not create a spare disk].

Click <Next>, and the message below appears.

Figure 5.21 Disk Selection



Selecting Disks to Mirror

Disks available for mirroring appear in [Selected Disk] field.

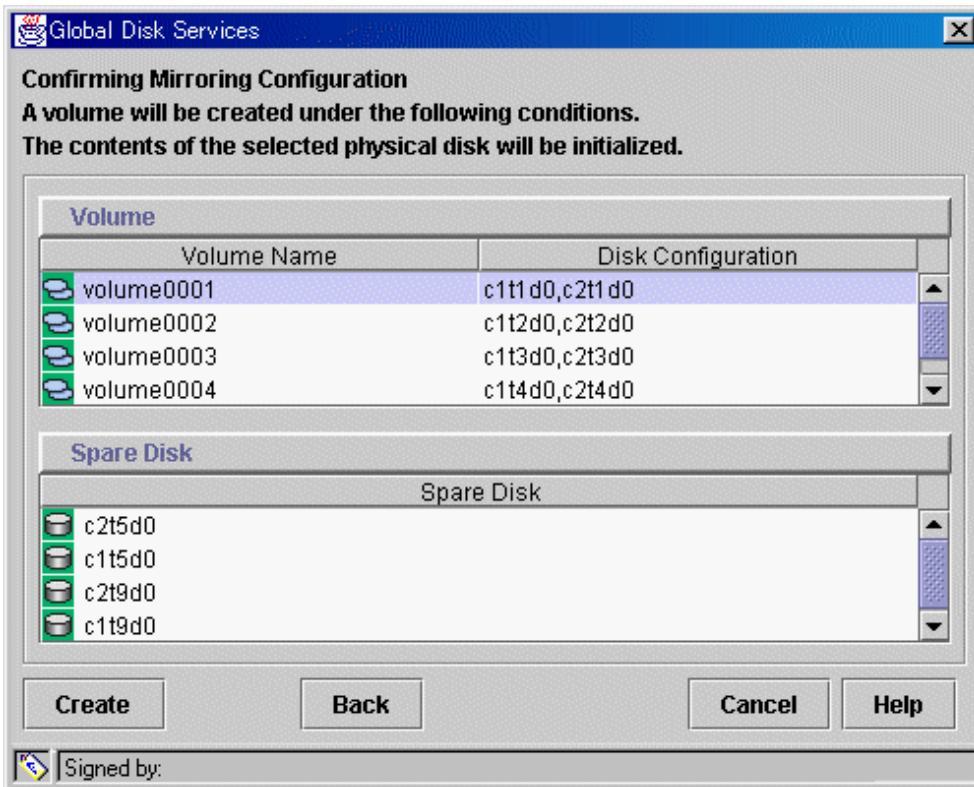
If there are disks you do not want to mirror, select the disk by clicking the icon, and click <Delete>. If there are disks you want to mirror in [Candidate Disk] field, select the disk by clicking the icon, and click <Add>.

When moving all disks in [Candidate Disk] field to [Selected Disk] field, click <Add All>.

When moving all disks in [Selected Disk] field to [Candidate Disk] field, click <Delete All>.

Click <Next>, and the message below appears.

Figure 5.22 Confirming Mirroring Configuration



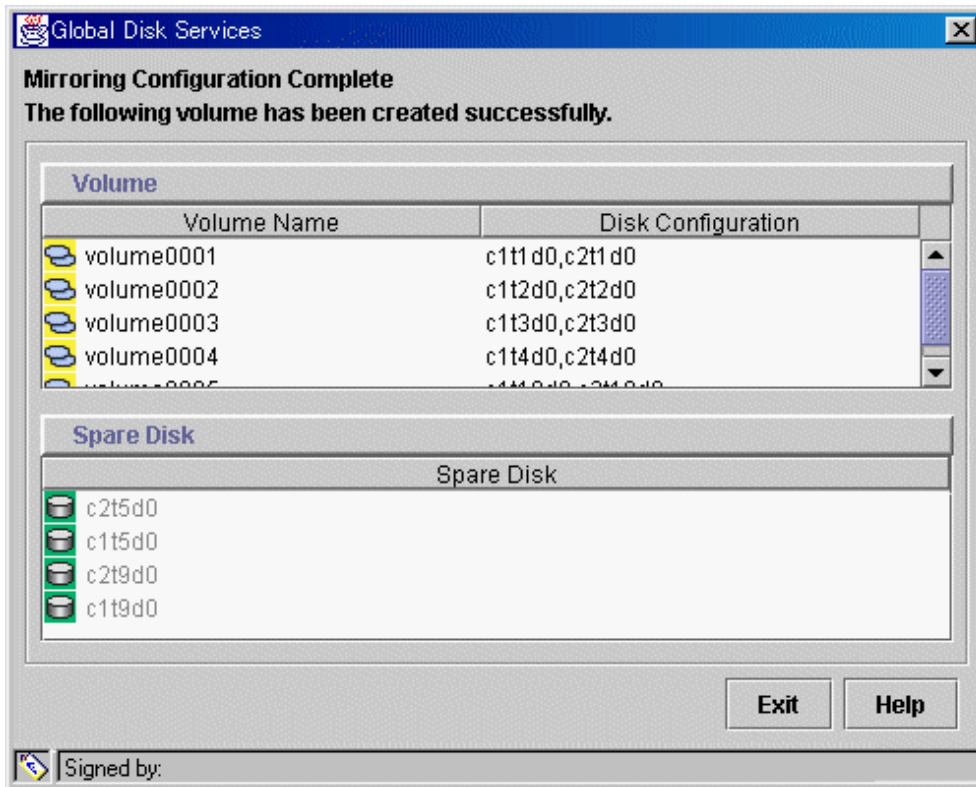
Confirming the Volume Configuration

GDS will automatically combine the disks selected in the previous screen and create a volume.

Volume name and its disk configuration, and the disk defined as spare disk will appear.

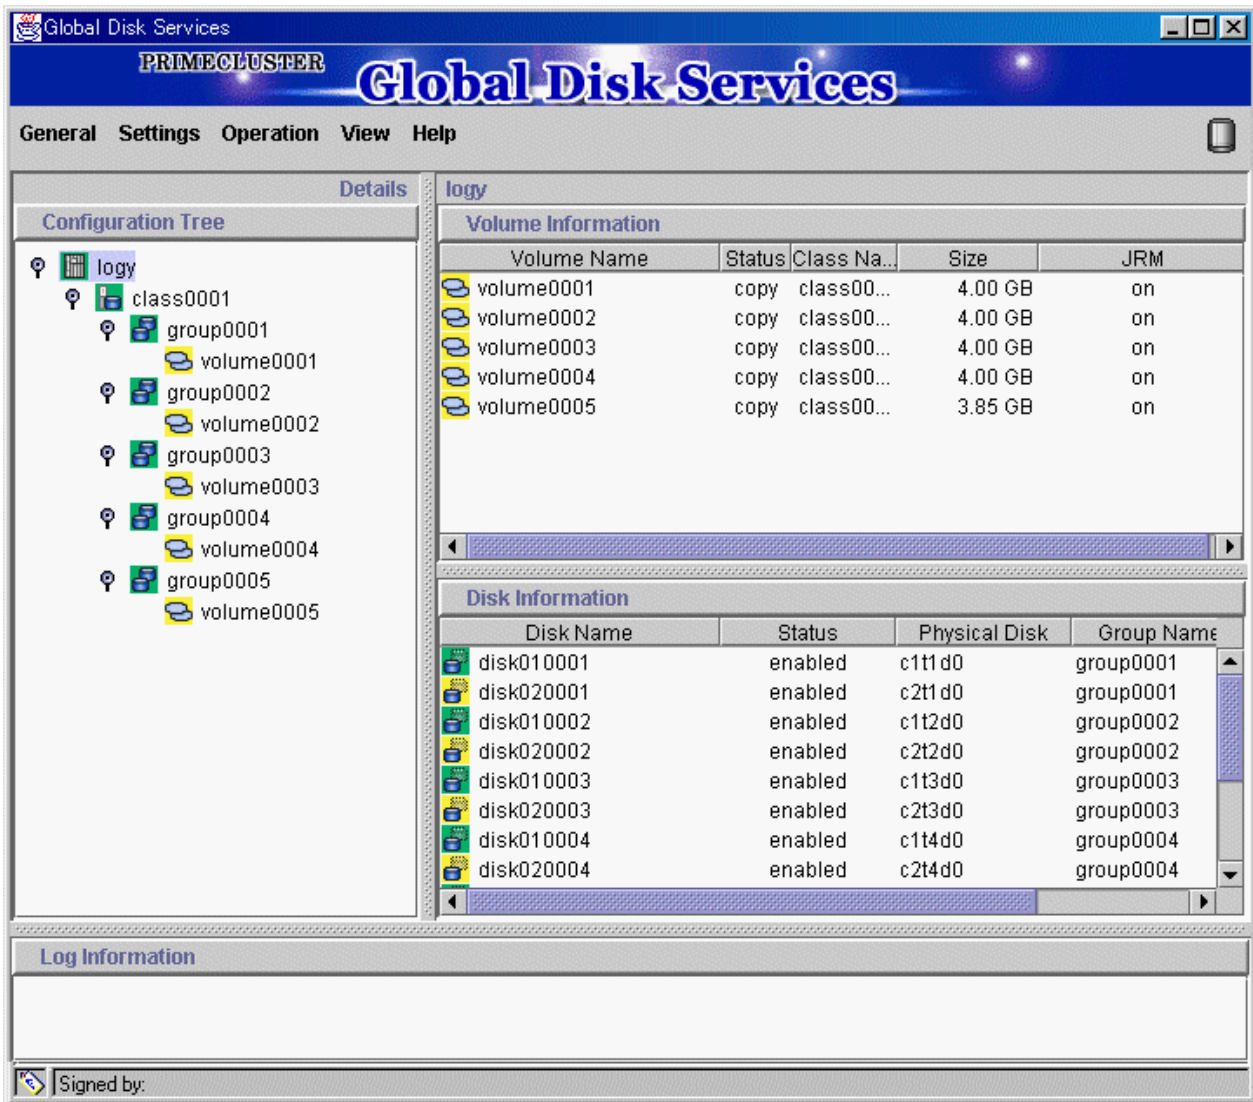
Click <Create> to continue.

Figure 5.23 Mirroring Configuration Complete



Confirm that the volume creation was successful, and click <Exit>.

Figure 5.24 Main Screen (after completing automatic configuration)



Volume information and disk information created using the Automatic Configuration Wizard will appear on the Main screen.

5.2.3 Operating from the Settings Menu

From the Main screen, select [XXXXX Configuration] in the [Settings] menu.

Settings screen for different features appears. Use the screen switch tab to switch between "Class Configuration," "Group Configuration" and "Volume Configuration."



Note

Configuring a System of High Reliability


In order to ensure a system of high reliability, mirroring disks connected to separate I/O adapter, controller and cable is recommended.

5.2.3.1 Class Configuration

In this section, how to create a new class is explained.

In the [Settings] menu, select [Class Configuration]. Class Configuration screen appears.

Figure 5.25 Class Configuration



1. Selecting class name

In the [Class Name] list, select "New."

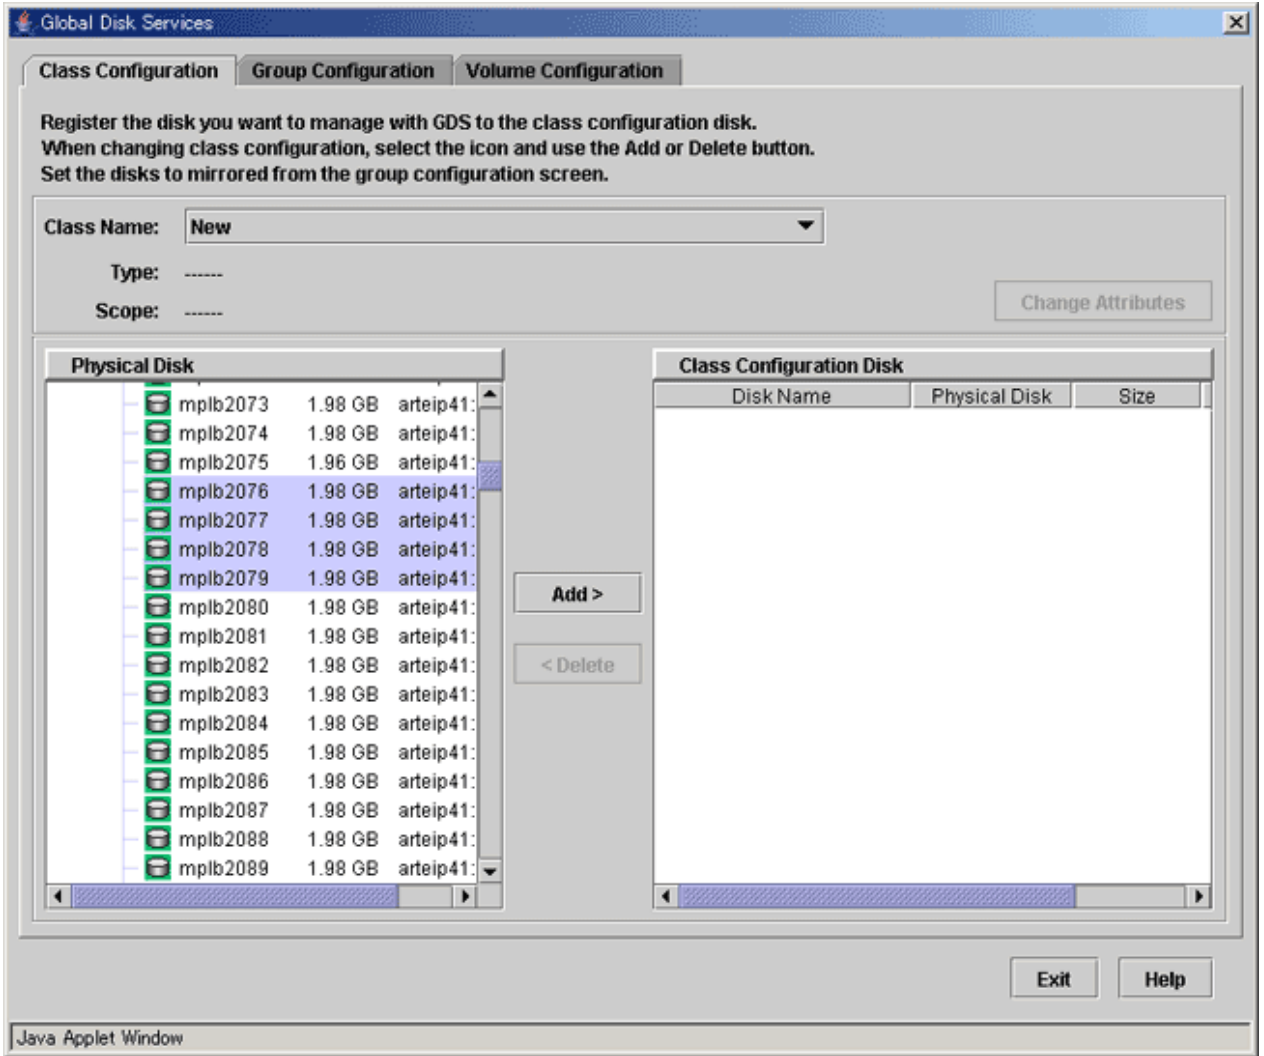
2. Selecting physical disk

In the [Physical Disk] field, select the physical disk you want to include in the disk class.

You can select more than one physical disk at a time.

Selecting a physical disk will make the <Add> button available.

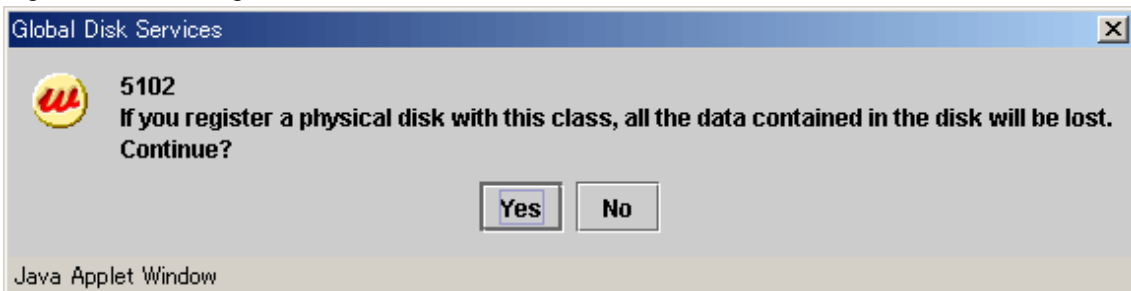
Figure 5.26 Selecting Physical Disk to Configure Class



3. Creating a class

<Add>, and the message below appears.

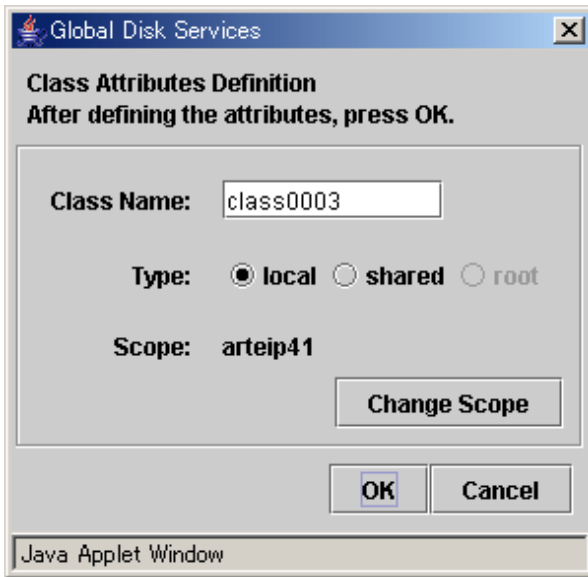
Figure 5.27 Warning



Click <Yes> to continue, and <No> to cancel.

Click <Yes>, and Class Attributes Definition screen appears.

Figure 5.28 Class Attributes Definition



In the [Class Name] of Class Attributes Definition screen, an automatically created disk class name appears as default. Change the [Class Name] if necessary, and click <OK>.

The [Type] is "shared" and the [Scope] indicates a node that can share disks with respect to the class. To accept the settings, click the <OK> button.

Information

Inputting Class Name

Class name will be used for the device path name.

/dev/sfdsk/class_name/rdisk/volume_name

/dev/sfdsk/class_name/dsk/volume_name

You must be careful when inputting the class name, as once the volume is created, it cannot be changed.

Note

Creating Local Type Class with Cluster System

When creating a local type class in cluster system, class name should be set differently for each node.

See

For information on assigning a class name, see "A.1.1 Object Name."

This operation determines the class name.

When using single node, [Type] is fixed to "local" and you cannot change it.

If you click <Cancel> in the Class Attributes Definition screen, registration of the physical disk itself will be canceled.

4. Setting disk attributes

Selecting the [Disk Name] in the [Class Configuration Disk] field allows you to set the disk attributes. From here, you can change the [Disk Name] and [Type].

a. Changing disk name

Double-click the [Disk Name] in the [Class Configuration Disk] field, and change the name.

b. Changing disk type

Display the [Disk Type] in the [Class Configuration Disk] field and select the disk type you wish to change from the list.
When specifying as a spare disk select "spare." When specifying as a single disk, select "single." Default is set to "undef."

5. Completing class creation

After creating all classes, click <Exit> and close Class Configuration screen.



Classes in Cluster Systems

- Creating a class adds a class resource, and removing a class removes a resource.
- When removing a class resource, remove the class without using the PRIMECLUSTER cldelrsc(1M) command.
- Cluster applications that use resources of a class should be set after the volume configuration is complete.



For information on how to create a shared type class in cluster systems, see "[5.2.3.2 Cluster System Class Configuration](#)."

5.2.3.2 Cluster System Class Configuration

In cluster systems, specify [Type] and [Scope] in the Class Attributes Definition screen.

1. **Setting [Type]**

Sets the class type.

When creating a new disk class, selecting a physical disk that is not shared by other nodes from the [Physical Disk] field, sets the type to "local" by default.

Selecting a shared physical disk, sets the type to "shared."

2. **Displaying [Scope]**

Displays connecting nodes that can share a class.

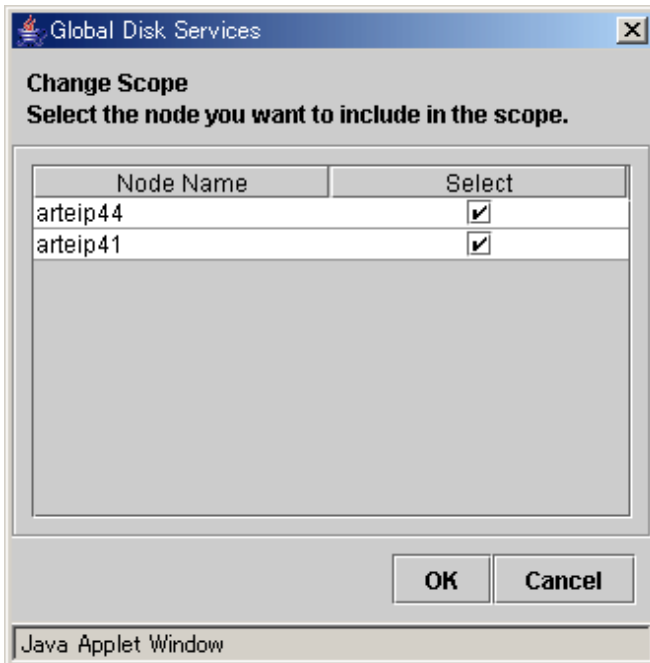
To change scope, click <Change Scope>.

3. <Change Scope> Button

Changes nodes connecting to class.

Click <Change Scope>, and Change Scope screen appears.

Figure 5.29 Change Scope Screen



- Changing the checkbox in the [Change Scope] dialog box specifies the connecting node. (Multiple specification possible).
- Clicking < OK > of the [Change Scope] dialog box determines the class connecting node.
- Clicking <Cancel> of the [Change Scope] dialog box cancels the changing of the connecting node.

Note

- In a cluster system, creating a class adds a class resource, and removing a class removes a resource.
- When removing a class resource, remove the class without using the PRIMECLUSTER cldelrsc(1M) command.
- Cluster applications that use resources of a class should be set after the volume configuration is complete.

See

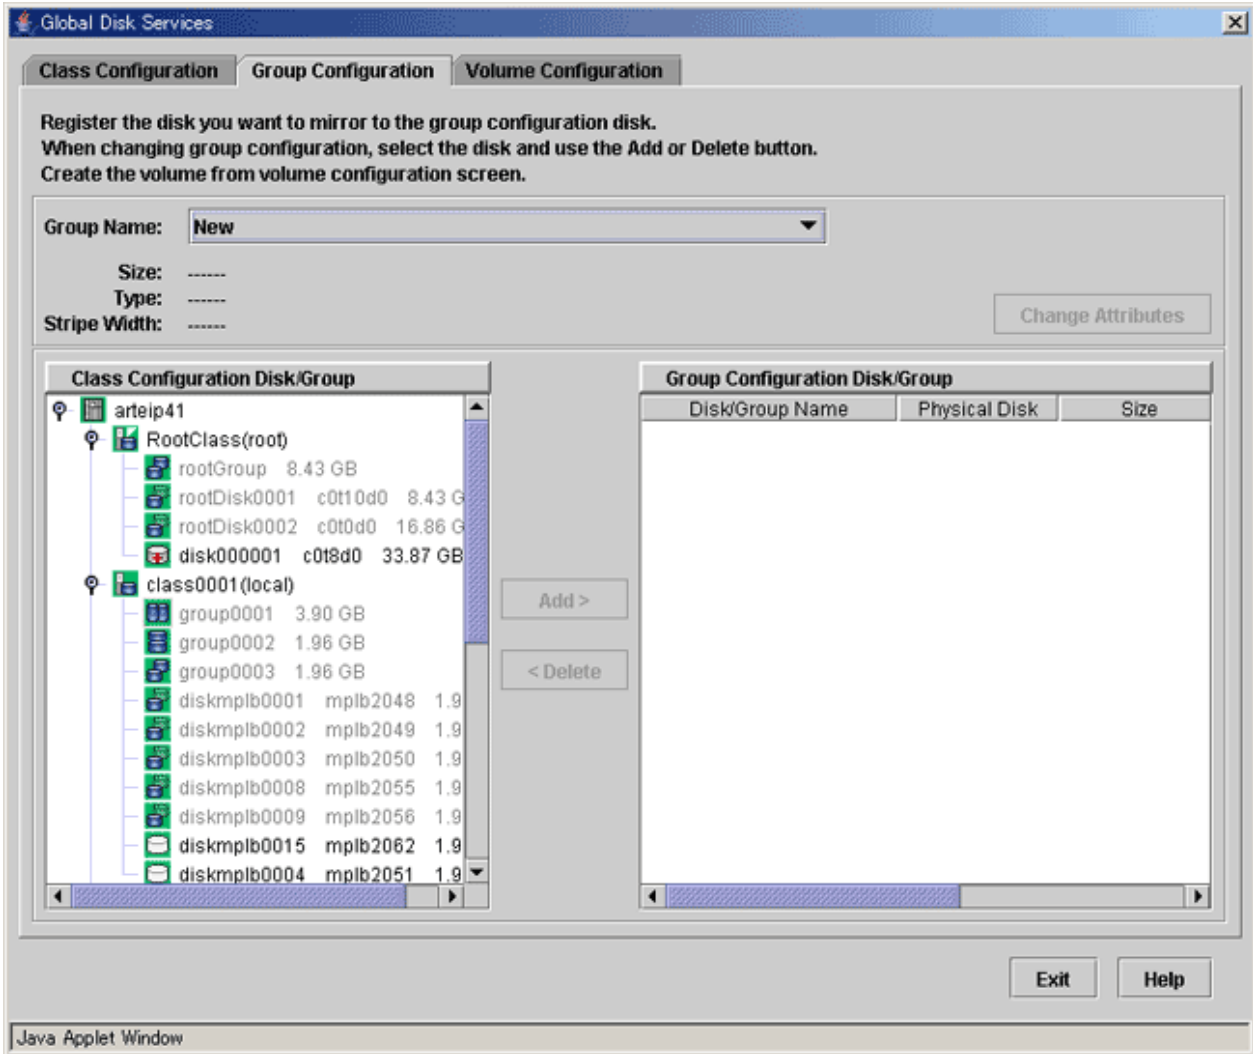
For information on installation and initial setting of cluster systems, refer to "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide."

5.2.3.3 Group Configuration

In this section, how to create a new group is explained.

In the [Settings] menu, select [Group Configuration]. Group Configuration screen appears.

Figure 5.30 Group Configuration



Follow the procedures below to create a new group.

1. Selecting group name

In the [Group Name] list, select "New."


2. Selecting disk/lower level group

In the [Class Configuration Disk/Group] field, select the disk/lower level group you want to include in the disk group.

You can select more than one disk/group at a time.

Selecting a disk/group will make the <Add> button available.

Figure 5.31 Selecting Disk/Lower level group to Configure Group



3. Creating a group

Click <Add>, and the Group Attributes Definition screen appears. You will be able to specify the group attributes such as group name, type and stripe width.

Figure 5.32 Group Attributes Definition Screen



a. **Setting [Group Name]**

Enter the group name.
Change the default group name if necessary.

b. **Setting [Type]**

Set the group type.
Select "mirror" for mirroring, "stripe" for striping, and "concat" for concatenating. The default setting is "mirror."

c. **Setting [Stripe Width]**

You will be able to enter this field only when you select "stripe" for the [Type].
For the stripe width, you can specify a value of two raised to the power. The default setting is "32."
After setting the attributes, click <Exit> and a new group will be created.
If you click <Cancel> in the Group Attributes Definition screen, connection of the disk itself will be canceled.



See

.....
For information on assigning a group name, see "[A.1.1 Object Name.](#)"
.....

4. **Completing group creation**

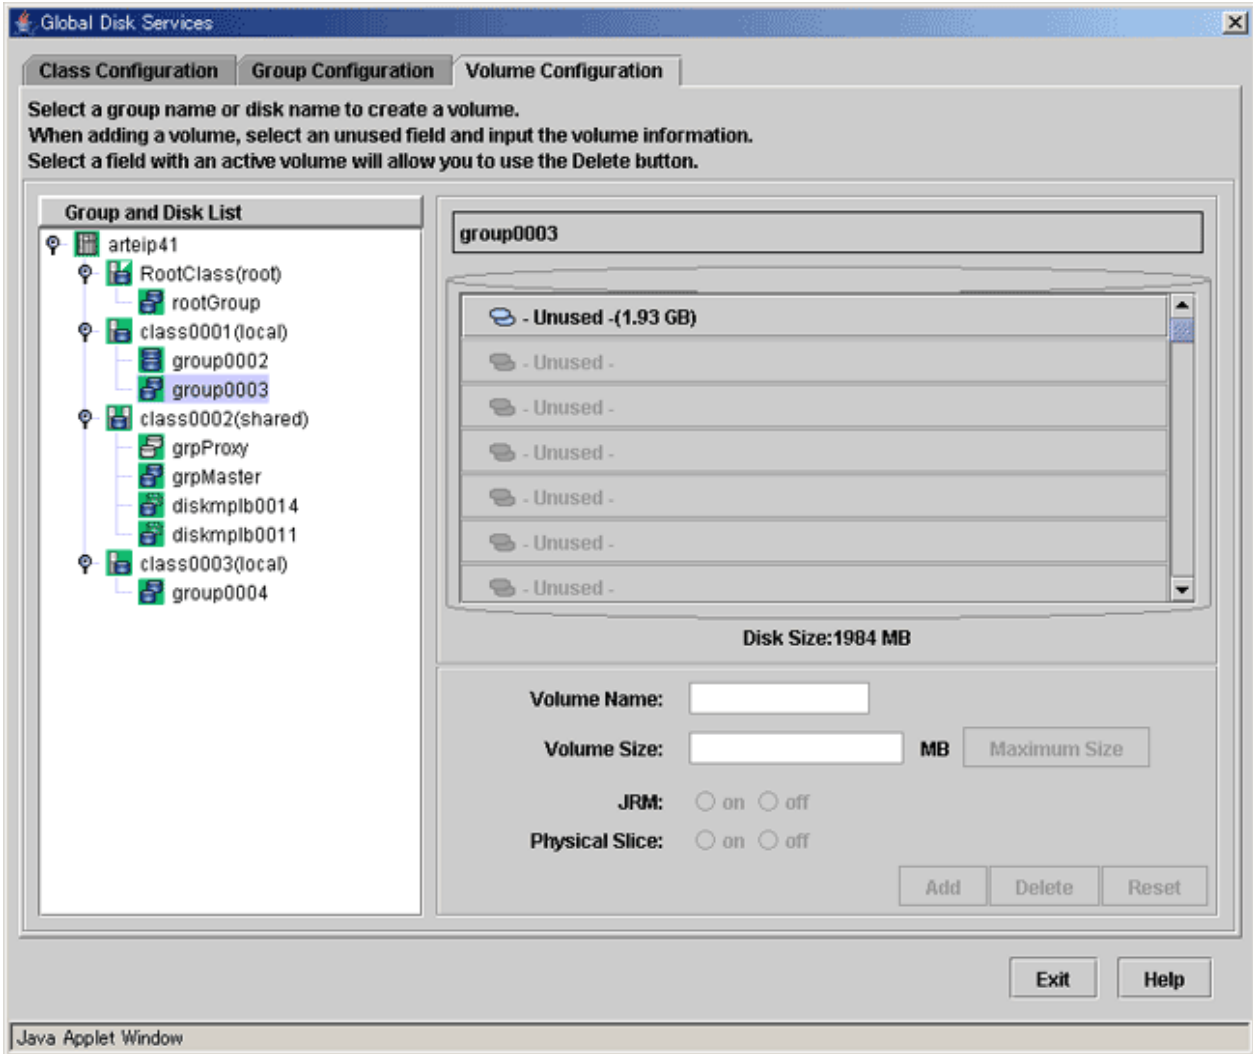
After creating all groups, click <Exit> and close Group Configuration screen.

5.2.3.4 Volume Configuration

In this section, how to create a new volume is explained.

In the [Settings] menu, select [Volume Configuration]. The Volume Configuration screen appears.

Figure 5.33 Volume Configuration



Follow the procedures below to create a volume.

1. Selecting group/disk

In the [Group and Disk List], select the group or the disk.

2. Selecting unused volume

Click <Unused> with a volume icon, and select an unused volume.

3. Setting volume attributes

Selecting an unused volume will allow you to type in the volume attributes field (Volume Name, Volume Size, JRM, Physical Slice).

a. Setting [Volume Name]

Type the volume name.



See

.....
For information on assigning a volume name, see "A.1.1 Object Name."
.....

b. Setting [Volume Size]

Type the volume size in MB units, using numbers only.



Note

Size and Cylinder Boundary

When the specified size does not match the cylinder boundary, it will automatically adjust by rounding up. For details, see "A.2.7 Volume Size."

c. [Maximum Size] button

Sets the value in the [Volume Size] field to the maximum available size.

d. Setting [JRM]

Default is set to "on." Change the setting when you want to disable the just resynchronization feature.

When you select a stripe group or concatenation group in step 1., the setting will fail.

e. Setting [Physical Slice]

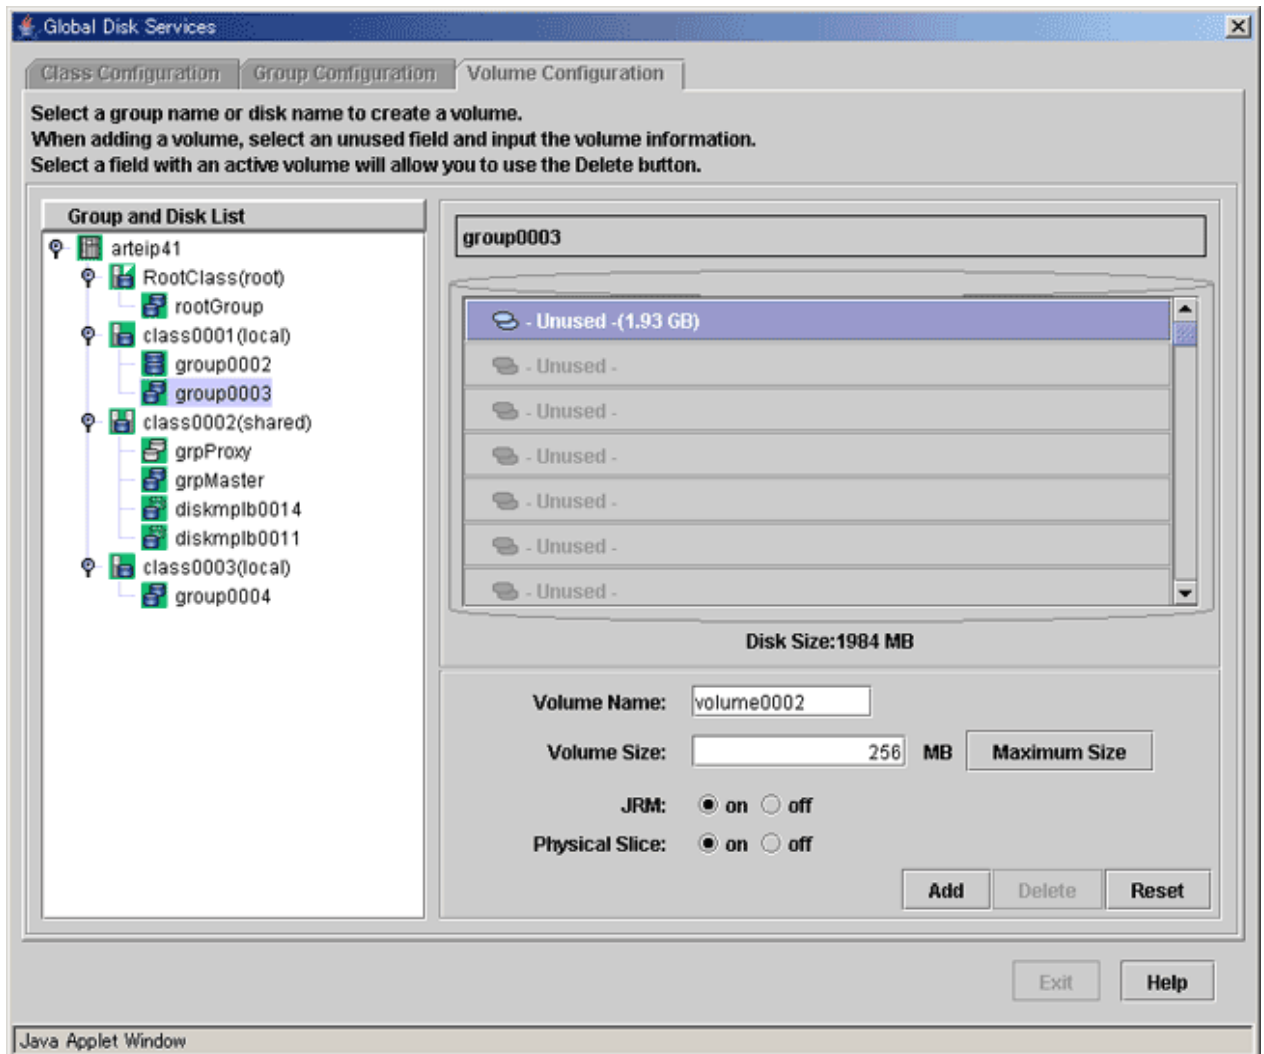
Sets the volume's physical slice attribute value. Default is the "on."

on : a volume which consists of physical slices is created.

off : a volume without physical slices is created.

When you select a stripe group or concatenation group in step 1., the setting will fail.

Figure 5.34 Setting Volume Attributes



4. Determining the new volume

After setting the attributes, click <Add>. A new volume will be created.

If you click <Reset>, creation of the new volume will be canceled.

After creating all volumes, click <Exit> and close Volume Configuration screen.

After creating the volume, the volume will be started. You can access to the volume using the following special files.

/dev/sfdsk/class_name/rdisk/volume_name

/dev/sfdsk/class_name/dsk/volume_name



Note

Shared Class Volume Operation

Volume created in a shared class cannot be used from other nodes immediately. If you wish to access from other nodes, you must activate the volume from the node you wish to access.

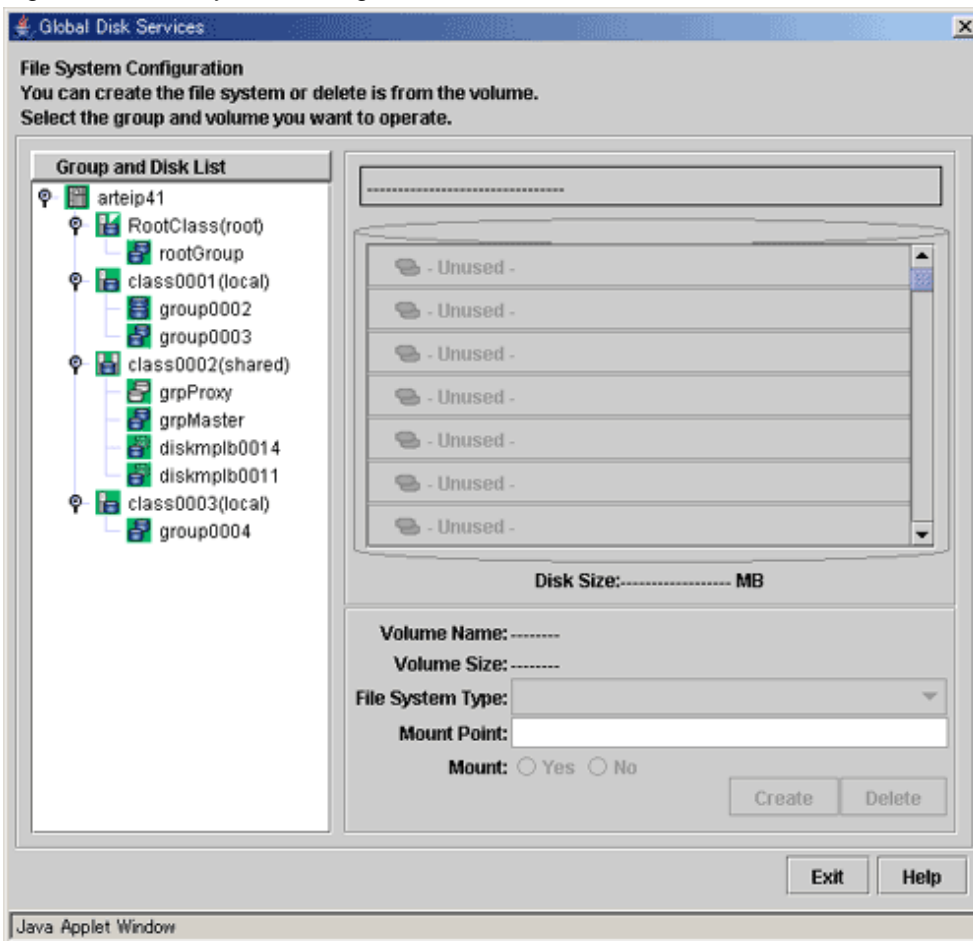
After activating the volume, the node which created the volume, and the node which activated the volume will gain access. However, since operating from two nodes could affect data consistency, you must be careful when operating the volume.

5.2.4 File System Configuration

In this section, how to create a file system in a volume is explained.

In the [Settings] menu, select [File System Configuration]. File System Configuration screen appears.

Figure 5.35 File System Configuration



1. **Selecting a group/disk**

In the [Group and Disk List] field, select a group or disk with which you want to perform an operation.

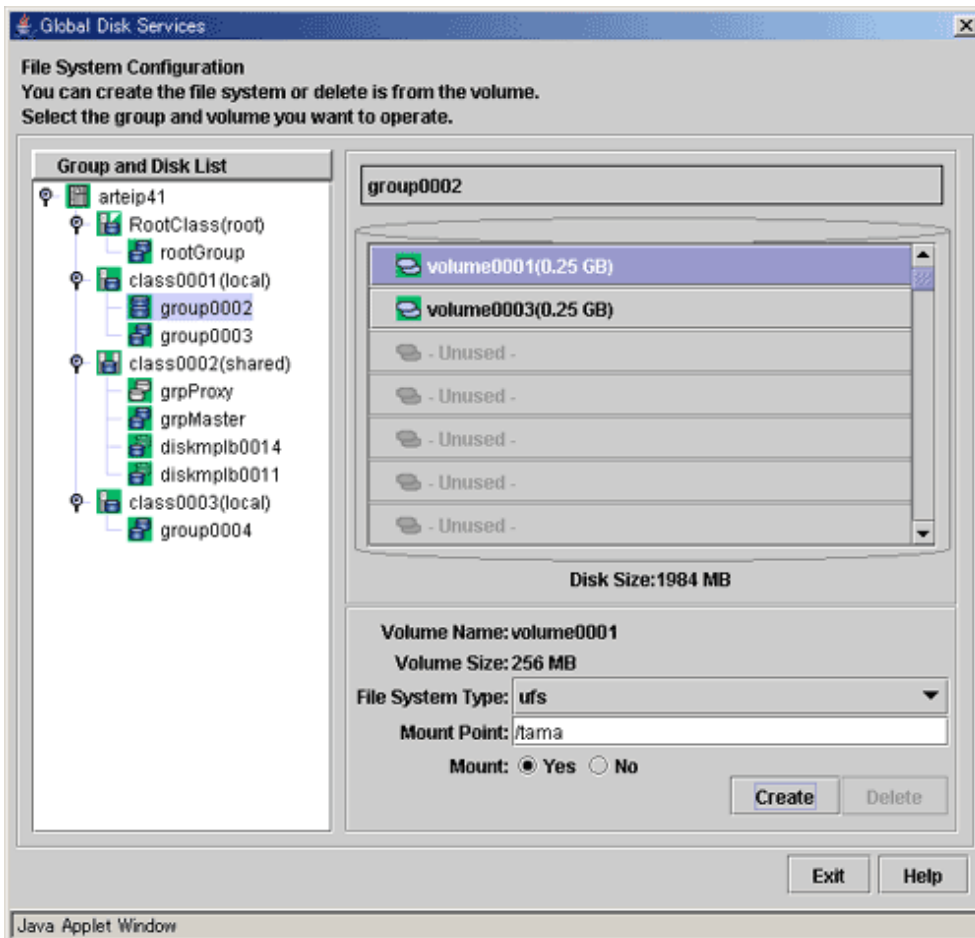
2. **Selecting a volume**

Select a volume in which you want to create a file system.

3. **Setting file system attributes**

Selecting a volume will allow you to type in the file system attributes field (File System Type, Mount Point, Mount).

Figure 5.36 Setting File System Attributes



a. **Setting [File System Type]**

Select the file system type. It is set to "ufs" by default.
When using GFS Local File System, select "sfxfs."

b. **Setting [Mount Point]**

Type the mount point you want to set for /etc/vfstab.

c. **Setting [Mount]**

Select "Yes" if you want to automatically mount when you start the volume, and "No" if not.

Selecting "Yes" will mount the file system to the specified mount point immediately after its creation.

However, you cannot select "Yes" when the volume in which you will create the file system exists on a shared disk unit.



Note

When Selecting "No" for "Mount"

When you select "No" for "Mount," [mount at boot] field in the /etc/vfstab file will be set to "No" and the new mount information will be added.

Change the mount information as necessary.

4. Creating the file system

After setting the attributes, click <Create>. A new file system will be created.

After creating all file systems, click <Exit>.



In Cluster Systems

For using volumes on shared disks as file systems in a cluster system, certain settings are required after creating the file systems.

For details on how to set file systems created on shared disks, see "PRIMECLUSTER Reliant Monitor Services (RMS) with Wizard Tools Configuration and Administration Guide."

5.2.5 Proxy Configuration

This section describes the procedures for relating proxy objects (proxy volumes or groups) to master objects (volumes or groups) in a system on which GDS Snapshot is installed.

The following two methods are available.

- Join

Relate a proxy to a master and join them. Synchronization copying from the master to the proxy is launched, and after the copying is complete, they are synchronized. When joining groups, proxy volumes are created in the proxy group and they are joined to their corresponding master volumes.

- Relate

Relate a proxy to a master and leave them parted. The statuses and contents of the master and the proxy remain unchanged.

5.2.5.1 Join

This subsection describes procedures for joining proxy volumes or proxy groups to volumes or groups in GDS Snapshot installed systems.

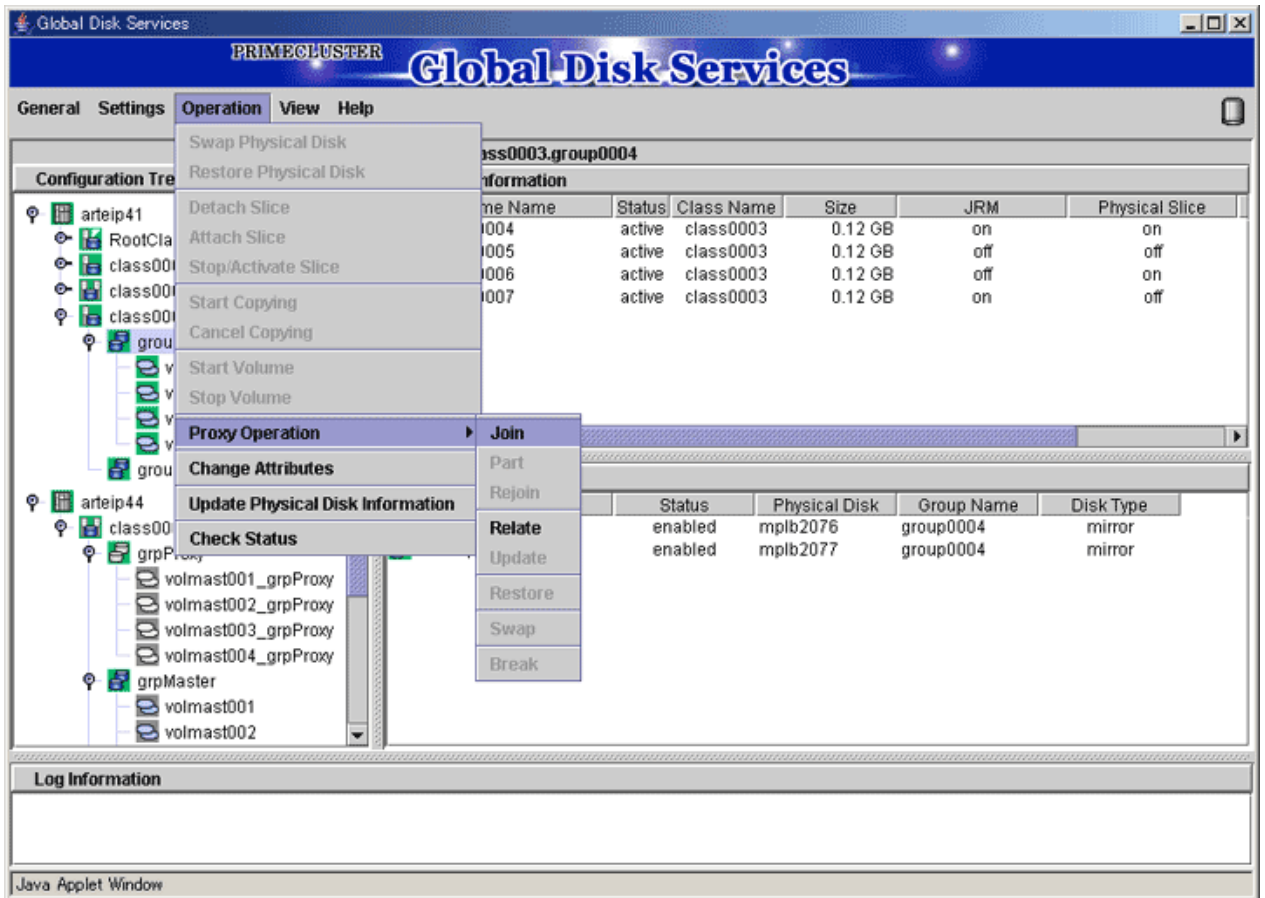
1. Selecting a master volume or a master group

Click the icon of a mirror volume, a single volume, or a mirror group to which a proxy is related in the Main Screen.

2. Selecting the [Join] menu

Select [Operation]:[Proxy Operation]:[Join] in the Main screen.

Figure 5.37 Join




3. Selecting a proxy to be joined

The Select Proxy dialog box appears.

Information

The following figure shows the window for group connection. On the window for volume connection, the <OK> button appears in the position of the <Next> button.

Figure 5.38 Select Proxy



Select a volume or a group to be joined to the master volume or the master group from the [Available Groups/Volumes].

Volumes or groups conforming to all of the following conditions are selectable.

Object	Volume	Group
Conditions	<ul style="list-style-type: none"> - Belongs to the class of the master volume - Equal to the master volume in size - Mirror type or single type (Volumes created in hierarchized mirror groups are also selectable. Any mirroring multiplicity is supported.) - Does not belong to the group or single disk of the master volume - Is not related to other master objects or proxy objects 	<ul style="list-style-type: none"> - Belongs to the class of the master group - Mirror type (Hierarchized mirror groups are also selectable. Any mirroring multiplicity is supported.) - Includes no volume - Is not related to other master objects or proxy objects

Note

Status of Proxy Volumes That Can Be Joined

"active" proxy volumes cannot be joined. To join an "active" proxy volume, stop the proxy volume in advance.

When joining volumes, select a volume to be joined to the master volume and click <OK>.

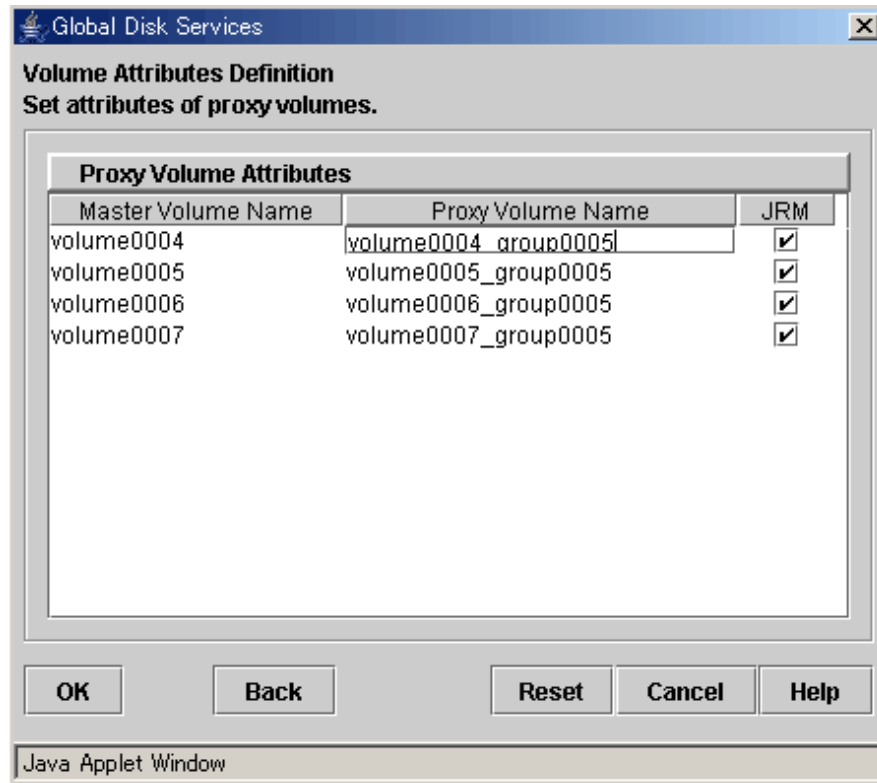
When joining groups, select a group to be joined to the master group and click <Next>.

Clicking <Cancel> cancels the join process.

4. Defining attributes of proxy volumes created in the proxy group

When joining groups, the Volume Attributes Definition dialog box appears.

Figure 5.39 Proxy Volume Attributes Definition



When a proxy group is joined to a master group, proxy volumes are created within the proxy group and joined to corresponding master volumes within the master group. In the Volume Attributes Definition dialog box, set attributes of such proxy volumes.

a. Proxy Volume Name

Assign volume names of proxy volumes. [Proxy Volume Name] shows default volume names. To change the default value, click and edit the volume name.



See

For the volume naming conventions, see "A.1.1 Object Name."



Information

Automatically Generated Proxy Volume Names

Proxy volume names are automatically generated as "master volume name" + "_" (underscore) + "proxy group name" (e.g. volume0001_group0002).

b. JRM

Set JRM (Just Resynchronization Mechanism) for volumes. The default value is "on." To turn "off", uncheck the [JRM] box.

Information

JRM for Volumes

The JRM setting in the Proxy Volume Attributes Definition dialog box is the "JRM for Volumes" of the proxy volume. Note that it is not the "JRM for Proxies." For details, see "[A.2.14 Just Resynchronization Mechanism \(JRM\)](#)."

When the settings are complete, click <OK>.

To initialize the proxy volume attribute settings, click <Reset>.

Clicking <Cancel> cancels the join process.

5. Confirming


A confirmation dialog box appears asking you whether to join the proxy.

Figure 5.40 Confirming "Join" (Volumes)



To continue the process, click <Yes>. Clicking <No> displays the Select Proxy dialog box shown in step 3. again.

Figure 5.41 Confirming "Join" (Groups)

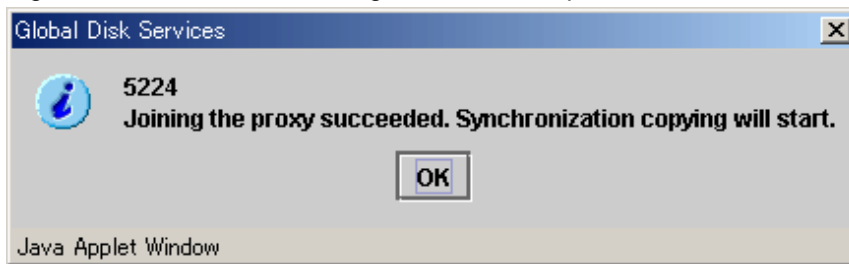


To continue the process, click <Yes>. Clicking <No> displays the Volume Attributes Definition dialog box shown in step 4. again.

6. Information message of the completion

A message window appears informing you that the join process is complete.

Figure 5.42 Information Message of "Join" Completion



Click <OK> to close the information message window.

5.2.5.2 Relate

This subsection describes the procedures for relating proxy volumes or proxy groups to volumes or groups in GDS Snapshot installed systems.

Information

Data of Related Masters and Proxies

Even if masters and proxies are related, data of the master objects and the proxy objects remain unchanged.

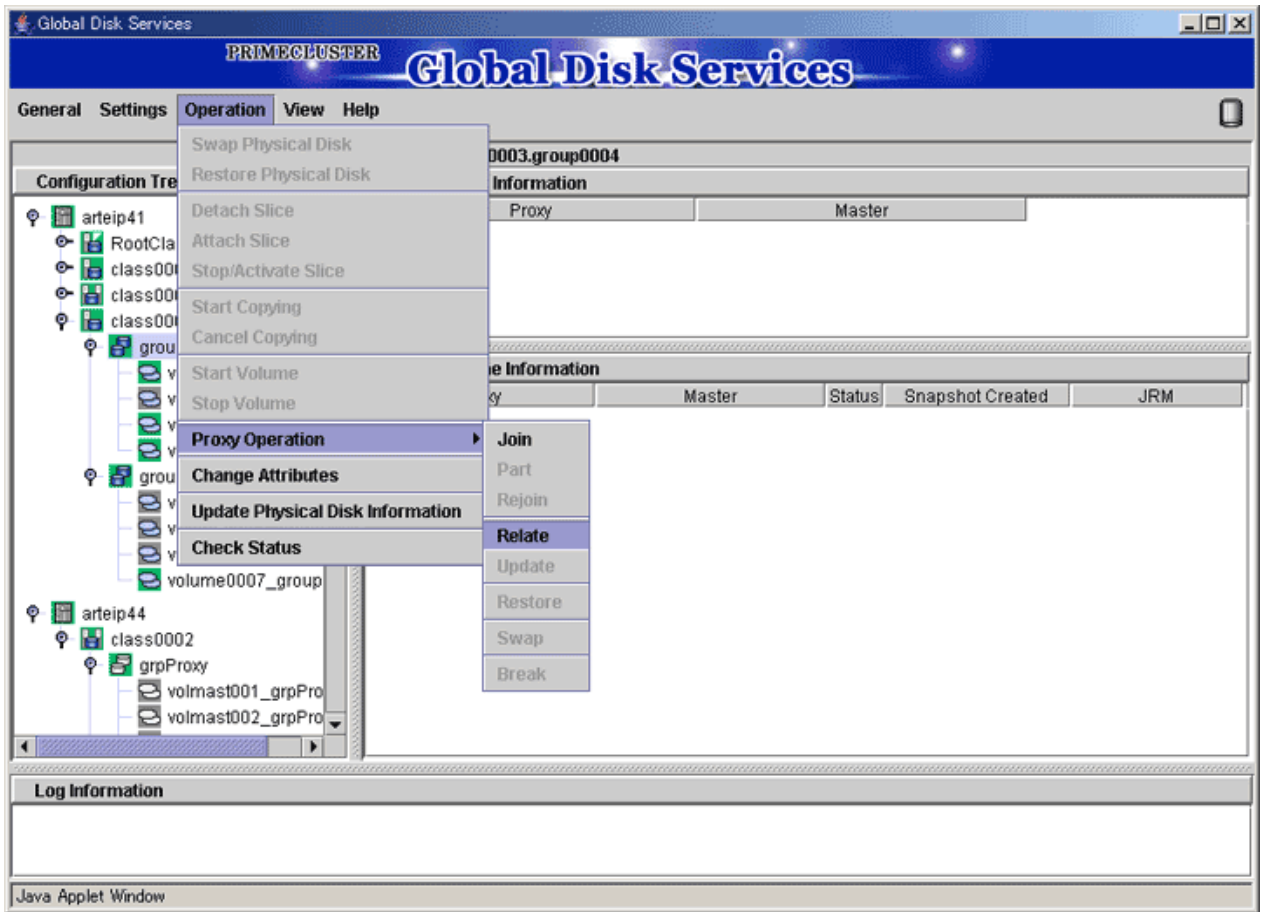
1. **Selecting a master volume or a master group**

Click the icon of a volume or a group to which a proxy is related in the Main screen.

2. Selecting the [Relate] menu

Select [Operation]:[Proxy Operation]:[Relate] in the Main screen.

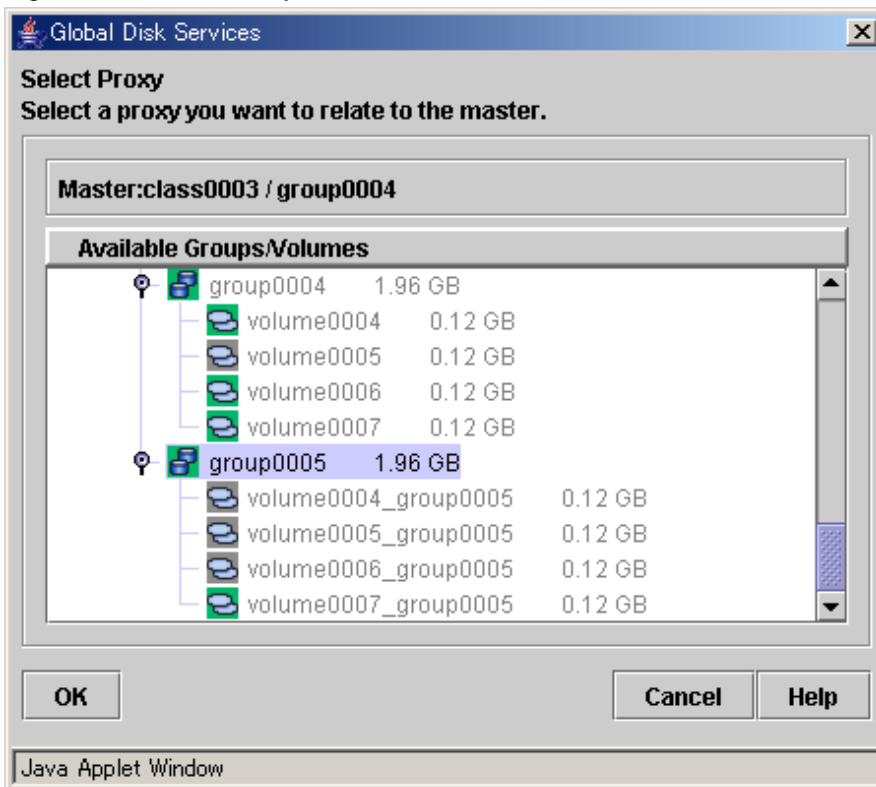
Figure 5.43 Relate



3. Selecting a proxy to be related

The Select Proxy dialog box appears.

Figure 5.44 Select Proxy



Select a volume or a group to be related to the master volume or the master group from the [Available Groups/Volumes].

Volumes or groups conforming to all of the following conditions are selectable.

Object	Volume	Group
Conditions	<ul style="list-style-type: none"> - Belongs to the class of the master volume - Equal to the master volume in size - Mirror type or single type (Volumes created in hierarchized mirror groups are also selectable. Any mirroring multiplicity is supported.) - Does not belong to the group or single disk of the master volume - Is not related to other master objects or proxy objects 	<ul style="list-style-type: none"> - Belongs to the class of the master group - Mirror type (Hierarchized mirror groups are also selectable. Any mirroring multiplicity is supported.) - Volume layout (offset and size) is same as that of the master group - Is not related to other master objects or proxy objects

Select a volume or a group to be related to the master volume or the master group and click <OK>. Clicking <Cancel> cancels the relating process.

4. Confirming

A confirmation dialog box appears asking you whether to relate the proxy.

Figure 5.45 Confirming "Relate"




To continue the process, click <Yes>. Clicking <No> displays the Select Proxy dialog box shown in step 3. again.

5. Information message of the completion

A message window appears informing you that the relating process is complete.

Figure 5.46 Information Message of "Relate" Completion



Click <OK> to close the information message window.

5.3 Operation in Use

GDS operates monitoring and maintenance from the Main screen.

This section explains the following operations:

- Viewing Configurations/Statuses and Monitoring Statuses
- Backup
- Restore
- Disk Swap
- Disk Migration
- Copying Operation

5.3.1 Viewing Configurations/Statuses and Monitoring Statuses

In the Main Screen, object configuration and status viewing and status monitoring can be performed.

5.3.1.1 Confirming SDX Object Configuration

Select [SDX Object] in the [View]:[Details] menu to view object configurations according to the following units.

- Object configuration within a node

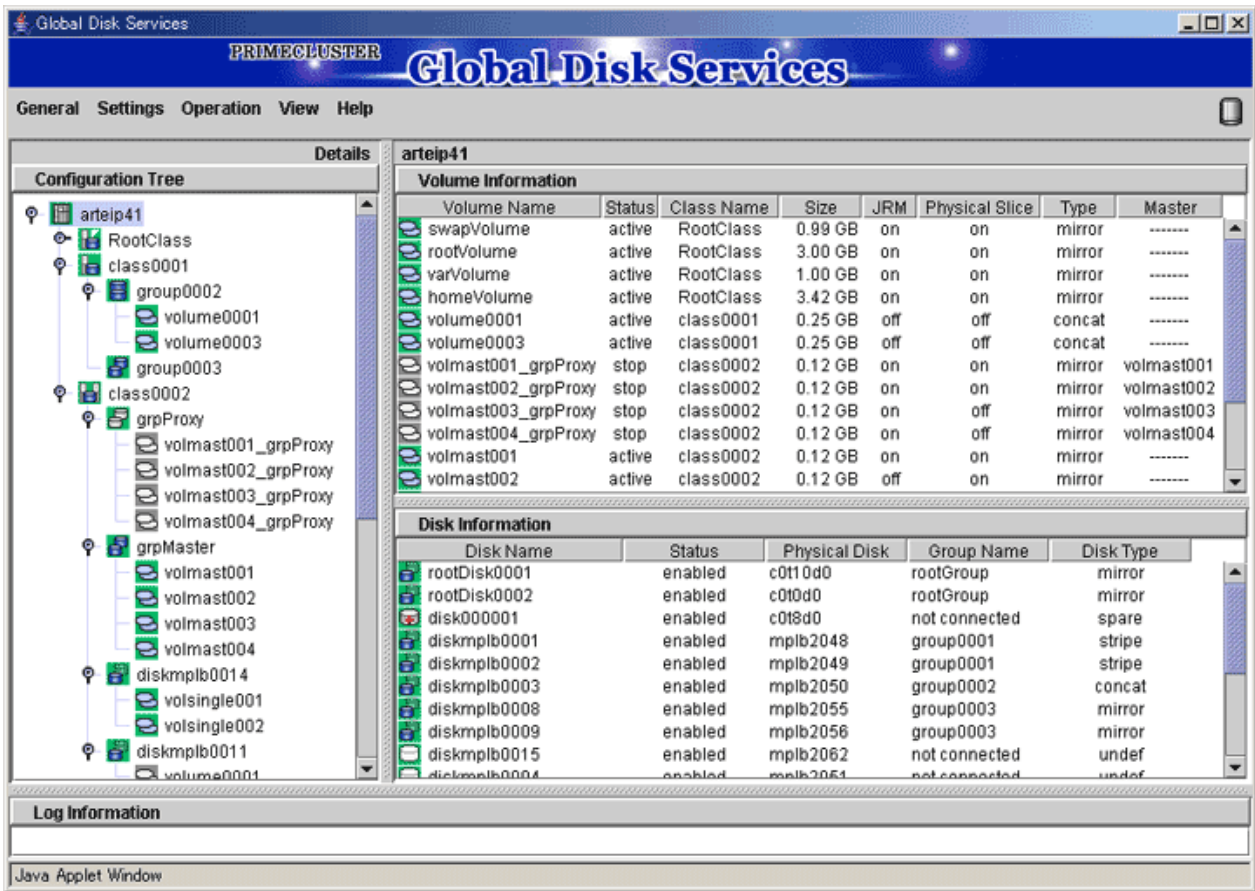
- Object configuration within a class
- Object configuration within a group
- Object configuration within a single disk
- Object configuration within a volume

For GDS Snapshot shadow objects, the object names, the status and so on are displayed in italics.

Object configuration within a node

Click the node icon in the Configuration Tree field, and all volumes and disks within the specified node appear.

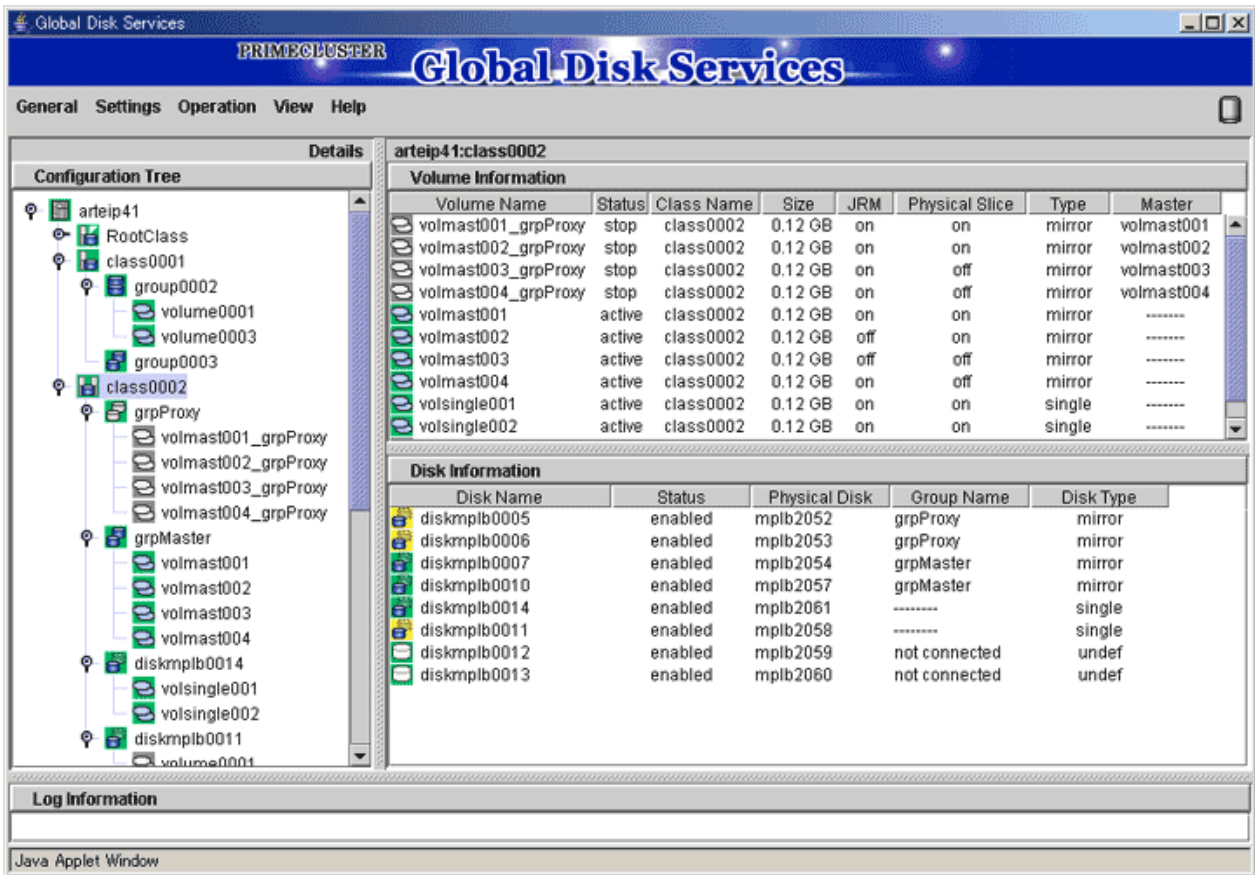
Figure 5.47 Main Screen (for SDX Objects of a Node)



Object configuration within a class

Click the class icon in the Configuration Tree field, and all volumes and disks within the class appear.

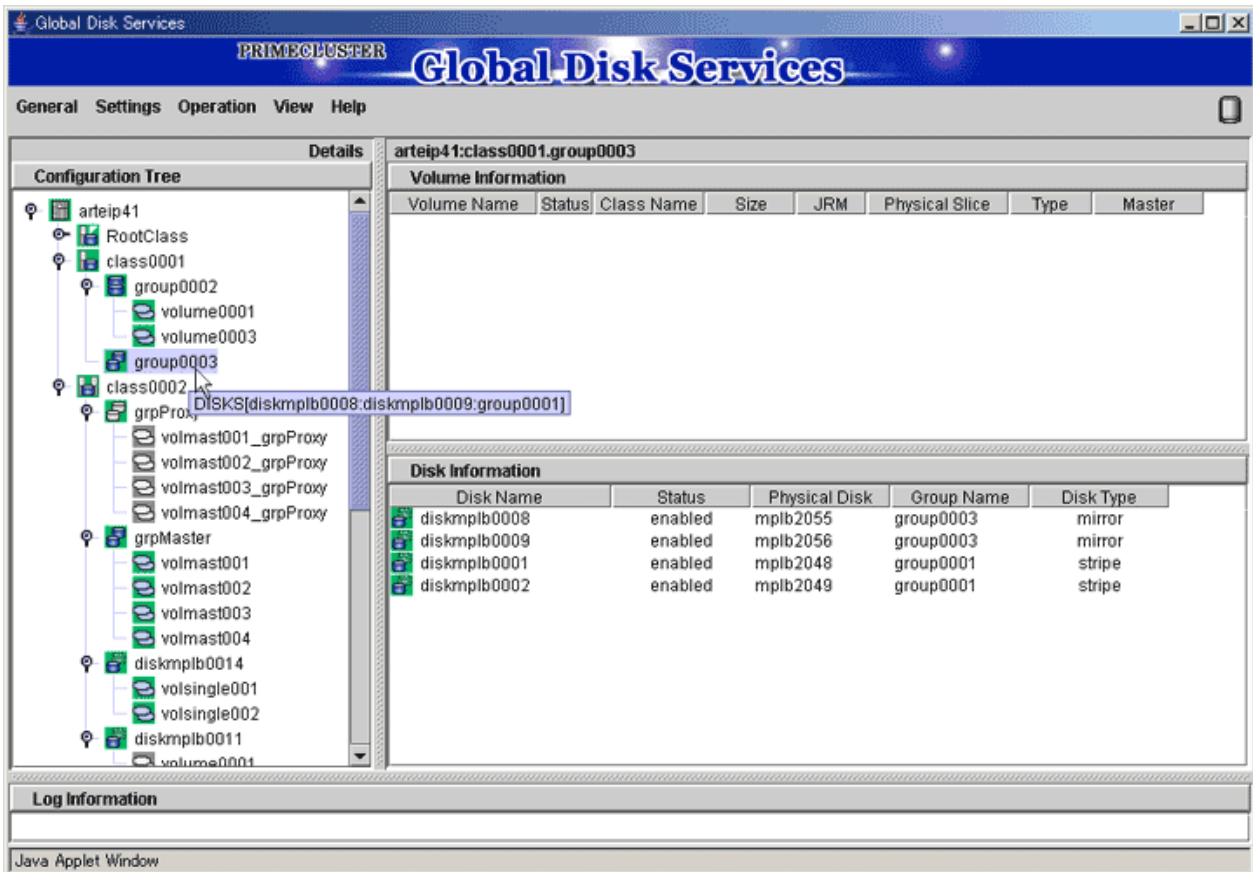
Figure 5.48 Main Screen (for SDX Objects of a Class)



Object configuration within a group

Click a group icon in the GDS configuration tree field, and all volumes and disks at any level within the specified group appear. Additionally, place a mouse pointer on a group icon, and the disks and lower level groups constituting the group appear.

Figure 5.49 Main Screen (for SDX Objects of a Group)



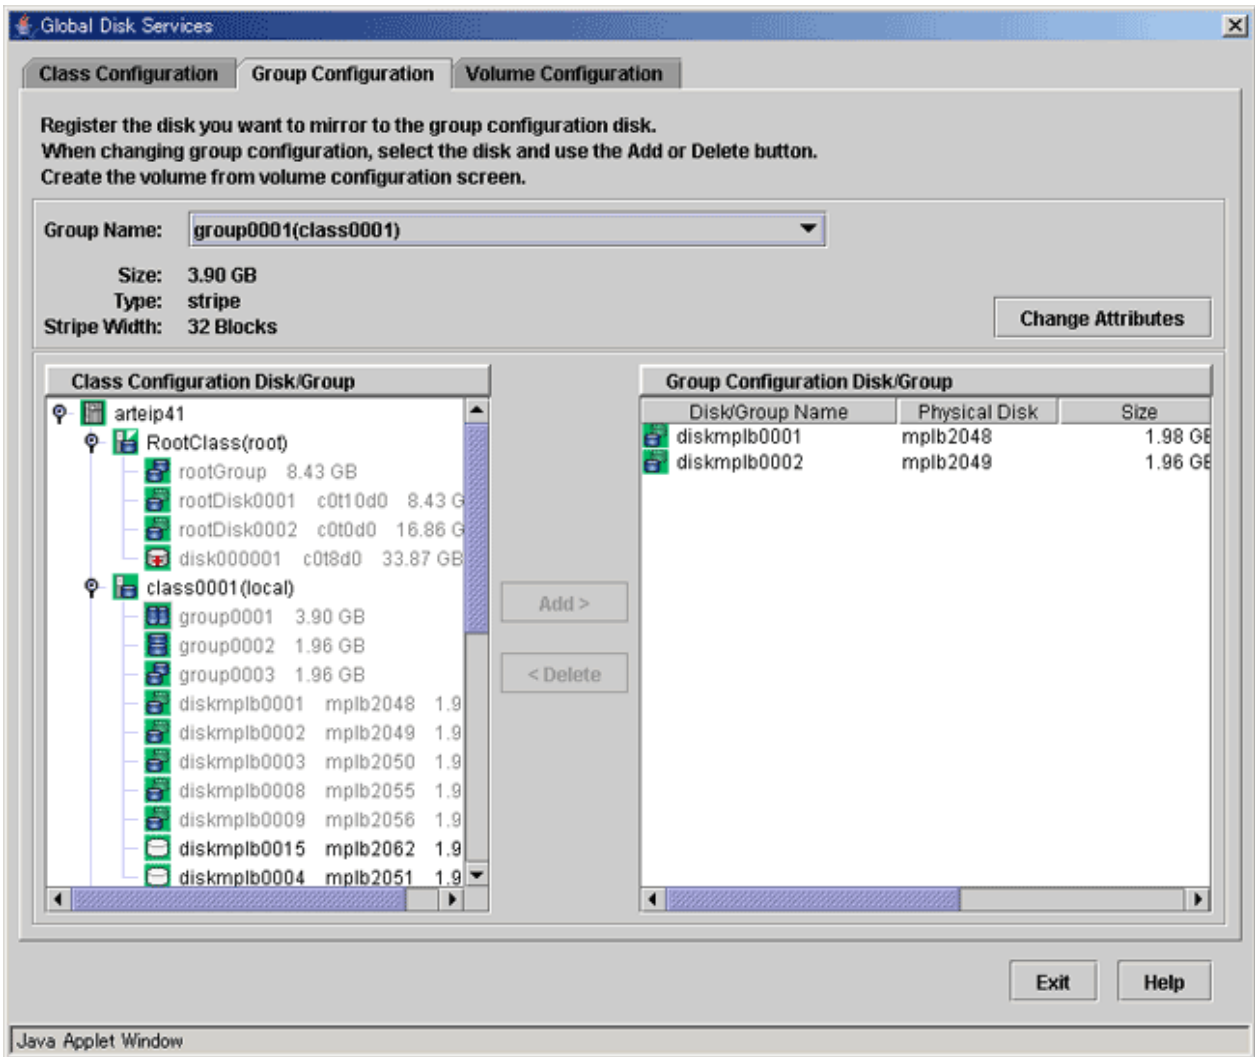
Groups displayed in the GDS configuration tree field are only the highest level groups. You can view the disks and lower level groups constituting their lower level groups in the Group Configuration screen.

Perform the following procedure.

1. Display the Group Configuration screen.
 - Select [Group Configuration] in the [Settings] menu on the Main screen, and the Group Configuration screen will appear.
2. Select the group you want to view the configuration in the [Group Name] list.
3. In the [Group Configuration Group/Disk] field, view the disks and lower level groups constituting the group.

In this example, group group0003 has disk diskmplb0008, diskmplb0009 and lower level group group0001. In a similar manner, view the configuration of lower group group0001.

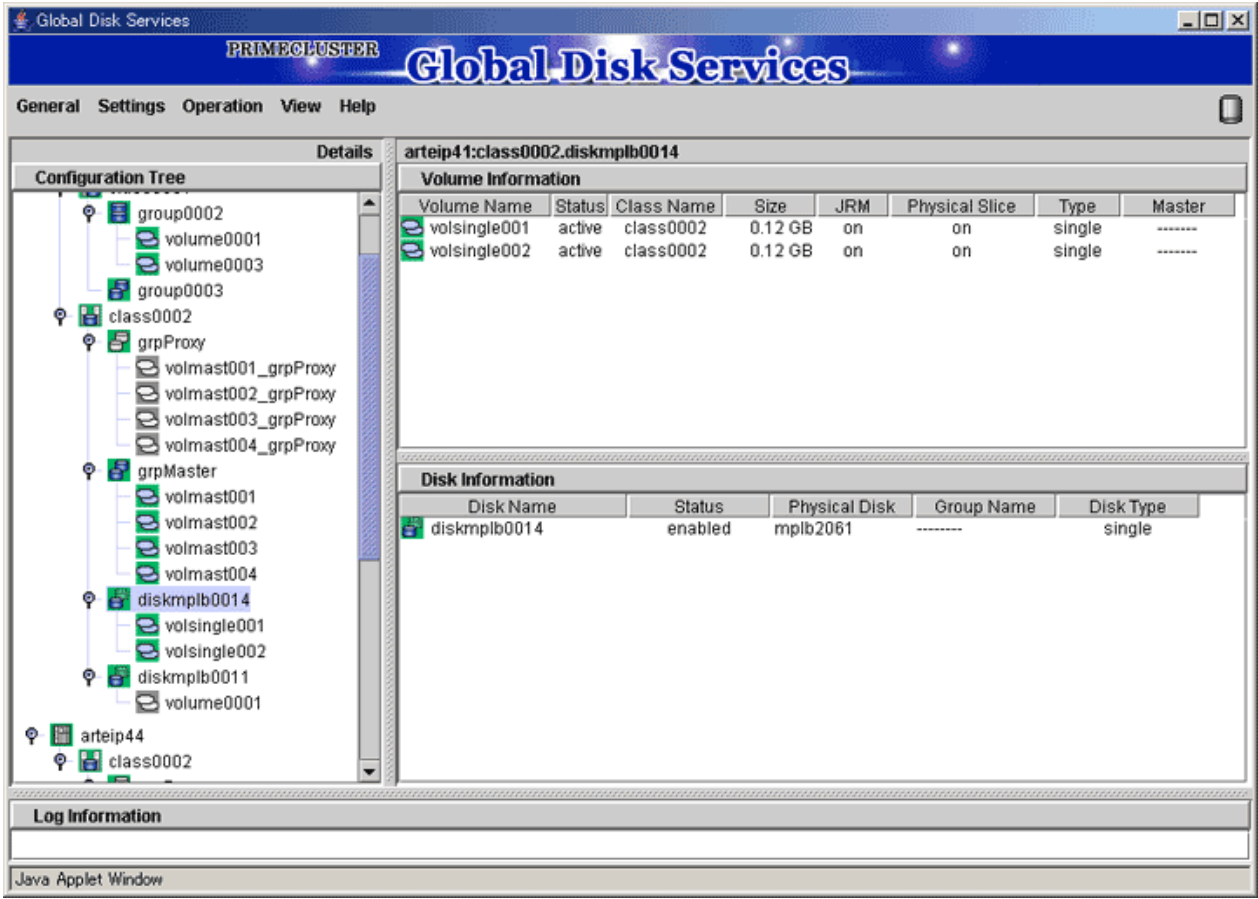
Figure 5.50 Confirming Group Configuration Group/Disk



Object configuration within a single disk

Click a single disk icon in the GDS Configuration Tree Field to view all the volumes within the single disk as well as the single disk.

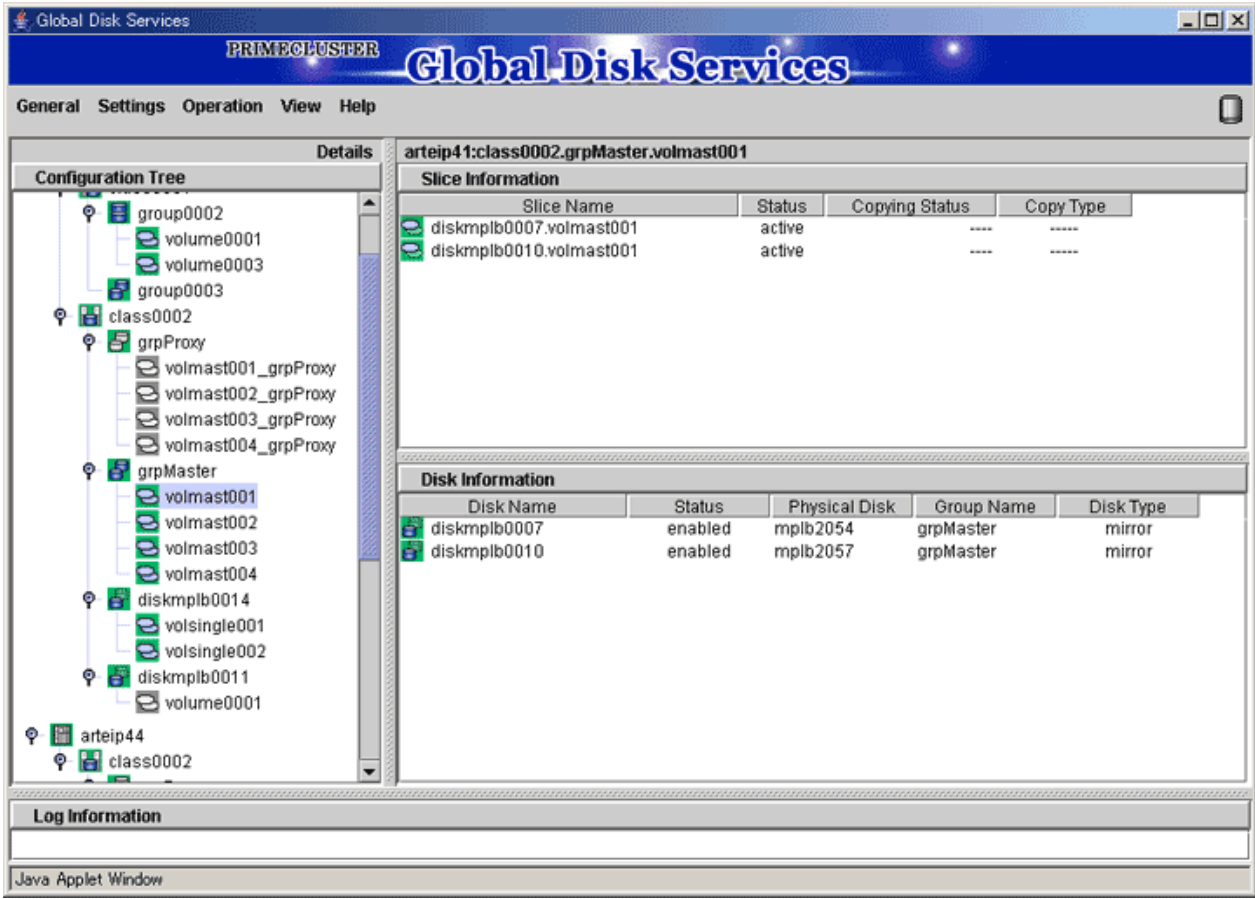
Figure 5.51 Main Screen (for SDX Objects of a Single Disk)



Object configuration within a volume

Click the volume icon in the Configuration Tree field, and all slices and disks within the specified volume appear.

Figure 5.52 Main Screen (for SDX Objects of a Volume)



5.3.1.2 Viewing Proxy Object Configurations

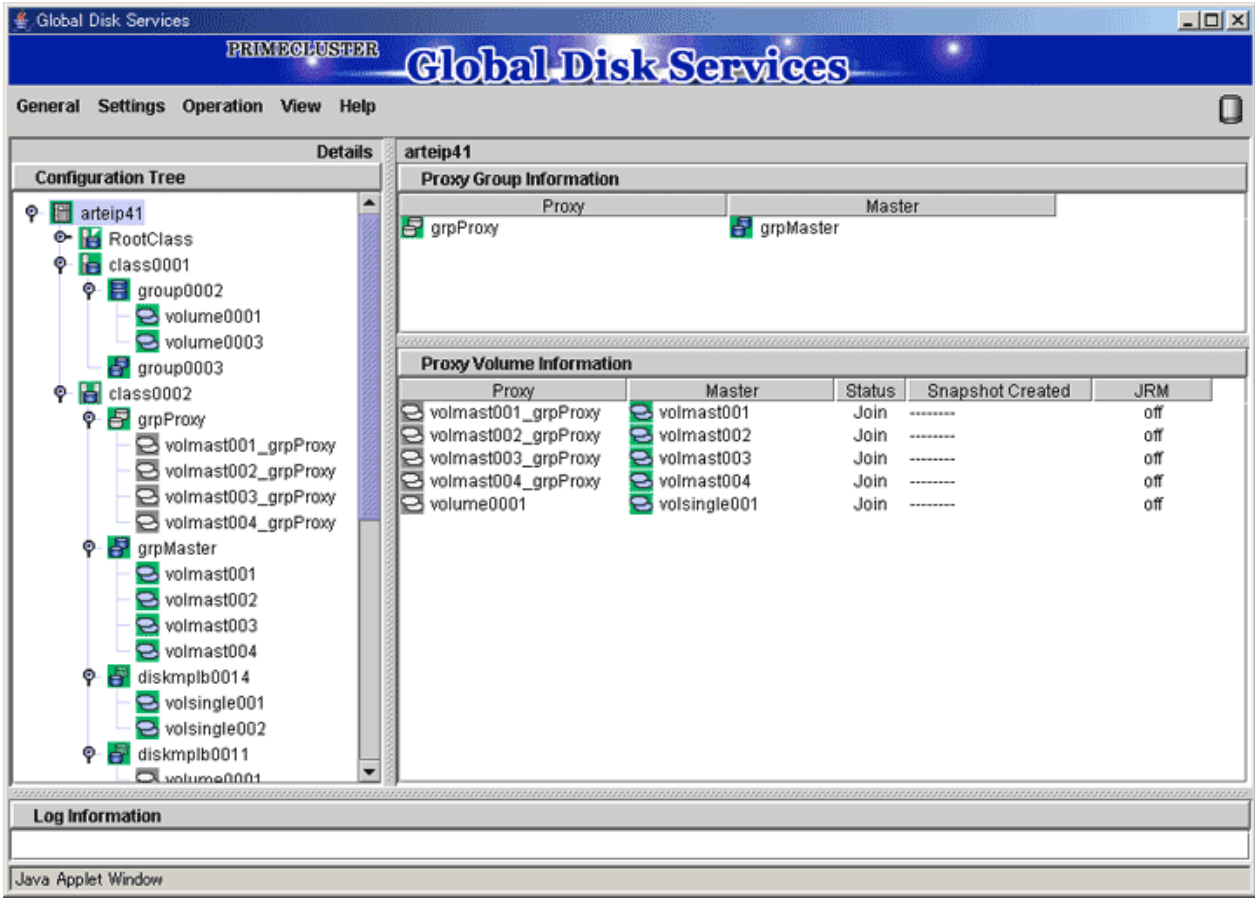
Select [Proxy Object] in the [View]:[Details] menu to view object configurations according to the following units.

- Proxy object configuration within a node
- Proxy object configuration within a class
- Proxy object configuration related to a group
- Proxy object configuration within a single disk
- Proxy object configuration related to a volume

Proxy object configuration within a node

Click the node icon in the Configuration Tree field, and all proxy groups and proxy volumes within the specified node appear.

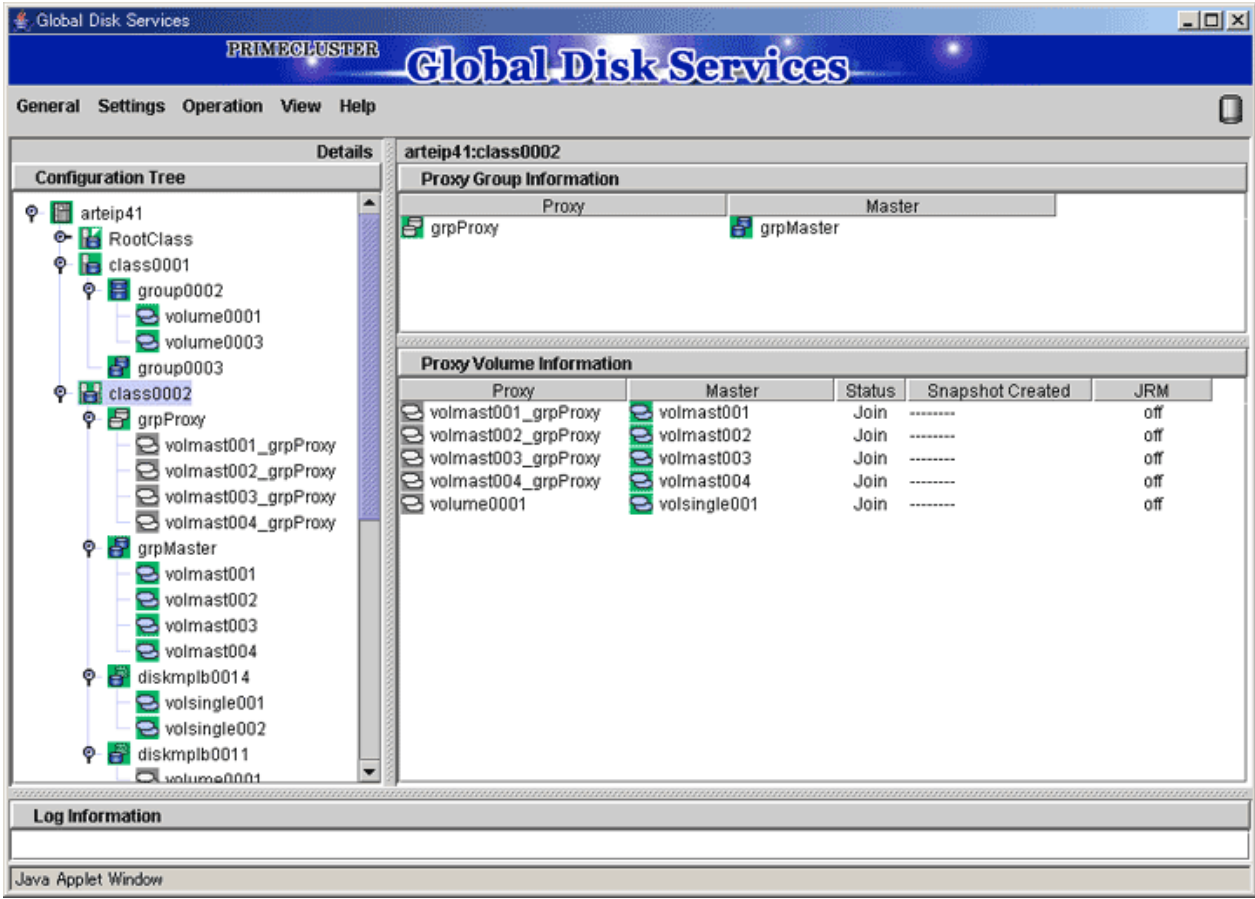
Figure 5.53 Main Screen (for Proxy Objects of a Node)



Proxy object configuration within a class

Click the class icon in the Configuration Tree field, and all proxy groups and proxy volumes within the class appear.

Figure 5.54 Main Screen (for Proxy Objects of a Class)

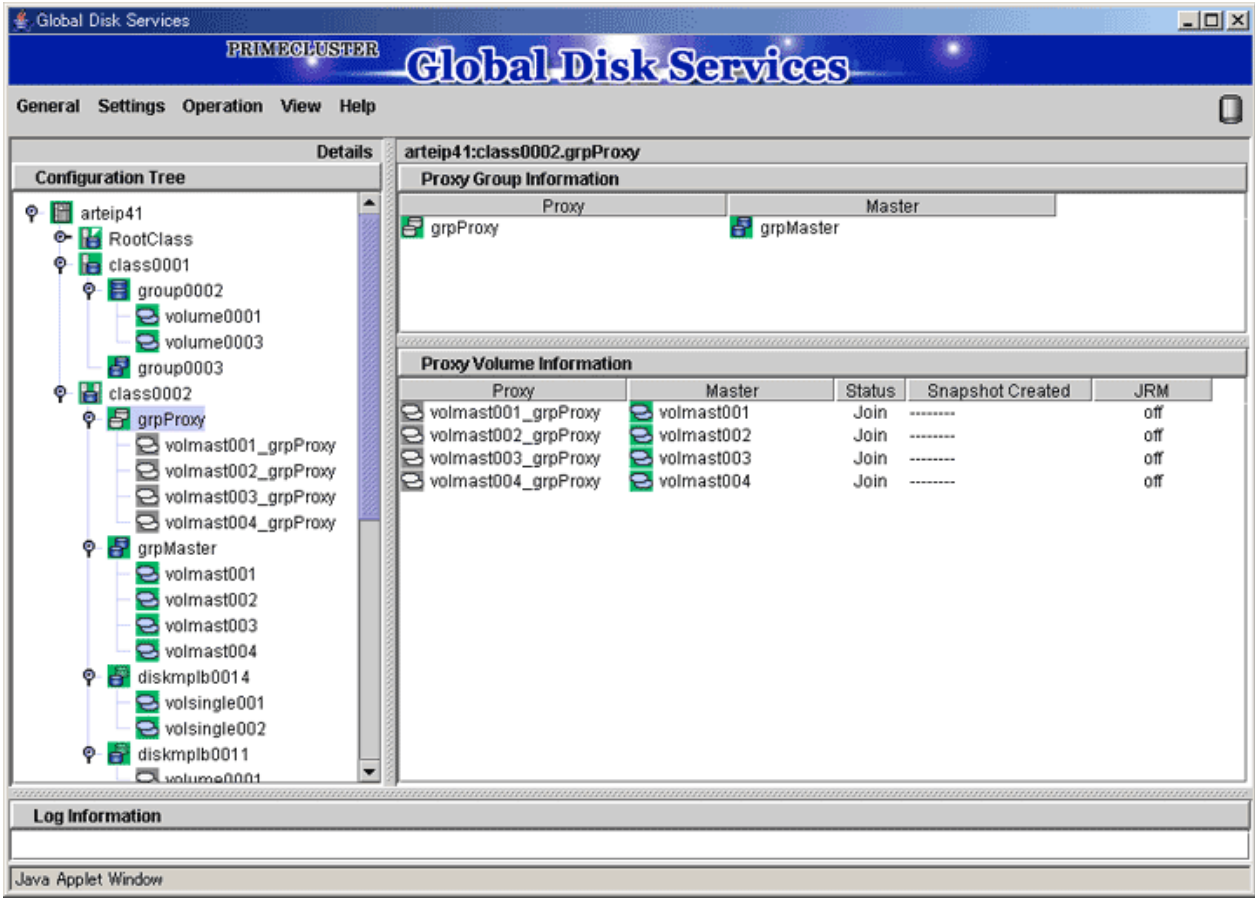


Proxy object configuration related to a group

Click a group icon in the GDS Configuration Tree Field to view the following configuration information.

- All the master groups or proxy groups related to that group
- All the proxy volumes within that group and within proxy groups related to it

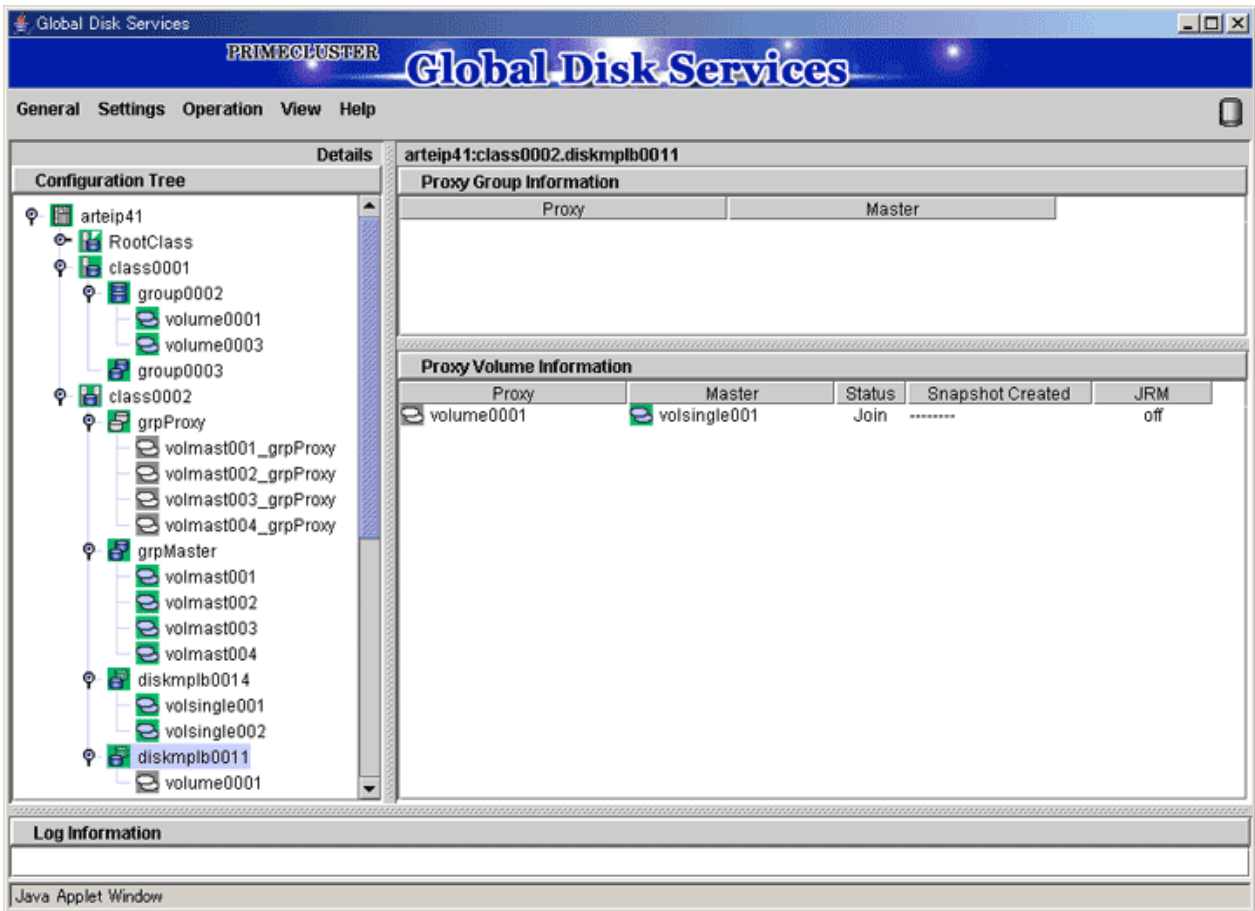
Figure 5.55 Main Screen (for Proxy Objects of a Group)



Proxy object configuration within a single disk

Click a single disk icon in the GDS Configuration Tree Field to view all the proxy volumes within that single disk.

Figure 5.56 Main Screen (for Proxy Objects of a Single Disk)

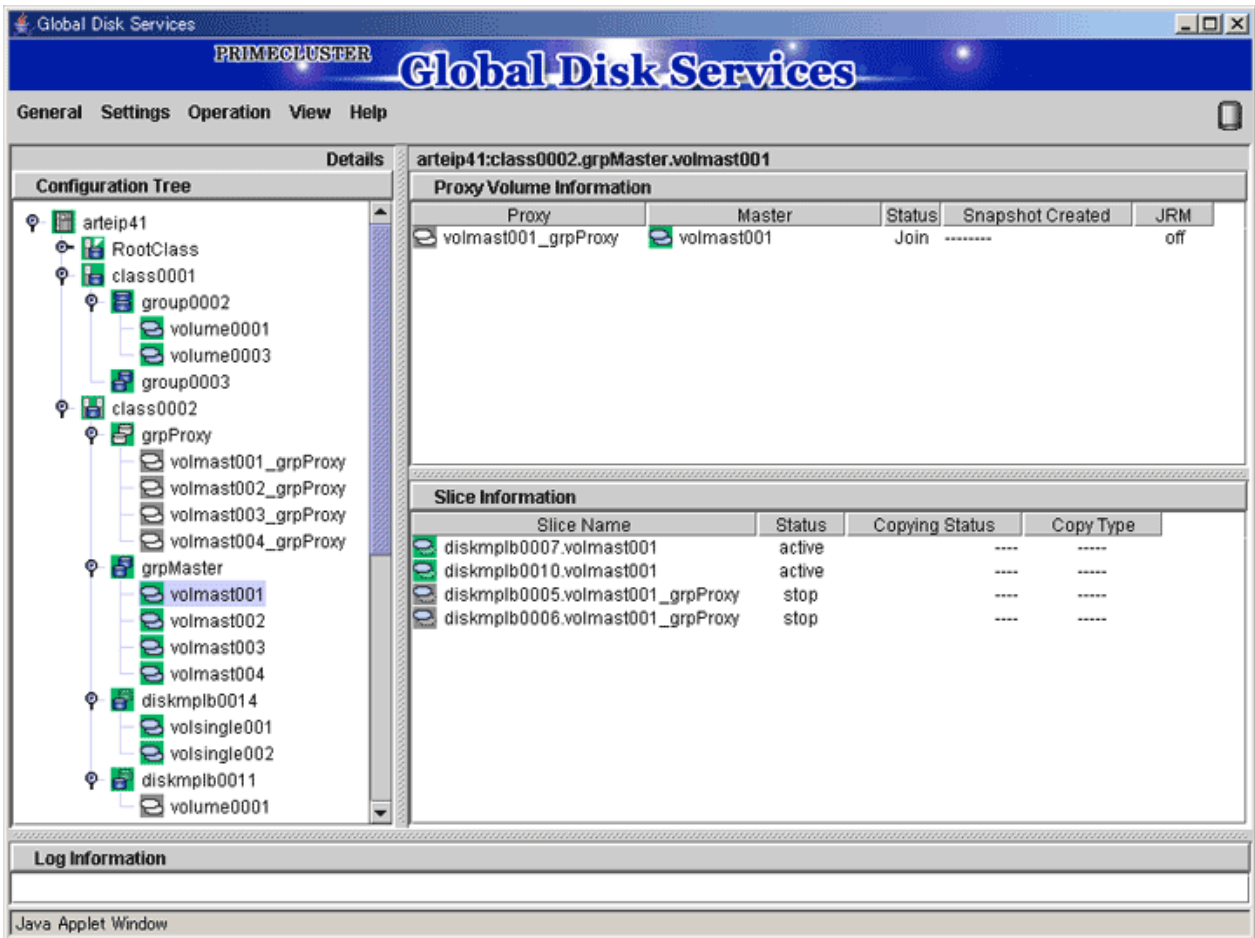


Proxy object configuration related to a volume

Click a volume icon in the GDS Configuration Tree Field to view the following configuration information.

- All the master volumes or proxy volumes related to that volume
- All the slices within that volume or within volumes related to it

Figure 5.57 Main Screen (for Proxy Objects of a Volume)



5.3.1.3 Monitoring Object Status

You can monitor the object status from the Main screen.

Object status will be updated at intervals specified in [Change Monitoring Intervals] in the [General] menu. You may also use [Update] in the [View] menu to update the status immediately.

When an abnormality is detected with objects, a warning lamp (yellow/red) will flash.

Clicking the flashing lamp will change the lamp to a lit-up lamp.



See

For details on warning lamps, see "4.1 Screen Configuration."

When the object status changes such as by detection of the failed disk, the icon color and description in the status field will change as well.



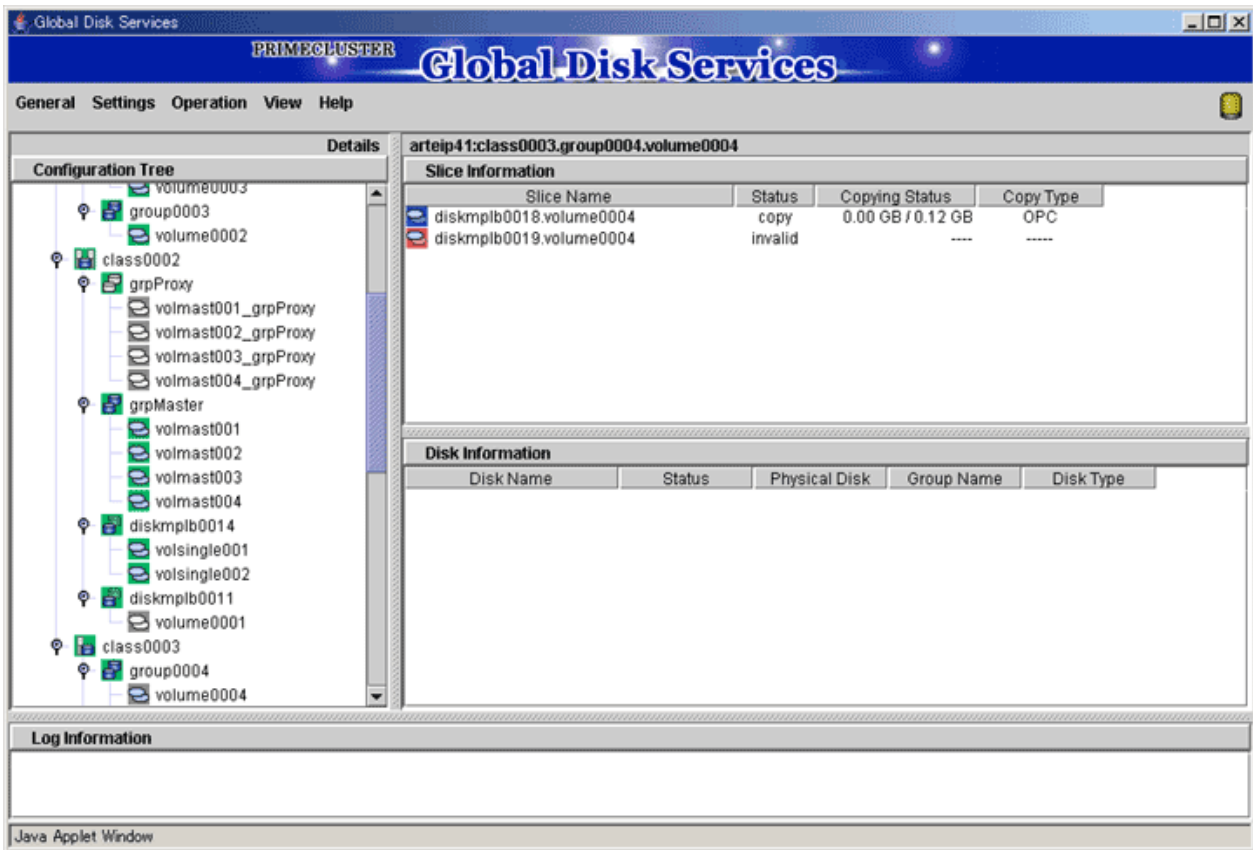
See

For details on icons, see "4.3 Icon Types and Object Status."

If a disk unit fails during operation, an icon of the disk in which an error is detected will turn red.

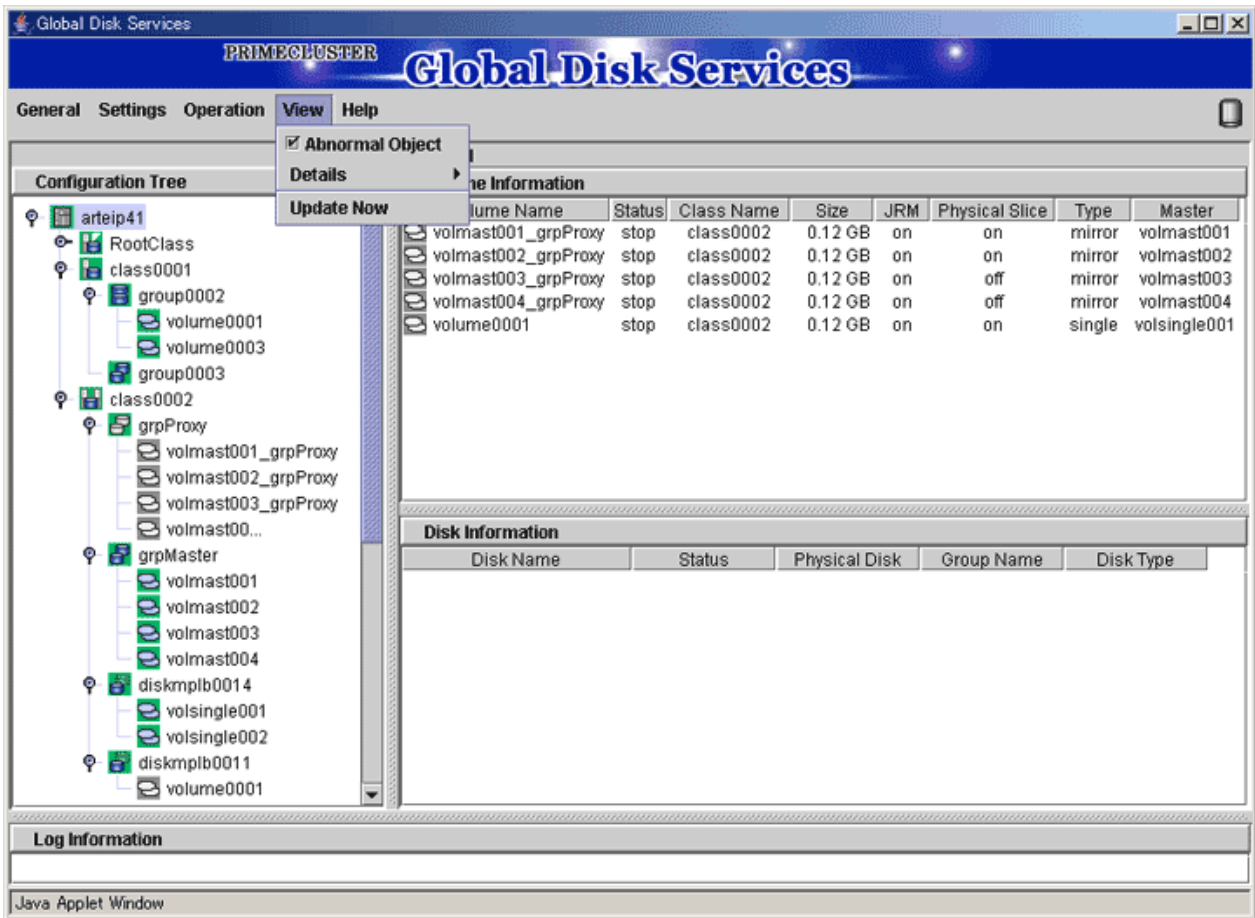
Perform recovery work, following procedures described in "5.3.4 Disk Swap."

Figure 5.58 Object Status Monitoring in Main Screen



Clicking [Abnormal Object] in the [View] menu will only display objects with abnormalities, making it easy to resolve the problem even when a number of disks are connected.

Figure 5.59 Main Screen when [View]: [Abnormal Object] is selected



Note

Abnormality Detected with GDS Management

What GDS Management displays as abnormal is limited to objects detected by GDS.

Therefore, even if the disk unit may have a hardware abnormality, it will be displayed as normal until it is accessed, and the abnormality is detected.

5.3.1.4 Viewing Object Statuses

In the Main Screen, the following two methods are provided for displaying a detailed description of the state of an object and the help to restore the object if it is faulty.

- [Check Status] in the popup menu
- [Check Status] in the drop-down menu

Information

Object Status Update

The state of objects is updated at regular intervals specified with [General]:[Change Monitoring Intervals]. To check the latest state, select [View]:[Update Now].

See

If the state of an object changes due to a disk failure and so on, the corresponding icon color and status field will change. For status indicator icons, see "4.3 Icon Types and Object Status."

1. Checking the state of an object

Using one of the following methods, check the state of an object.

- a. [Check Status] in the popup menu

In the Main Screen, right click the desired object and select [Check Status] from the popup menu.

- b. [Check Status] in the drop-down menu

In the Main Screen, click the desired object and select [Operation]:[Check Status].

2. Object Status Window

View the state of the object in the displayed dialog box.

Figure 5.60 Object Status Window



To close the window, click <OK>.

If you see <Help> in the window, click <Help> for a detailed description and the help to restore the object.

5.3.2 Backup

GDS is software that provides a highly reliable system environment at times of failure, allowing you to continue normal service. However, using a mirrored system environment does not guarantee absolute safety. Computer problems could occur due to hardware, software, operation error and poor environment. A reliable method for working around to such problems is creating "backups." To minimize the damage caused by trouble, periodical backup is strongly recommended.

Information

Backing Up and Restoring a System Disk

Among system disk volumes, volumes (such as /opt and /home) other than system volumes (/ , /usr, /var and swap area) can be backed up by following the procedures described in this section. For system volume backup and restore, see "6.1 Backing Up and Restoring a System Disk" and "6.2 Backing Up and Restoring a System Disk through an Alternative Boot Environment."

5.3.2.1 Backup (by Slice Detachment)

This section describes the procedure for creating backup of a mirror volume making use of a snapshot by detaching a slice. This method requires the following "slice operations."

- Detach Slice
- Attach Slice



See

- For snapshot by slice detachment, see "[1.3.8 Snapshots by Slice Detachment](#)."
- For the operation flow, see "[5.1.4.1 Backup \(by Slice Detachment\)](#)."



Note

Use Conditions on Snapshot by Slice Detachment

Slices can be detached only from mirror volumes with physical slices. In other words, if disks are not connected directly to mirror groups, creation of snapshots by slice detachment is impossible.



Information

Restore

When data is restored back to a volume using data backed up in this procedure, data is restored for the access path to that volume, `/dev/sfdsk/class_name/rdisk/volume_name`. For details, see "[5.1.5 Restore](#)."

Slice Detachment

In order to create a backup, you must temporarily exclude one of the slices from volume, and make it accessible as a separate volume.

The procedures are explained below.

1. Displaying the volume status including the slice

In the Main screen, display the volume containing the slice for which you want to create a backup. Click the icon and [Slice Information] appears.

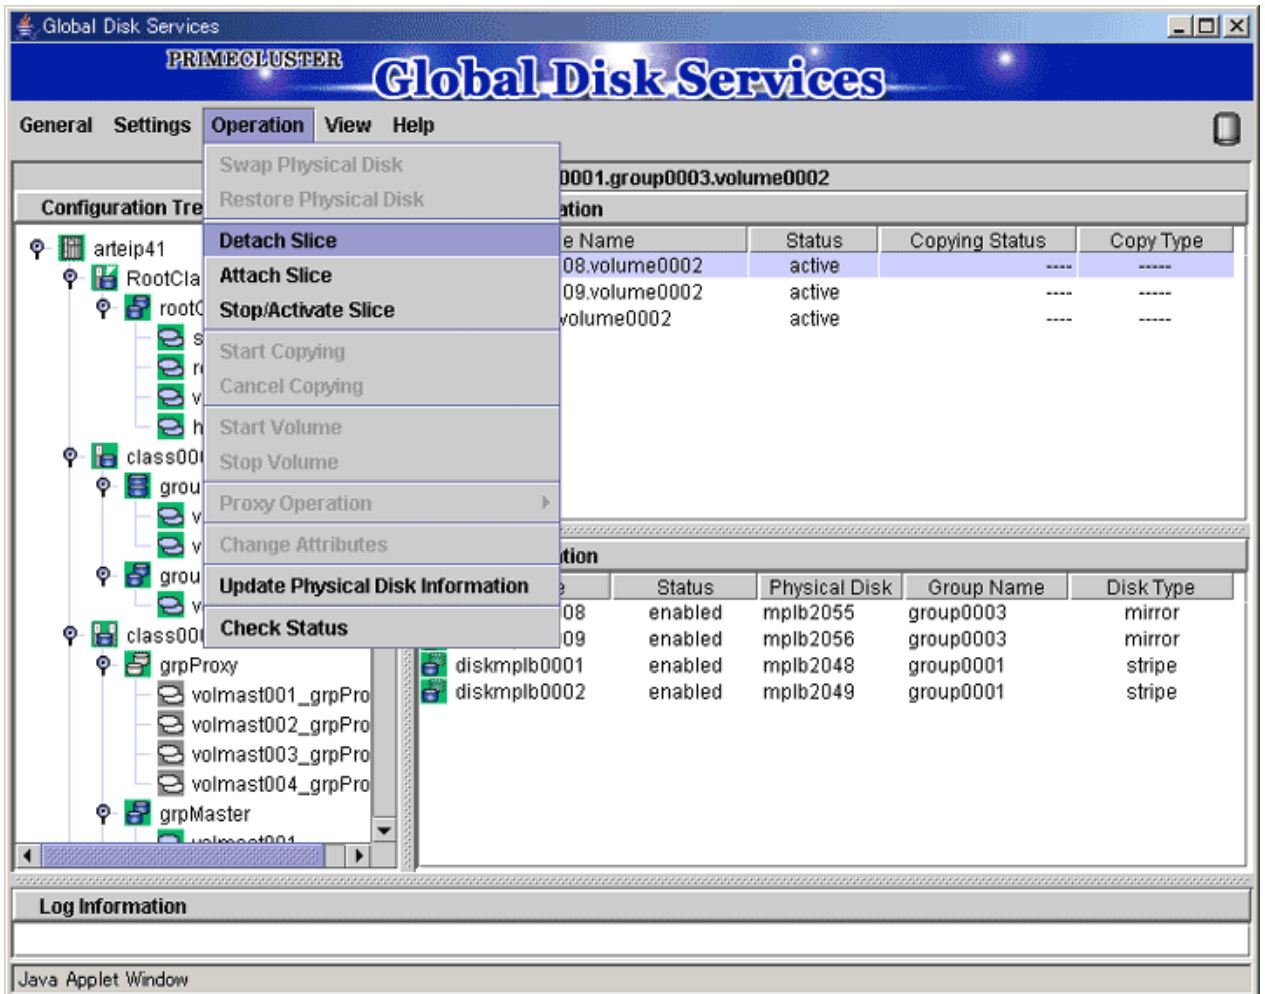
2. Selecting the slice to detach

In the Slice List, select the slice you want to detach by clicking its icon.

3. Selecting [Detach Slice] menu

In the Main screen [Operation] menu, select [Detach Slice].


Figure 5.61 Detach Slice



4. Setting the environment of the detaching slice

We will set the environment of the detaching slice.

Figure 5.62 Setting the environment of the detaching slice



a. Access Mode

Set the access mode of the detached slice.

The initial value is "Read and write possible."

When you specify "Read only possible," the detached mirror slice will be available for read only. Opening a read-only slice in write mode will result in an error.

Click <OK> after setting the environment. If you click <Cancel>, the detaching of the slice will be cancelled.




Slice status available for [Detach Slice]

You can only perform [Detach Slice] operation to slices that are in either "active" or "stop" status.

The confirmation screen shown below appears.

To continue the process, click <Yes>. Clicking <No> cancels the slice detachment process.

Figure 5.63 Confirming Slice Detach




5. Backup Using Access Path

Click <Yes>, and message notifying the completion of detach slice appears.

Use the access path specified in the message to proceed with backup.

Figure 5.64 Notifying Detach Slice Complete



Attach Slice

After the backup is complete, the slice that was temporarily detached will be attached to the volume again.

If the volume is activated, synchronization copying will begin.

The procedures are explained below.

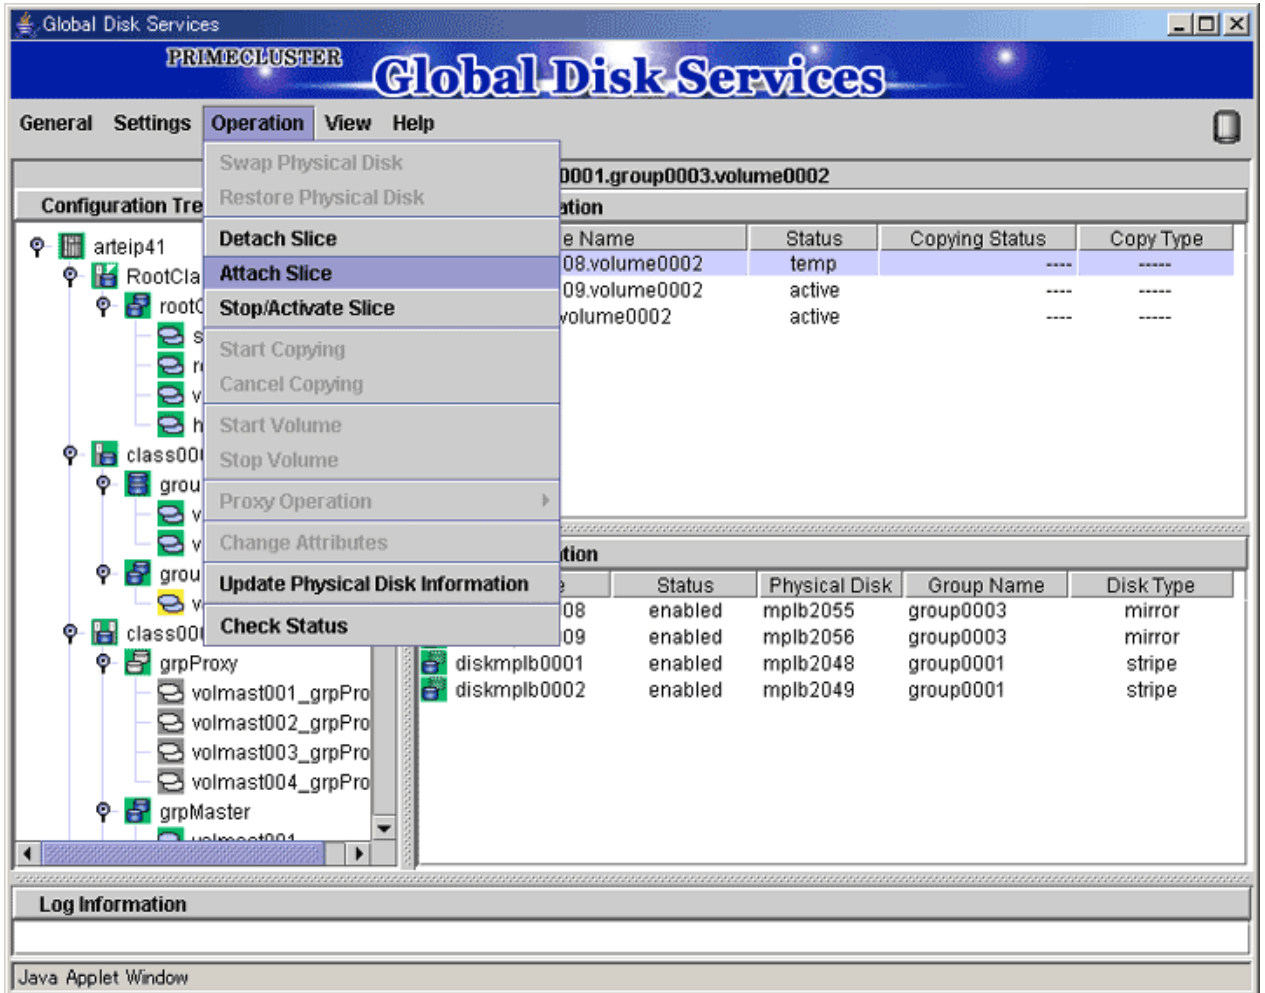
1. **Selecting the slice to attach**

In the Slice Information field, select the mirror slice you want to attach by clicking its icon.

2. **Selecting [Attach Slice] menu**

In the Main screen [Operation] menu, select [Attach Slice].


Figure 5.65 Attach Slice



The confirmation screen shown below appears.

To continue the process, click <Yes>. Clicking <No> cancels the slice attachment process.


Figure 5.66 Confirming "Attach Slice"



3. Notifying the completion of Attach Slice

Click <OK>, and a message notifying the completion of Attach Slice appears.

Figure 5.67 Notifying Attach Slice Completion



Reliable Backup Procedures

Although the procedure above allows you to use the volume as it is after preparing backup, reliability is affected since one of the slices is excluded.

To ensure reliability, temporarily add a backup disk and perform synchronization copying. After its completion, exclude the slice.

For safe and reliable backup, follow the procedures below.

1. Register backup disk with class.
2. Connect backup disk to group.
3. After completing synchronization copying, stop service.
4. Detach slice.
5. Resume service.
6. Perform backup using access path for backup.
7. Attach slice.
8. Disconnect backup disk from group.
9. Remove backup disk from class.

Completing procedures 1. and 2. in advance will save the waiting time for synchronization copying in procedure 3., therefore reducing the time required for backup.

Note

Perform [Detach Slice] after Stopping Services

To ensure integrity of backup data, always stop service before excluding the slice.

You may resume service once [Detach Slice] is complete. You do not have to suspend service during the actual backup process.

When excluding the slice without stopping service, run host programs such as fsck (in case of file system) as necessary.

Stop/Activate Slice

Stop Slice

In order to protect data of the slice which has been detached to prepare for backup, a slice with "temp" status will temporarily become inaccessible.


1. Selecting the slice to stop

In the Slice Information Field, select the slice you want to stop by clicking the "temp" status slice icon.

2. Selecting [Stop/Activate Slice] menu

In the Main screen [Operation] menu, select [Stop/Activate Slice].

Figure 5.68 Stop Slice



To continue the process, click <Yes>. Clicking <No> cancels the slice stop process.

Activate Slice

Reactivate the slice that has become inaccessible ("temp-stop" status) as a result of performing [Stop Slice] operation or switching of the nodes, and make it accessible.

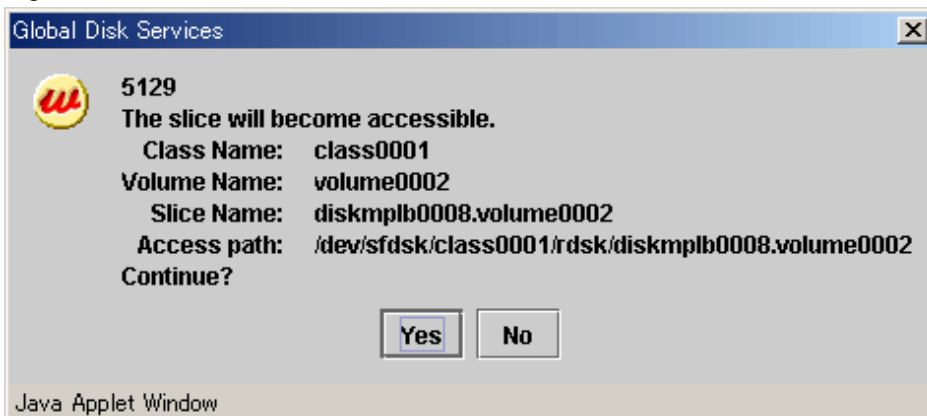
1. Selecting the slice to activate

In the Slice Information Field, select the slice you want to reactivate by clicking the "temp-stop" status slice icon.

2. Selecting [Stop/Activate Slice] menu

In the Main screen [Operation] menu, select [Stop/Activate Slice].

Figure 5.69 Activate Slice



To continue the process, click <Yes>. Clicking <No> cancels the slice activation process.

5.3.2.2 Backup (by Synchronization)

This subsection describes the procedures for backing up volumes through use of snapshots of GDS Snapshot by synchronization. This method requires the following "Proxy Operations."

- Join
- Part
- Rejoin
- Break



See

- For snapshot by synchronization, see "[1.5.1 Snapshot by Synchronization](#)."
- For the operation flow, see "[5.1.4.2 Backup \(by Synchronization\)](#)."



Point

Keep Proxies Parted If Possible

If proxies are kept parted, they can be used for master data restoration. Therefore, it is recommended to keep proxies parted if possible. For the restore procedures see "[5.3.3 Restore](#)."



Note

"Part" Proxies after Stopping Services

To ensure integrity of backup data, always stop services before executing "Part." You may resume the services once "Part" is complete. It is unnecessary to stop services when backing up data to tape and so on.

For details, see "[A.2.22 Ensuring Consistency of Snapshot Data](#)."



Note

Snapshot by Synchronization Use Conditions

See the following sections for points of concern.

- "[A.1.8 Proxy Configuration Preconditions](#)"
- "[A.1.9 Number of Proxy Volumes](#)"
- "[A.1.10 Proxy Volume Size](#)"
- "[A.1.11 Proxy Group Size](#)"



Information

Use Conditions for Copy Functions of Disk Units

See "[A.2.18 Using the Advanced Copy Function in a Proxy Configuration](#)" and "[A.2.21 Using EMC TimeFinder or EMC SRDF in a Proxy Configuration](#)."

Join

Join a volume for backup (a proxy volume) to a volume to be backed up (a master volume).

To back up all the volumes within a group simultaneously, join a group for backup (a proxy group) to a group to be backed up (a master group).

For the "Join" procedures see "[5.2.5.1 Join](#)."

Part

Make sure that synchronization copying from a master to a proxy is complete in the Main screen, and part the proxy from the master.

Follow the procedures below.

1. Selecting a proxy to be parted

Click a master volume icon on the GDS Configuration Tree in the Main screen.

To back up all the master volumes within a master group simultaneously, click the master group icon.

Select [View]:[Details]:[Proxy Object] to view all the proxy objects related to the master object selected on the GDS Configuration Tree, in the Object Information Field.

In this field, click an icon of a proxy volume or a proxy group used for backup.



Note

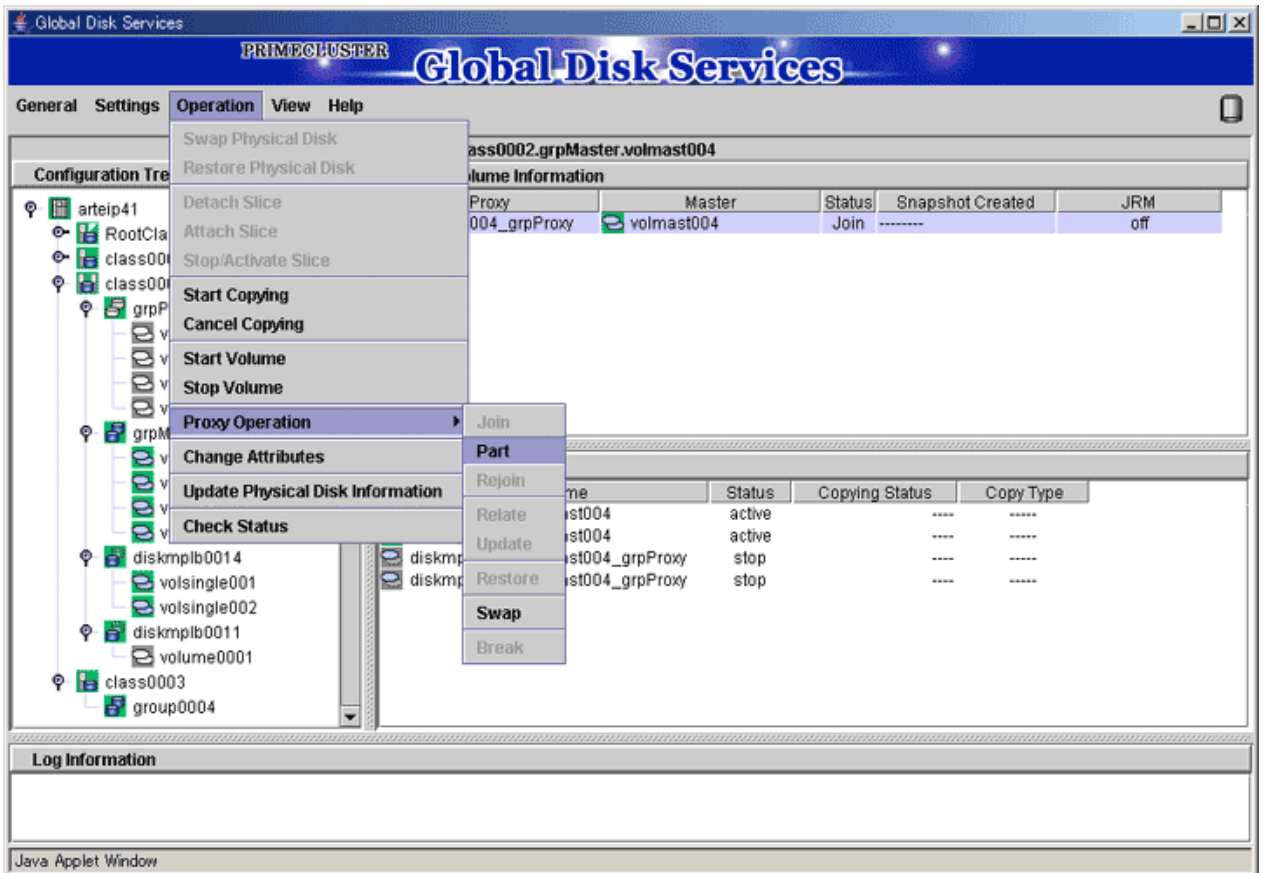
Proxy Objects That Can Be Parted

Proxy volumes can be parted if they are joined and if copy is complete.

2. Selecting the [Part] menu

Select [Operation]:[Proxy Operation]:[Part].


Figure 5.70 Part



3. Setting the environment for parting the proxy

The Part Proxy dialog box appears.

Figure 5.71 Setting the Part Environment



Set the environment for parting the proxy.

a. Instant Snapshot

Specify whether to change the synchronization mode to the OPC mode.

The default value is "No." If synchronization copying from the master to the proxy is incomplete, the part process will fail.

If this option is set to "Yes", instant snapshots are created with the OPC function. Even if synchronization copying from the master to the proxy is in progress, the proxy will be parted and then background copying from the master to the proxy will be executed with the OPC function. If the OPC function is unavailable, the part process will fail.



Instant Snapshot by OPC Use Conditions

See the following sections.

- ["A.2.18 Using the Advanced Copy Function in a Proxy Configuration"](#)
- ["A.2.19 Instant Snapshot by OPC"](#)

b. Just Resynchronization Mechanism

Set the mode of Just Resynchronization Mechanism (JRM) for proxies. The default value is "on."



For details on JRM for proxies, see ["A.2.14 Just Resynchronization Mechanism \(JRM\)."](#)

c. Access Mode

Set the access mode of the parted proxy volume.

The default value is "Read Only." The parted proxy volume will be read-only and an error occurs if it is opened in write mode.


To permit write access to the parted proxy volume, set this option to "Read/Write."

After the settings are complete, click <OK>. Clicking <Cancel> cancels the part process.

4. Information message of the completion

A message window appears informing you that the part process is complete.

Figure 5.72 Information Message of "Part" Completion



Click <OK> to close the information message window.

Back up data through use of the proxy volume.

Rejoin

To re-execute backup, rejoin a parted proxy to a master.

Follow the procedures below.

1. Selecting a proxy to be rejoined

Click an icon of a master volume to be backed up on the GDS Configuration Tree in the Main screen.

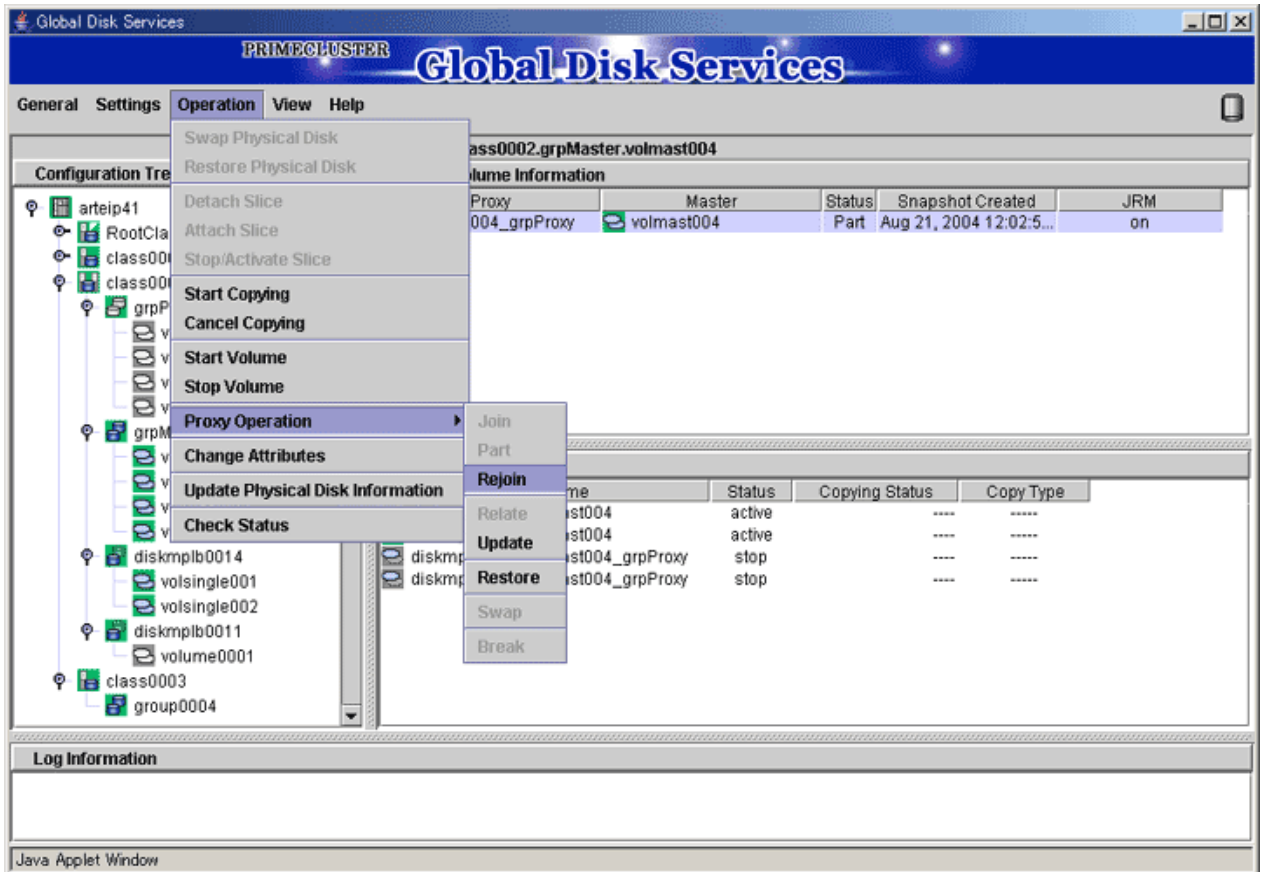
To back up all the master volumes in a master group simultaneously, click the master group icon.

Select [View]:[Details]:[Proxy Object] to view all the proxy objects related to the master object selected on the GDS Configuration Tree, in the Object Information Field. In this field, click an icon of a proxy volume or proxy group used for backup.

2. Selecting the [Rejoin] menu

Select [Operation]:[Proxy Operation]:[Rejoin] in the Main screen.


Figure 5.73 Rejoin



3. Confirming

A confirmation dialog box appears asking you whether to rejoin the proxy.

Figure 5.74 Confirming "Rejoin"




To continue the process, click <Yes>. Clicking <No> cancels the proxy rejoin process.

4. Information message of the completion

A message window appears informing you that the rejoin process is complete.

Figure 5.75 Information Message of "Rejoin" Completion



Click <OK> to close the information message window.

Break

If no more backup is necessary, break the relationship between the master and the proxy.

For "Break Proxy" procedures see "[5.5.6 Breaking a Proxy](#)."

5.3.2.3 Backup (by OPC)

This subsection describes the procedures for backing up volumes through use of snapshots of GDS Snapshot by OPC. This method requires the following "Proxy Operations."

- Relate
- Update
- Break



See

- For snapshot by OPC, see "[1.5.3 Instant Snapshot by OPC](#)."
- For the operation flow, see "[5.1.4.3 Backup \(by OPC\)](#)."



Note

"Update" Proxies after Stopping Services

To ensure integrity of backup data, always stop services before executing "Update." You may resume the services once "Update" is complete. It is unnecessary to stop services when backing up data to tape and so on.

For details, see "[A.2.22 Ensuring Consistency of Snapshot Data](#)."



Note

Instant Snapshot by OPC Use Conditions

See the following sections for points of concern.

- "[A.1.8 Proxy Configuration Preconditions](#)"
- "[A.1.9 Number of Proxy Volumes](#)"
- "[A.1.10 Proxy Volume Size](#)"

- "A.1.11 Proxy Group Size"
 - "A.2.18 Using the Advanced Copy Function in a Proxy Configuration"
 - "A.2.19 Instant Snapshot by OPC"
-

Relate

Relate a volume for backup (a proxy volume) to a volume to be backed up (a master volume).

To back up all the volumes within a group simultaneously, relate a group for backup (a proxy group) to a group to be backed up (a master group).

For the "Relate Proxy" procedures see "[5.2.5.2 Relate](#)."

Update

Copy (overwrite) data from a master to a proxy with the OPC function.

Follow the procedures below.

1. Selecting a proxy to be updated

Click an icon of a master volume to be backed up on the GDS Configuration Tree in the Main screen.

To back up all the master volumes in a master group, click the master group icon.

Select [View]:[Details]:[Proxy Object] to view all the proxy objects related to the master object selected on the GDS Configuration Tree, in the Object Information Field. In this field, click an icon of a proxy volume (or a proxy group) used for backup.



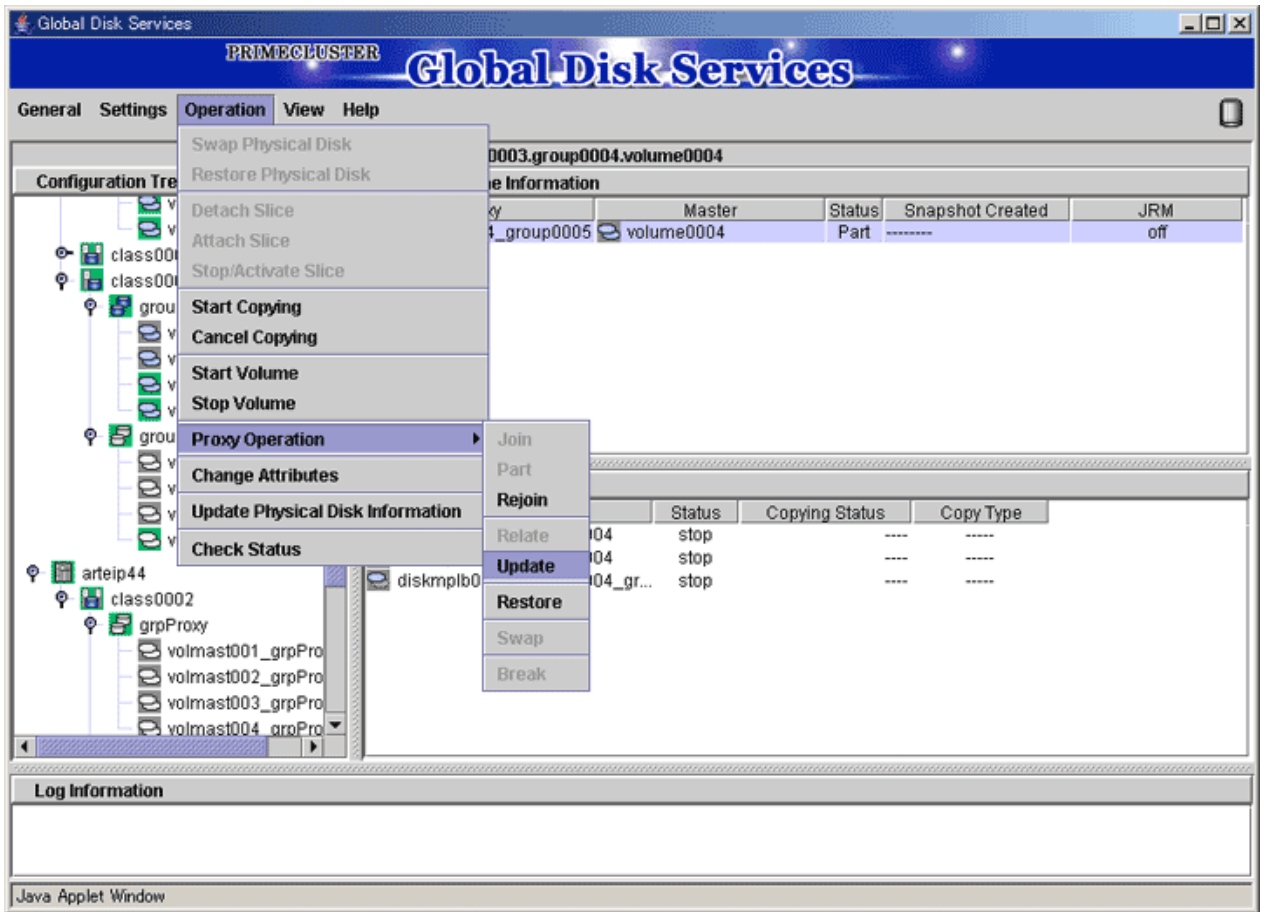
Proxy Objects That Can Be Updated

Proxy volumes can be updated if they are parted and in "stop" status.

2. Selecting the [Update] menu

Select [Operation]:[Proxy Operation]:[Update] in the Main screen.

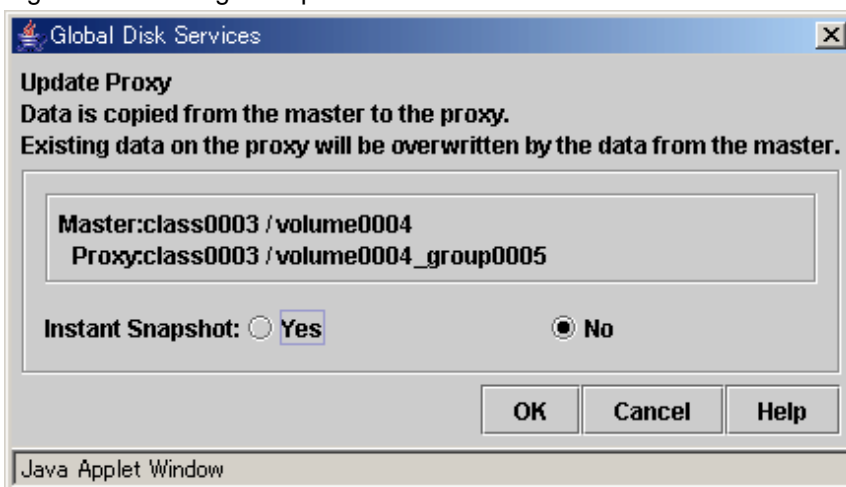
Figure 5.76 Update



3. Setting the environment for updating the proxy

The Update Proxy screen appears.

Figure 5.77 Setting the Update Environment



Set the environment for updating the proxy.

a. Instant Snapshot

Specify whether to apply instant snapshot.

The default value is "No." Wait until copying from the master volume to the proxy volume is complete and start the proxy volume.


To start the proxy volume immediately without waiting until copying from the master volume to the proxy volume is complete, set this option to "Yes."

After the settings are complete, click <OK>. Clicking <Cancel> cancels the update process.

4. Information message of the completion


A message window appears informing you that the update process is complete.

Figure 5.78 Information Message of "Update" Completion (No to Instant Snapshot)



If "No" to [Instant Snapshot] was selected when setting the update environment in step 3., check the copy status in the Main screen, and after the copy is complete, start the proxy volume and execute backup.

Figure 5.79 Information Message of "Update" Completion (Yes to Instant Snapshot)



If "Yes" to [Instant Snapshot] was selected when setting the update environment in step 3., you may start the proxy volume and execute backup immediately without waiting until copying is complete.

Break

If no more backup is necessary, break the relationship between the master and the proxy.

For "Break Proxy" procedures see "[5.5.6 Breaking a Proxy](#)."

5.3.3 Restore

This subsection describes the procedures for restoring volumes through use of GDS Snapshot proxy volumes. This method requires the following "Proxy Operations."

- Restore

Stop services using a volume to be restored (a master volume) to stop the master volume and then perform the following procedures.



See

- For restore with proxy volumes, "[1.5.4 Instant Restore](#)."
- For the operation flow, see "[5.1.5 Restore](#)."



Note

System Volume Restoration

The system volumes currently running as file systems such as /, /usr, and /var cannot be stopped, and such volumes cannot be restored through this procedure. For the system volume restoration methods, see "[6.1 Backing Up and Restoring a System Disk](#)" or "[6.2 Backing Up and Restoring a System Disk through an Alternative Boot Environment](#)."

Restore

Copy (overwrite) data from a proxy to a master.

Follow the procedures below.

1. Selecting a proxy as a restore copy source

Click an icon of a master volume to be restored on the GDS Configuration Tree in the Main screen.

To restore all the master volumes within a master group simultaneously, click the master group icon.

Select [View]:[Details]:[Proxy Object] to view all the proxy objects related to the master object selected on the GDS Configuration Tree, in the Object Information Field.

In this field, click an icon of a proxy volume (or a proxy group) as a restore copy source.



Point

Backup Generation

Proxy volume data is a replica of master volume data at the moment of snapshot creation. For snapshot creation time, check [Snapshot Created] time in the proxy volume information field.



Note

Proxy Volumes That Can Be Restore Copy Sources

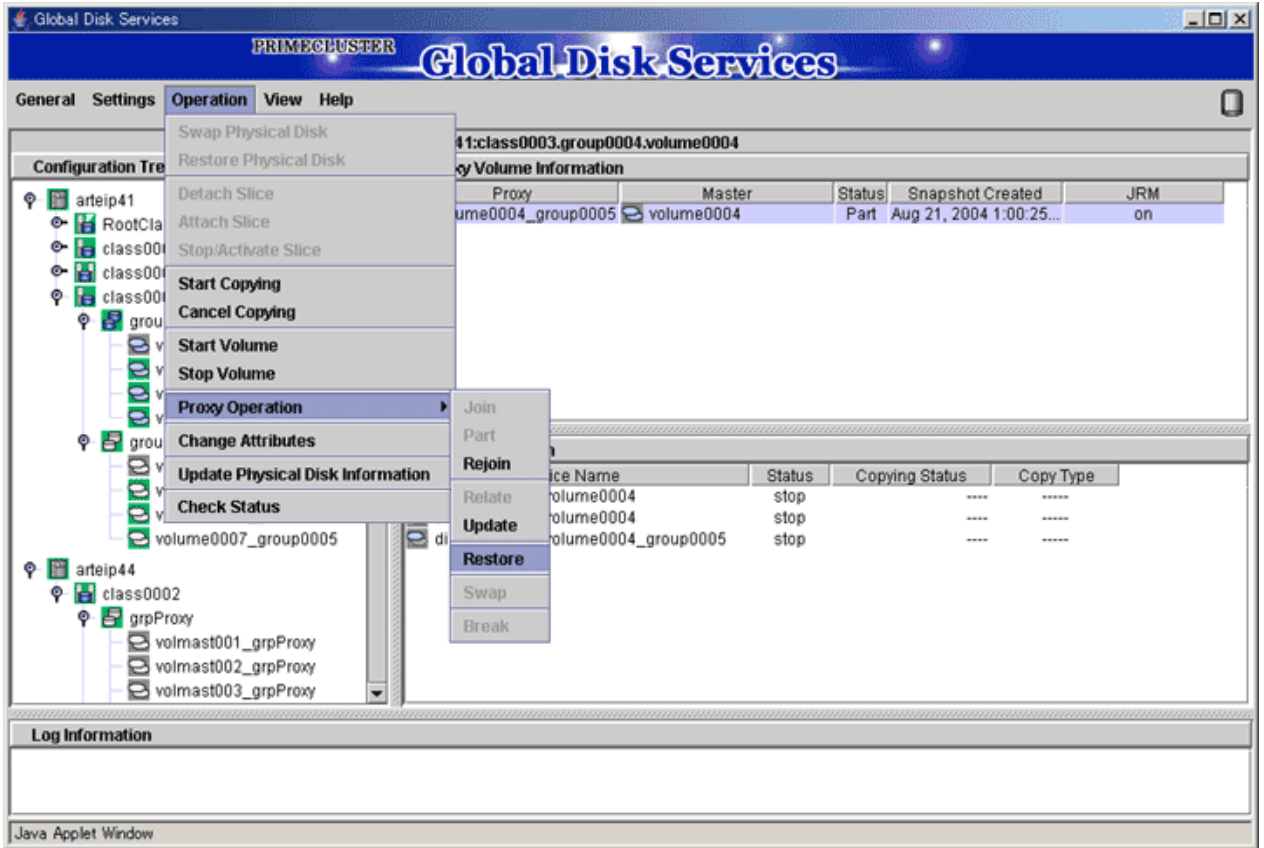
Proxy volumes can be restore copy sources if they are parted and in "active" or "stop" status.

However, when selecting "Yes" to "Rejoin" for setting the restore environment in step 3., copy sources must be proxy volumes that are parted and are in "stop" status.

2. Selecting the [Restore] menu

Select [Operation]:[Proxy Operation]:[Restore] in the Main screen.


Figure 5.80 Restore



3. Setting the environment for restoring the master

The Restore Master screen appears.

Figure 5.81 Setting the Restore Environment



Set the environment for restoring the master.

a. Rejoin

Specify whether to rejoin the master and the proxy.

The default value is "Yes." The master and the proxy will be joined, and after copy is complete they will be synchronized. If the OPC function is unavailable, select "Yes."

To leave the master and the proxy parted and execute restore with the OPC function, set this option to "No."

See

For the OPC function use conditions, see the following sections.

- "A.2.18 Using the Advanced Copy Function in a Proxy Configuration"
- "A.2.19 Instant Snapshot by OPC"

Note

Master Volumes with Mirroring Multiplicity of Two and Higher

Even if the device supports the OPC function, selecting "Yes" to [Rejoin] disables the OPC function. To use the OPC function for restoration, select "No" to [Rejoin]. However, slices other than the OPC copy destinations will be excluded from mirroring and the data statuses will be invalid. To restore the master volume mirroring status, select the master volume and execute [Operation]:[Start Copying] in the Main screen. If not executing [Start Copying], resynchronization copying automatically starts when the master volume starts.

b. Instant Restore

Specify whether to apply instant restore

The default value is "No." Wait until copying from the proxy volume to the master volume is complete and start the master volume.

To start the master volume immediately without waiting until copying from the proxy volume to the master volume is complete, set this option to "Yes."

Note

If "Yes" to [Rejoin] and "Yes" to [Instant Restore] Are Selected


Even if synchronization copying from the proxy to the master is in progress, the master volume can be started and accessed. Note, however, that the master and the proxy are joined and data written to the master is also written to the proxy. To prevent proxy data from being updated, wait until copying is complete and execute "Part" before starting the master volume.

After the settings are complete, click <OK>. Clicking <Cancel> cancels the restore process.

4. Information message of the completion


A message window appears informing you that the restore process is complete.

Figure 5.82 Information Message of "Restore" Completion (Yes to Rejoin/No to Instant Restore)




If "Yes" to [Rejoin] and "No" to [Instant Restore] were selected when setting the restore environment in step 3., the master and the proxy are joined. Wait until resynchronization copying from the proxy to the master is complete and start the master volume.

Figure 5.83 Information Message of "Restore" Completion (Yes to Rejoin/Yes to Instant Restore)




If "Yes" to [Rejoin] and "Yes" to [Instant Restore] were selected when setting the restore environment in step 3., the master and the proxy are joined. You may start the master volume immediately without waiting until resynchronization copying from the proxy to the master is complete.

Figure 5.84 Information Message of "Restore" Completion (No to Rejoin/No to Instant Restore)



If "No" to [Rejoin] and "No" to [Instant Restore] were selected when setting the restore environment in step 3., the master and the proxy are left parted. Wait until OPC copying from the proxy to the master is complete and start the master volume. If the OPC function is unavailable, the restore process fails.

Figure 5.85 Information Message of "Restore" Completion (No to Rejoin/Yes to Instant Restore)



If "No" to [Rejoin] and "Yes" to [Instant Restore] were selected when setting the restore environment in step 3., the master and the proxy are left parted. You may start the master volume immediately without waiting until OPC copying from the proxy to the master is complete.

If the OPC function is unavailable, the restore process fails.

Break

If no more backup is necessary after the restore, break the relationship between the master and the proxy.

For "Break Proxy" procedures see "[5.5.6 Breaking a Proxy.](#)"

5.3.4 Disk Swap

When a disk unit abnormality occurs, contact your local customer support to swap the disk units.

In GDS, regardless of hot swap or not, the following procedures are necessary before and after the disk swap.

- swap physical disk
- restore physical disk

Note

Identifying a Failed Disk Unit

You should pinpoint a hardware error on a disk unit based on, for example, log messages for the disk driver output in the /var/adm/messages file. For details, see "[F.1.11 Disk Unit Error.](#)"

Note

Swapping FC-AL Internal Disks

FC-AL internal disks cannot be swapped through these procedures. For details, see "[A.2.31 Swapping Internal Disks \(FC-AL\).](#)"

Note

Swapping Internal Disks of SPARC T3-1/T3-2/T3-4

Internal disks of SPARC T3-1/T3-2/T3-4 cannot be swapped through these procedures. For details, see "[A.2.32 Swapping Internal Disks \(SPARC T3-1/T3-2/T3-4\).](#)"

Note

Notes on Physical Disk Swap

See "[A.2.16 Swapping Physical Disks.](#)"

Swap Physical Disk

In order to swap the disk units, you must take the physical disk offline.

The procedures are explained below.

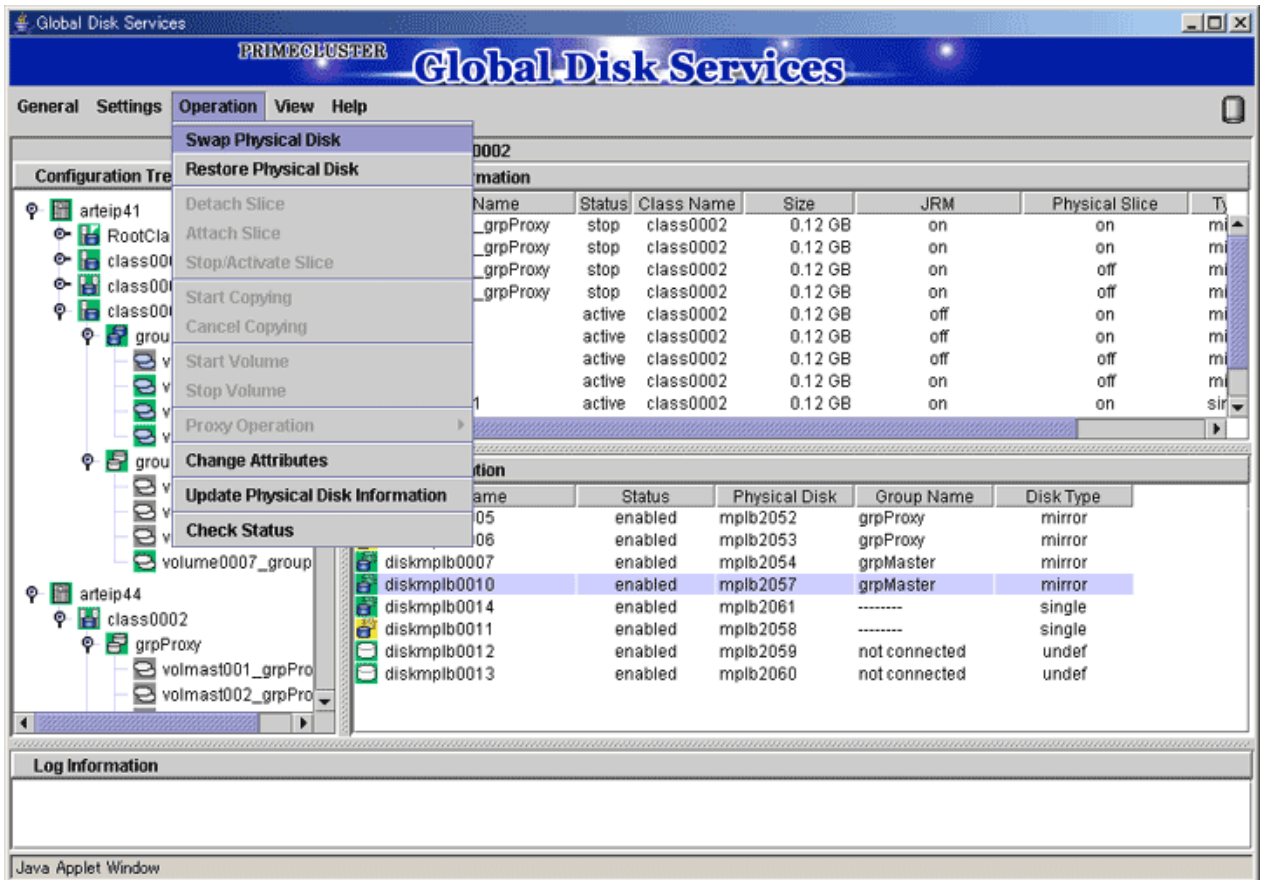
1. Displaying the status of physical disk

In the Main screen, display the physical disk to be swapped. Click the icon and select the physical disk.

2. Selecting [Swap Physical Disk]

In the Main screen [Operation] menu, select [Swap Physical Disk].


Figure 5.86 Swap Physical Disk



The confirmation screen shown below appears.

To continue the process, click <Yes>. Clicking <No> cancels the physical disk swapping process.

Figure 5.87 Confirming Swap Physical Disk



3. Requesting the swap of physical disks

Click <Yes>, and a message notifying the offline status appears.

Click <OK>, and request your local customer support to swap the disk unit.

Figure 5.88 Notifying Offline Completion



Note

To swap with a brand-new disk

If a label is not created for the swapped disk, it is necessary to create a label using the `format(1M)` command before performing "Restore Physical Disk." You can check whether the label exists by the `verify` menu of the `fomat(1M)` command.

Restore Physical Disk

After swapping the disk units, you must put the swapped physical disk back online.

The procedures are explained below.

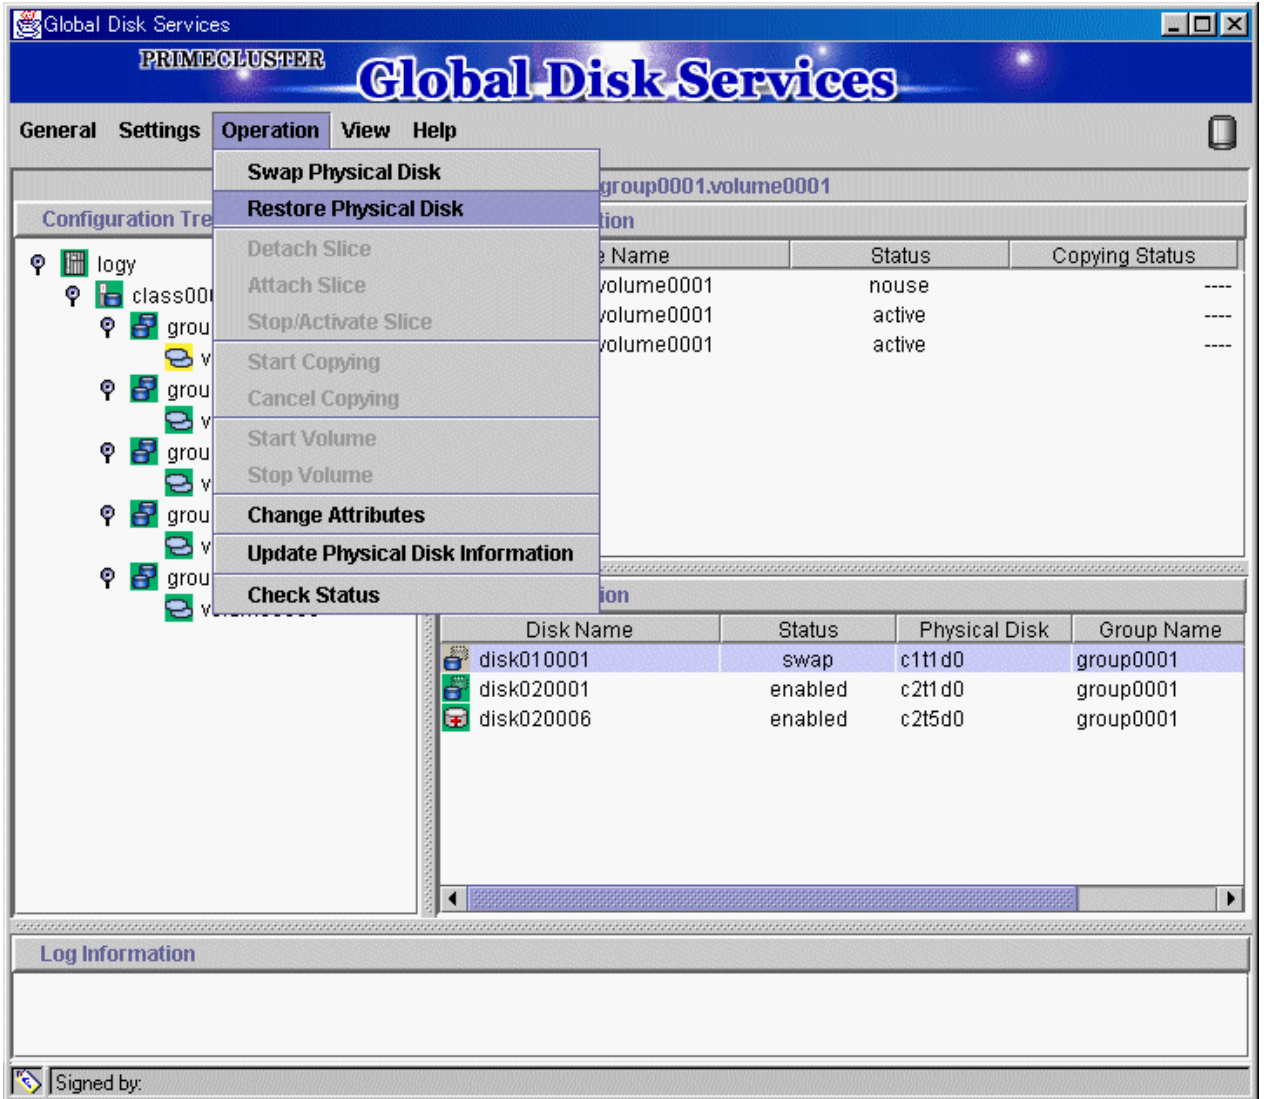
1. Selecting the physical disk to restore

Select the physical disk you want to restore.

2. Selecting [Restore Physical Disk] menu

In the Main screen [Operation] menu, select [Restore Physical Disk].

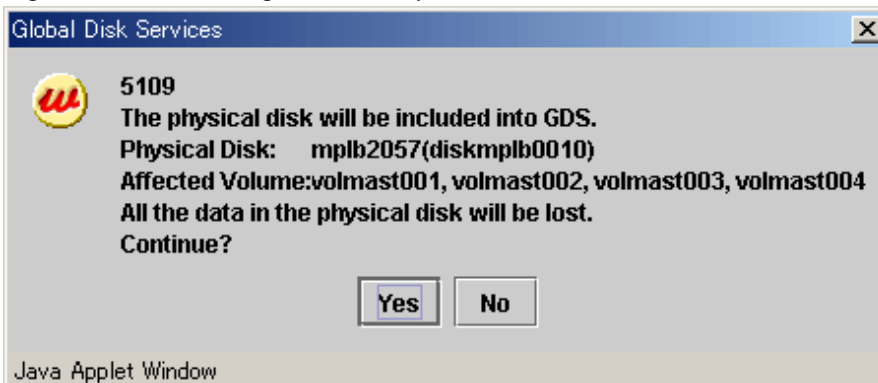
Figure 5.89 Restore Physical Disk



The confirmation screen shown below appears.

To continue the process, click <Yes>. Clicking <No> cancels the physical disk restore process.


Figure 5.90 Confirming Restore Physical Disk



3. Notifying online status

Click <OK>, and a message notifying the online status appears.

Figure 5.91 Notifying Online Completion



Note

In Cluster Systems

When physical disks for which disk resources are registered with the PRIMECLUSTER resource database are swapped, after recovering the physical disks, conduct automatic resource registration regardless of the class types. It is required to change VTOC volume names of original disks to those of the replacing disks, which are registered with the PRIMECLUSTER resource database. If automatic resource registration is not performed, a message indicating a disk unit connection error may be output, or failover may occur when nodes are started. For details on automatic resource registration, see "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide."

5.3.5 Disk Migration

This subsection describes the procedures for transferring volumes to other disks through use of GDS Snapshot proxy volumes. This method requires the following "Proxy Operations."

- Join
- Swap Slice
- Break

See

- For disk migration with proxy volumes, see "[1.5.5 Online Disk Migration](#)."
- For the operation flow, see "[5.1.7 Disk Migration](#)."

Note

Proxy Volume Use Conditions

See the following sections for points of concern.

- "[A.1.8 Proxy Configuration Preconditions](#)"
- "[A.1.9 Number of Proxy Volumes](#)"
- "[A.1.10 Proxy Volume Size](#)"
- "[A.1.11 Proxy Group Size](#)"

Join

Join a destination volume (a proxy volume) to a volume for disk migration (a master volume).

To perform disk migration on all the volumes within a group simultaneously, join a destination group (a proxy group) to a group for disk migration (a master group).

For "Join Proxy" procedures see ["5.2.5.1 Join."](#)

Swap Slice

Make sure that synchronization copying from the master to the proxy is complete in the Main screen and then swap slices comprising the master and slices comprising the proxy.

Follow the procedures below.

1. Selecting a destination proxy

Click an icon of a master volume for disk migration on the GDS Configuration Tree in the Main screen.

To perform disk migration on all the master volumes within a master group, click the master group icon.

Select [View]:[Details]:[Proxy Object] to view all the proxy objects related to the master object selected on the GDS Configuration Tree, in the Object Information Field.

In this field, click an icon of a destination proxy volume (or a proxy group).



Exchangeable Proxy Objects

"Swap Slices" is possible with proxy objects if they are joined and copy is complete.



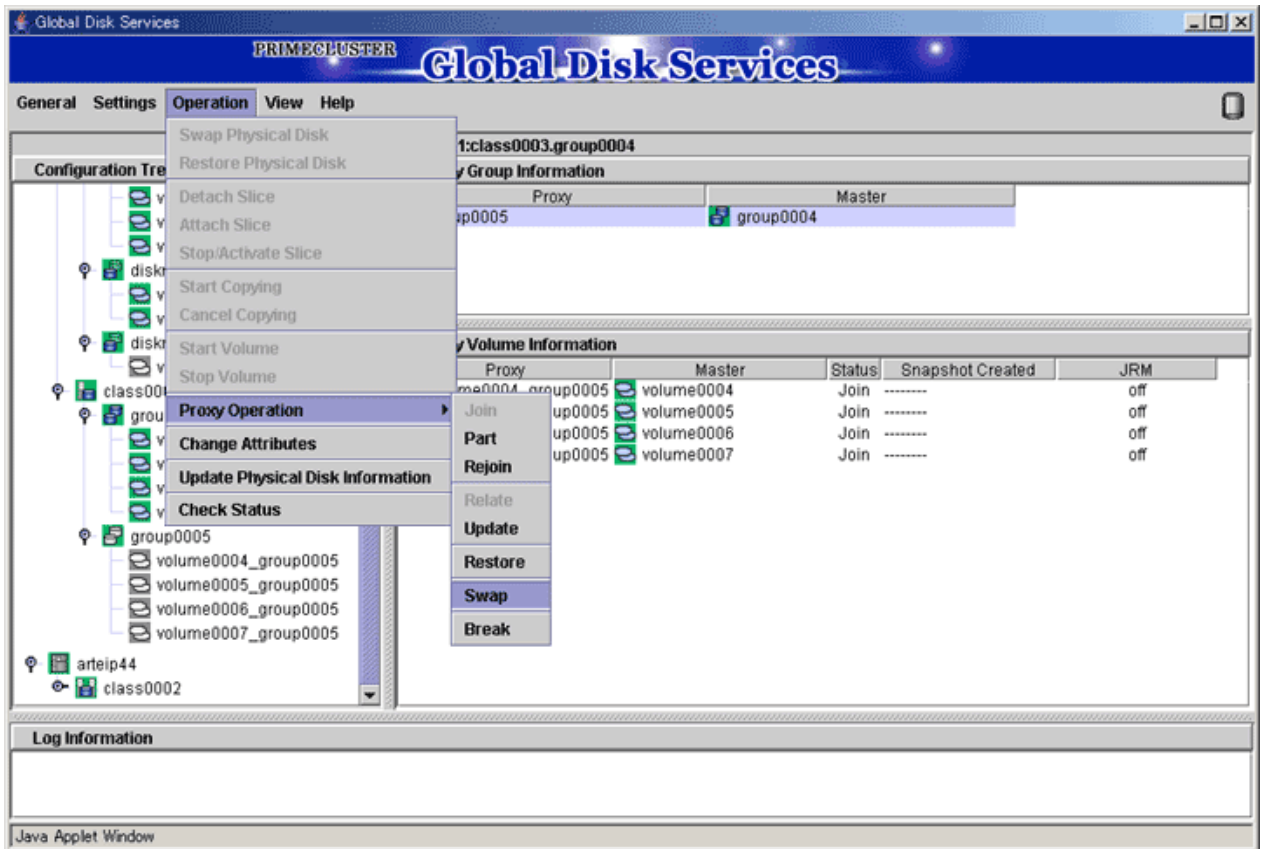
If There Is a Session by Disk Unit's Copy Function

If there is a session by a disk unit's copy function between the master and the proxy, slice swapping fails. For existing sessions, use the `sdxinfo -S -e` long command and check the FUNC field of the results. If the FUNC field for any slice of the master or the proxy shows a value other than the asterisk (*), a session exists between the master and the proxy. In this event, cancel the session with the `sdxproxy Cancel` command to swap the slices. For details, see ["D.6 sdxinfo - Display object configuration and status information"](#) and ["D.14 sdxproxy - Proxy object operations."](#)

2. Selecting the [Swap Slice] menu

Select [Operation]:[Proxy Operation]:[Swap Slice] in the Main screen.


Figure 5.92 Swap Slice



3. Confirming

A confirmation screen appears asking you whether to swap the slices.

Figure 5.93 Confirming "Swap Slice"

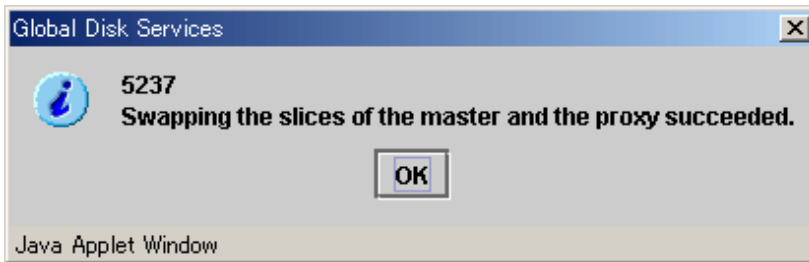


To continue the process, click <Yes>. Clicking <No> cancels the slice swapping process.

4. Information message of the completion

A message window appears informing you that the swap process is complete.

Figure 5.94 Information Message of "Swap Slice" Completion



Click <OK> to close the information message window.

Break

Break the relationship between the master and the proxy.

For "Break Proxy" procedures see "[5.5.6 Breaking a Proxy.](#)"

5.3.6 Copying Operation

The Copying Operation function controls synchronization copying of mirror volumes.

GDS provides the following copying operation.

- Start Copying
- Cancel Copying

Start Copying

Synchronization copying will be performed after attaching the slice that is in "invalid" or "copy-stop" status as a result of [Cancel Copying] operation.

For slices in "copy-stop" status, copying will resume from the point where copying was interrupted.

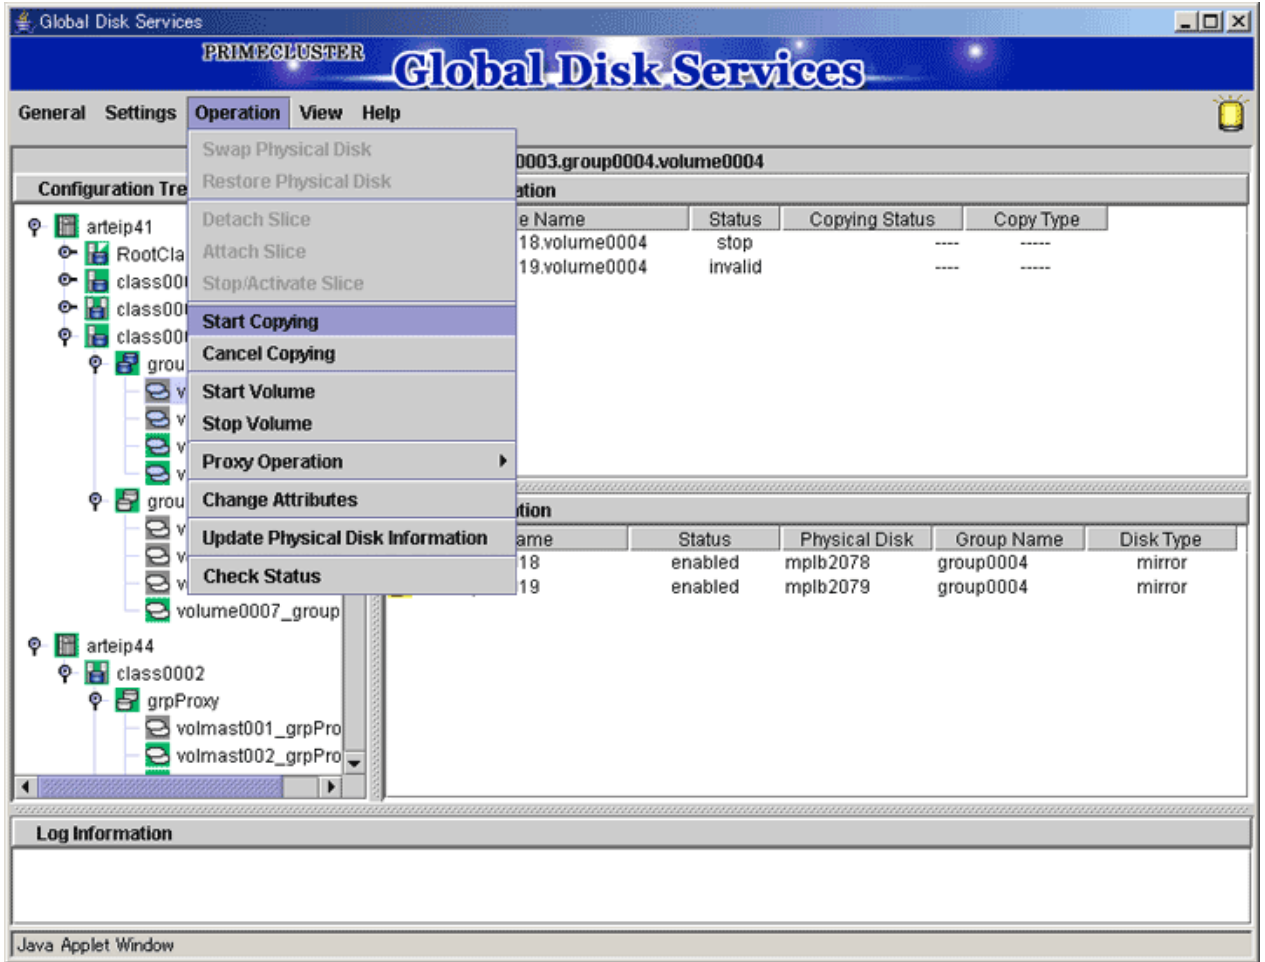
1. Select volume for synchronization copying

In the GDS Configuration tree field or Volume Information field, select the volume you want to copy by clicking the icon.

2. Selecting [Start Copying] menu

In the Main screen [Operation] menu, select [Start Copying].


Figure 5.95 Start Copying



The confirmation screen shown below appears.

To continue the process, click <Yes>. Clicking <No> cancels the copy start process.


Figure 5.96 Confirming Start Copying



3. [Start Copying] completion screen

Click <OK>, and the message notifying the [Start Copying] completion screen shown below appears.

Figure 5.97 Notifying Start Copying Completion



Note

[Start Copying] operation unavailable

When more than one volume is selected, [Start Copying] operation cannot be performed.

Cancel Copying

Execution of copying will be stopped to avoid effects caused by accessing the disk in the process of synchronization copying.

The slice whose copying has been stopped will be in "invalid" status. Perform [Start Copying] operation to resume its normal status

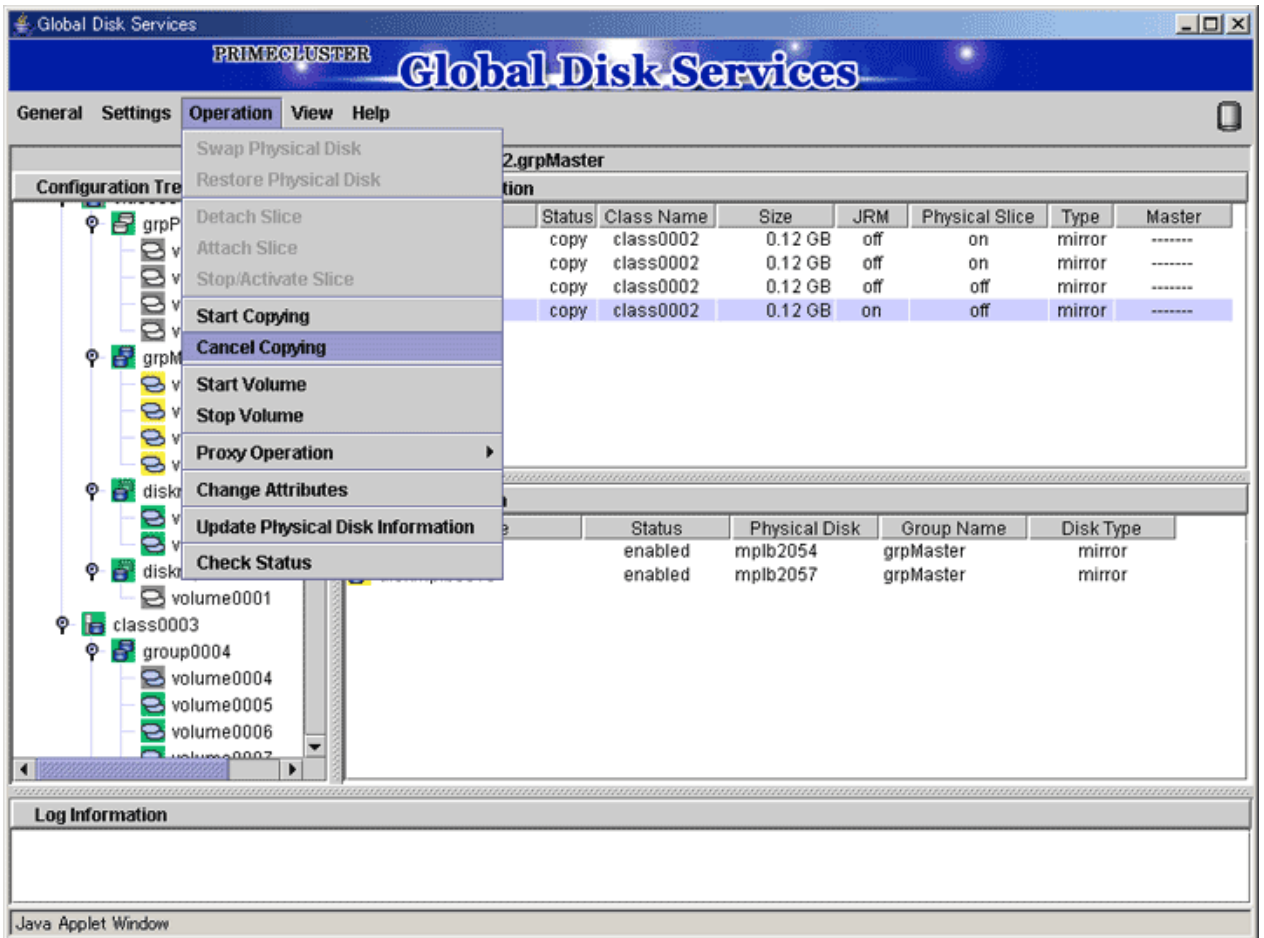
1. Select the volume to cancel copying

In the GDS Configuration tree field or Volume Information field, select the volume in "copy" status by clicking the icon.

2. Selecting [Cancel Copying] menu

In the Main screen [Operation] menu, select [Cancel Copying].


Figure 5.98 Cancel Copying



The confirmation screen shown below appears.

To continue the process, click <Yes>. Clicking <No> cancels the copy cancellation process.


Figure 5.99 Confirming Cancel Copying



3. [Cancel Copying] completion screen

Click <OK>, and the message notifying the [Cancel Copying] completion screen shown below appears.

Figure 5.100 Notifying Cancel Copying Completion



Note

[Cancel Copying] operation unavailable

When more than one volume is selected, [Cancel Copying] operation cannot be performed.

5.4 Changes

Class configurations, group configurations, and volume configurations can be changed through two types of operation: reconfiguration and attribute change.

This section explains the changing procedures for each configuration.

5.4.1 Class Configuration

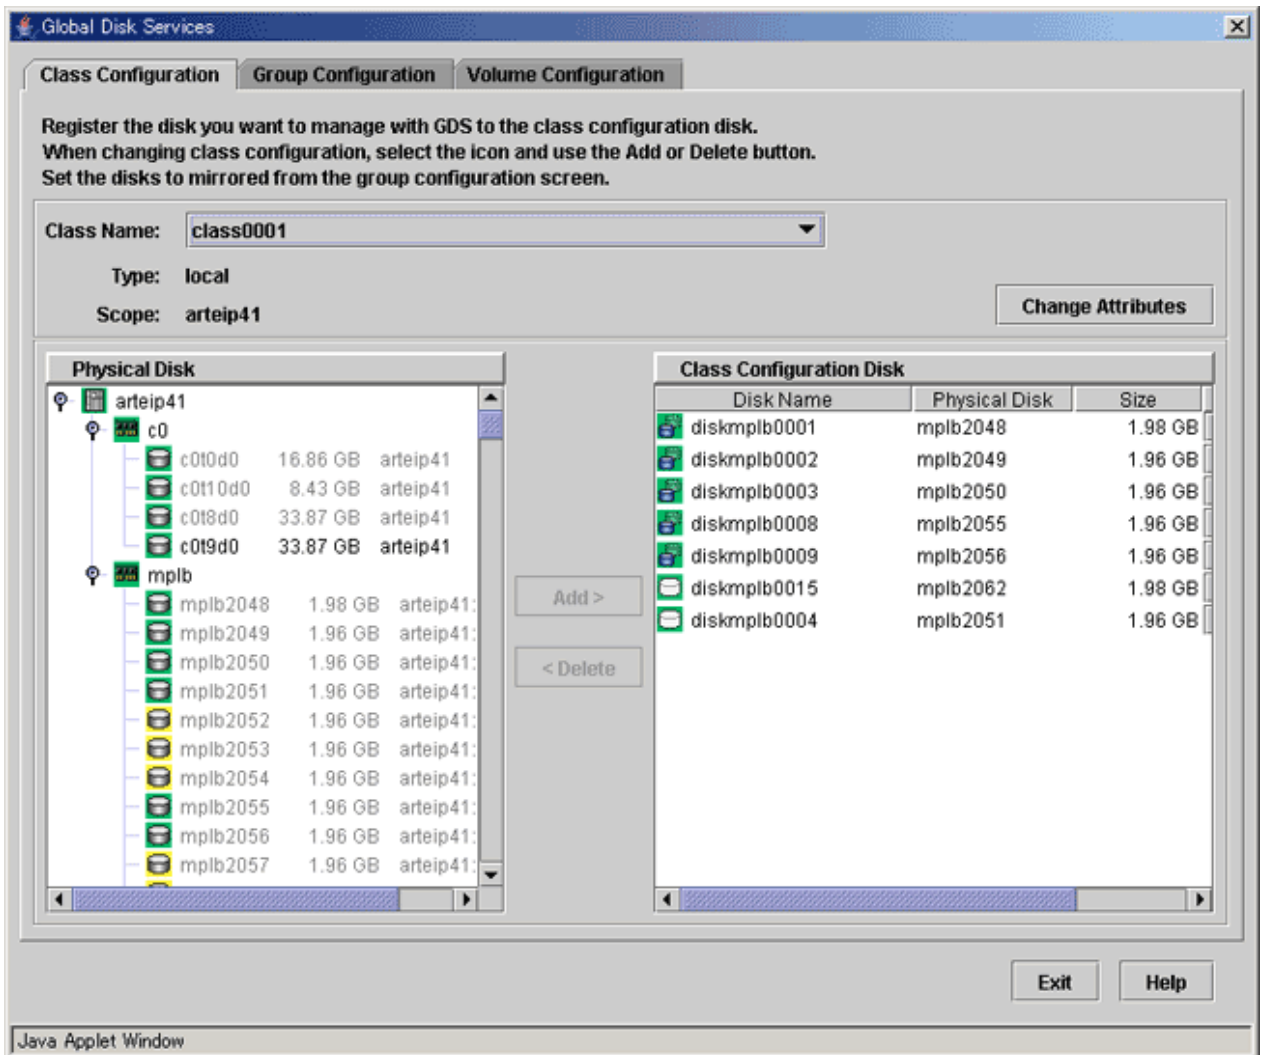
The procedures for changing class configuration are explained below.

Changing Class Configuration

1. Displaying Class Configuration screen

In the Main screen [Settings] menu, select [Class Configuration]. Class Configuration screen appears.

Figure 5.101 Class Configuration



2. Selecting class to change

Select the class you want to change from Class Configuration screen [Class Name].

3. Changing configuration

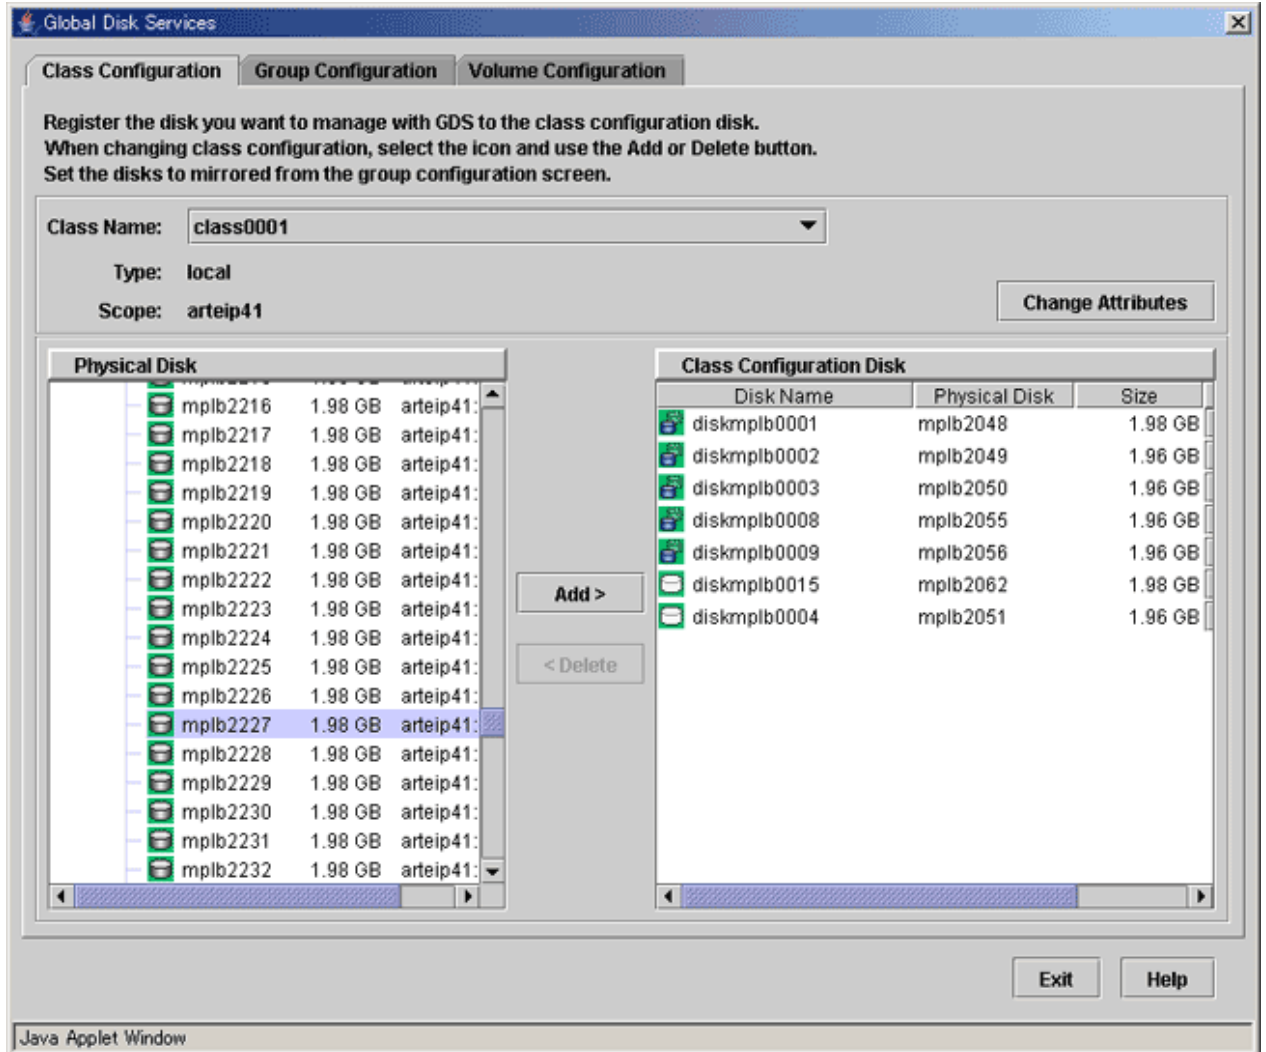
Follow the procedures below to register a physical disk (create a disk), or to remove a disk.

a. Registering a physical disk (Creating a disk)

1. Select the physical disk you want to register as disk from [Physical Disk] field.
2. Click <Add>.

- When changing the disk attributes from the initial values, change the disk type by selecting the appropriate disk from the [Class Configuration Disk] field.

Figure 5.102 Physical Disk Registration



b. Removing a disk

- Select the disk you want to remove from the [Class Configuration Disk] field.
- Click <Delete>.

4. Completing the change

If you have no other changes to make, click <Exit>.

Changing the Class Attributes

Change the class attributes using the following procedures.



Note

Preconditions for Changing Class Attributes

- When the class has active volumes, the following class attributes cannot be changed.
To change these class attributes, first stop all the volumes within the class on all the nodes within the class scope.
- Class name

- Type (from "shared" to "local")
 - Scope (node deletion)
- When the class has proxy objects, the class attributes cannot be changed.
To change attributes of such a class, break the proxy objects within the class.
-

1. Invoking the Class Attributes Definition screen

Invoke the Class Attributes Definition screen using one of the following methods.

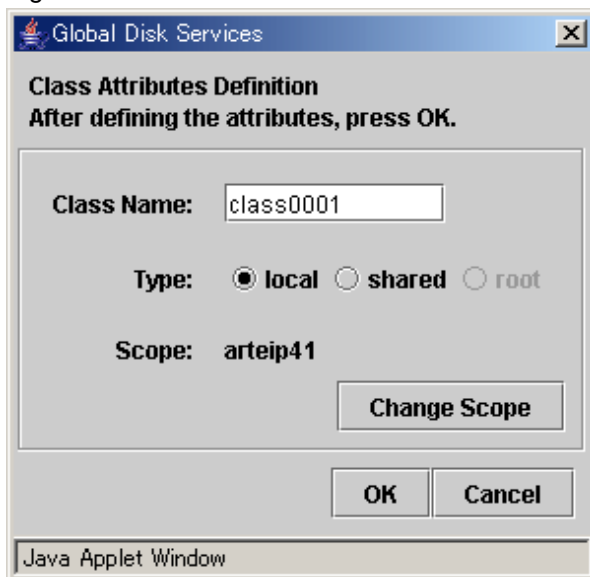
a. Operation menu in the Main screen

Click the target class's icon on the GDS configuration tree in the Main screen, and select [Operation]:[Change Attributes] to invoke the Class Attributes Definition screen.

b. Change Attributes button in the Class Configuration screen

Select the target class from the [Class Name] in the Class Configuration screen, and click the <Change Attributes> button to invoke the Class Attributes Definition screen.

Figure 5.103 Class Attributes Definition Screen



2. Changing attributes

a. For a single node

You can only change the class name.

b. For a cluster

You can change the following attributes:

- Class Name
- Type
- Scope



For information on assigning a class name, see "A.1.1 Object Name."

3. Implementing changes

Click <OK> to implement the attributes change, or <Cancel> to cancel.



Changes in Special File Path Names

Changing the class name will also change the special file path name used to access the volume.

Changing the Disk Attributes

When changing the disk attributes, there are two procedures to choose from:

- Using Main screen Operation menu.
- Using the Class Configuration Setting screen.



Preconditions for Changing Disk Attributes

Changes cannot be made in the attributes of disks that are connected to groups and disks that have volumes.

Using Main screen Operation menu


1. Selecting disk to change

Display the disk you want to change in the Main screen [Disk Information] field. Click the icon and select the disk you want to change.

2. Displaying Disk Attributes Definition screen

In the Main screen [Operation] menu, select [Change Attributes]. Disk Attributes Definition screen appears.

Figure 5.104 Disk Attribute Definition Screen



3. Changing attributes

You can change the following attributes.

- Disk Name

- Disk Type



See

.....
For details on assigning a disk name, see "[A.1.1 Object Name.](#)"
.....

4. **Applying changes**

Click <OK> to apply changes in the attributes change, or <Cancel> to cancel.

Using the Class Configuration screen

1. **Displaying the disk to change attributes**

Select the disk of which you want to change attributes in the [Class Configuration Disk] field.

2. **Changing attributes**

You can change the following attributes in the [Class Configuration Disk] field.

- Disk Name

Double-click the disk name and edit.

- Disk Type

Scroll to the right to get the [Disk Type] column, and select the disk type from the list.



See

.....
For details on assigning a disk name, see "[A.1.1 Object Name.](#)"
.....

5.4.2 Group Configuration

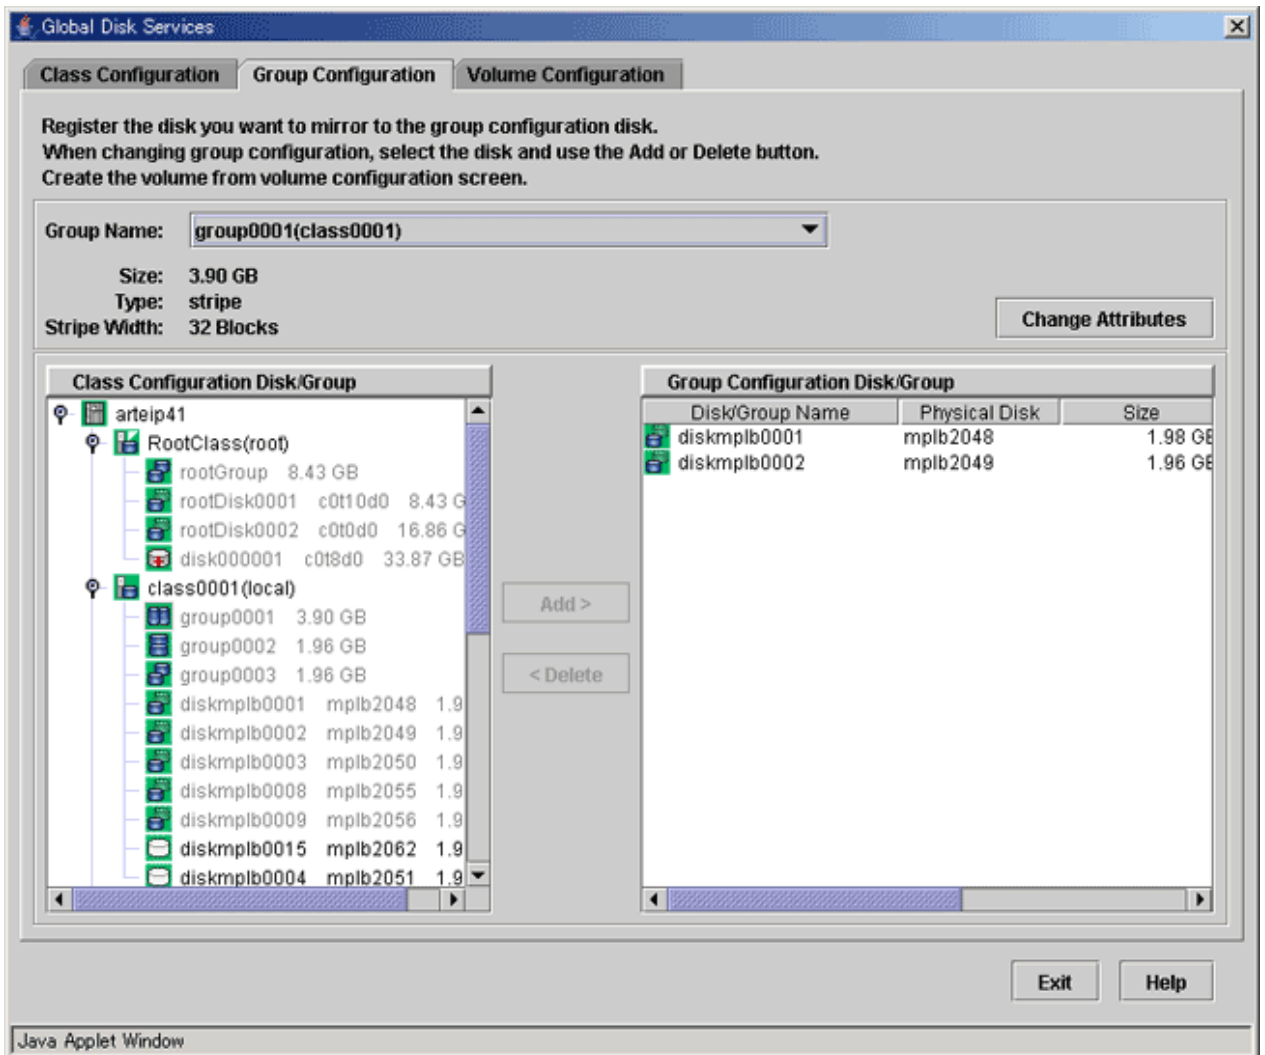
The procedures for changing group are explained below.

Changing Group Configuration

1. Displaying Group Configuration screen

In the Main screen [Settings] menu, select [Group Configuration]. Group Configuration screen appears.

Figure 5.105 Group Configuration



2. Selecting group to change

Select the group you want to change from Group Configuration screen [Group Name].

3. Changing configuration


Follow the procedures below to connect a disk/lower level group to a group, or to disconnect a disk/lower level group from a group.

a. Connecting a disk/lower level group

1. Select the disk/lower level group you want to add to group from [Class Configuration Disk/Group] field.

2. Click <Add>.

Figure 5.106 Connecting disk/lower level group



b. Disconnecting a disk/lower level group

1. Select the disk/lower level group you want to disconnect from [Group Configuration Disk/Group] field.
2. Click <Delete>.

4. Completing the change

If you have no other changes to make, click <Exit>.

Changing Group Attributes

You can only change the [group name]. Change the group attribute using the following procedures.



Note

Preconditions for Changing Group Attributes

- The attributes of lower level groups cannot be changed.
- The attribute of the highest level group with active volumes cannot be changed. To change the attribute of such a group, first stop all the volumes within the highest level group on all the nodes within the class scope.
- When the group has master volumes or proxy volumes, the attribute of the group cannot be changed. To change the attribute of such a group, first break the proxy volumes.

1. Invoking the Group Attributes Definition screen

Invoke the Group Attributes Definition screen using one of the following methods.

a. Operation menu in the Main screen

Click the target group's icon on the GDS configuration tree in the Main screen, and select [Operation]:[Change Attributes] to invoke the Group Attributes Definition screen.

b. Change Attributes button in the Group Configuration screen

Select the target group from the [Group Name] in the Group Configuration screen, and click the <Change Attributes> button to invoke the Group Attributes Definition screen.

Figure 5.107 Group Attributes Definition Screen



2. Changing group name

You can only change the group name.

Click <OK> to apply the change, or <Cancel> to cancel.



For information on assigning a group name, see "[A.1.1 Object Name.](#)"

5.4.3 Volume Configuration

The procedures for changing volume configuration are explained below.

Changing Volume Configuration

Change volume attributes using the following procedures.



Preconditions for Changing Volume Attributes

- When the volume is active, the following volume attributes cannot be changed. To change these volume attributes, first stop the volume on all the nodes within the class scope.
 - Volume name
 - Physical slice attribute
- When a volume's slice is detached temporarily, the physical slice attribute of the volume cannot be changed. To change this attribute, first attach the temporarily detached slice to the volume.

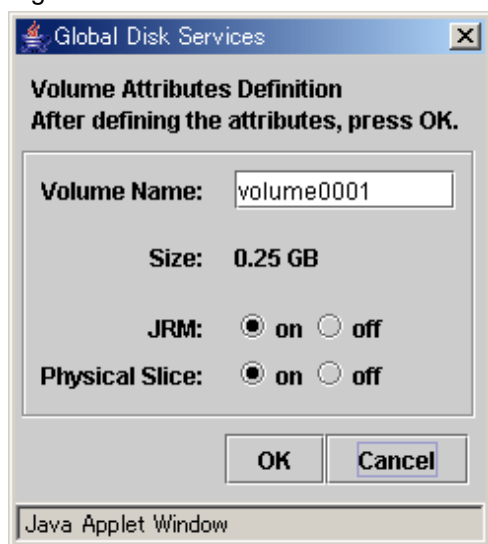
1. Selecting volume to change

On the GDS Configuration Tree in the Main screen, go to the object to be changed, and click the icon to select the volume to be changed.

2. Displaying Volume Attributes Definition screen

In the Main screen [Operation] menu, select [Change Attributes]. Volume Attributes Definition screen appears.

Figure 5.108 Volume Attributes Definition Screen



3. Changing attributes

You can change the following attributes:

- Volume Name
- JRM (on/off)
- Physical Slice (on/off)



See

For information on assigning a volume name, see "A.1.1 Object Name."

4. Applying changes

Click <OK> to apply changes in the attributes change, or <Cancel> to cancel.



Note

Changes in Special File Path Name

Changing the volume name will also change the special file path name used to access the volume.

5.5 Removals

When file system is not created, you can start the unmirroring process from removing the volume.

5.5.1 Removing a File System

The procedures for removing a file system are explained below.

1. **Displaying the File System Configuration screen**

In the Main screen [Settings] menu, select [File System Configuration]. File System Configuration screen appears.

2. **Selecting group/disk**

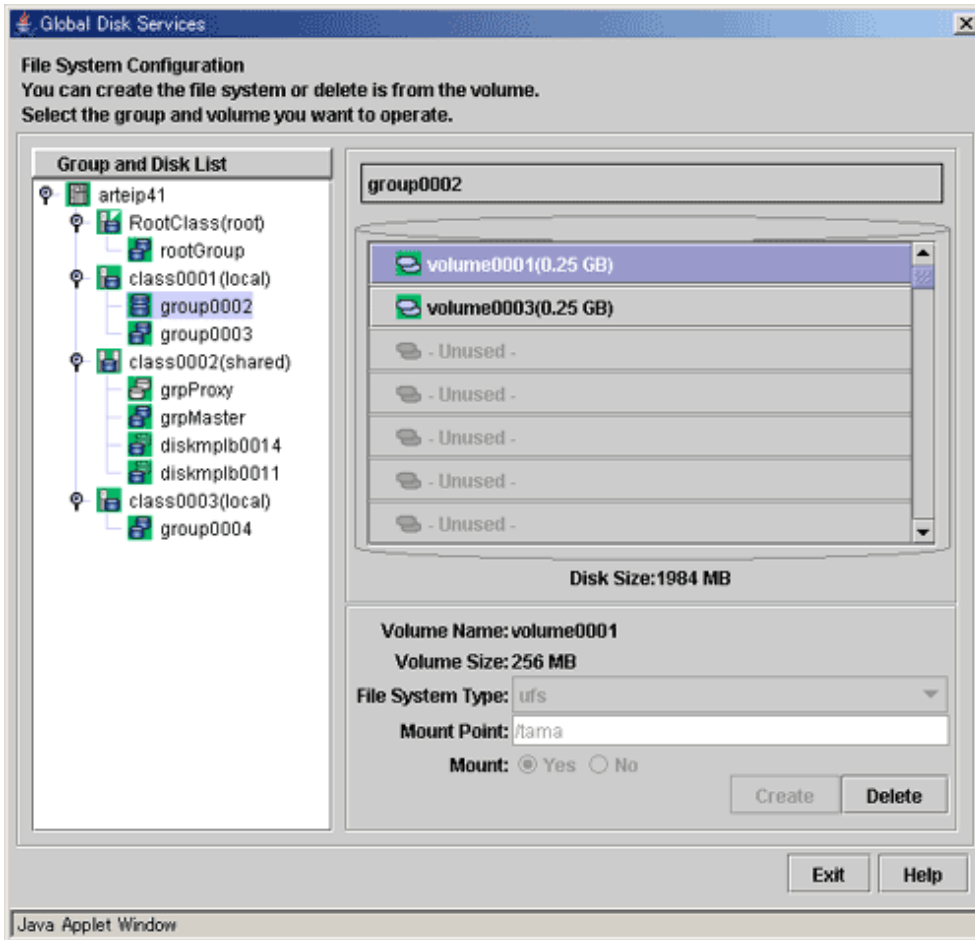
In the Group/Disk List, select a group or disk with which you want to perform an operation.

3. **Removing file system**

Select the volume from which you want to remove the file system, and click <Delete>.

When removing more than one file system, select the next volume, and repeat the same process.

Figure 5.109 Removing File System



 **Note**

In Cluster Systems

File systems on shared disk units used in a cluster system cannot be removed from the File System Configuration screen. For details on how to unset file systems on shared disk units used in a cluster system, see the "PRIMECLUSTER Reliant Monitor Services (RMS) with Wizard Tools Configuration and Administration Guide."

5.5.2 Removing a Volume

The procedures for removing a mirror volume are explained below.

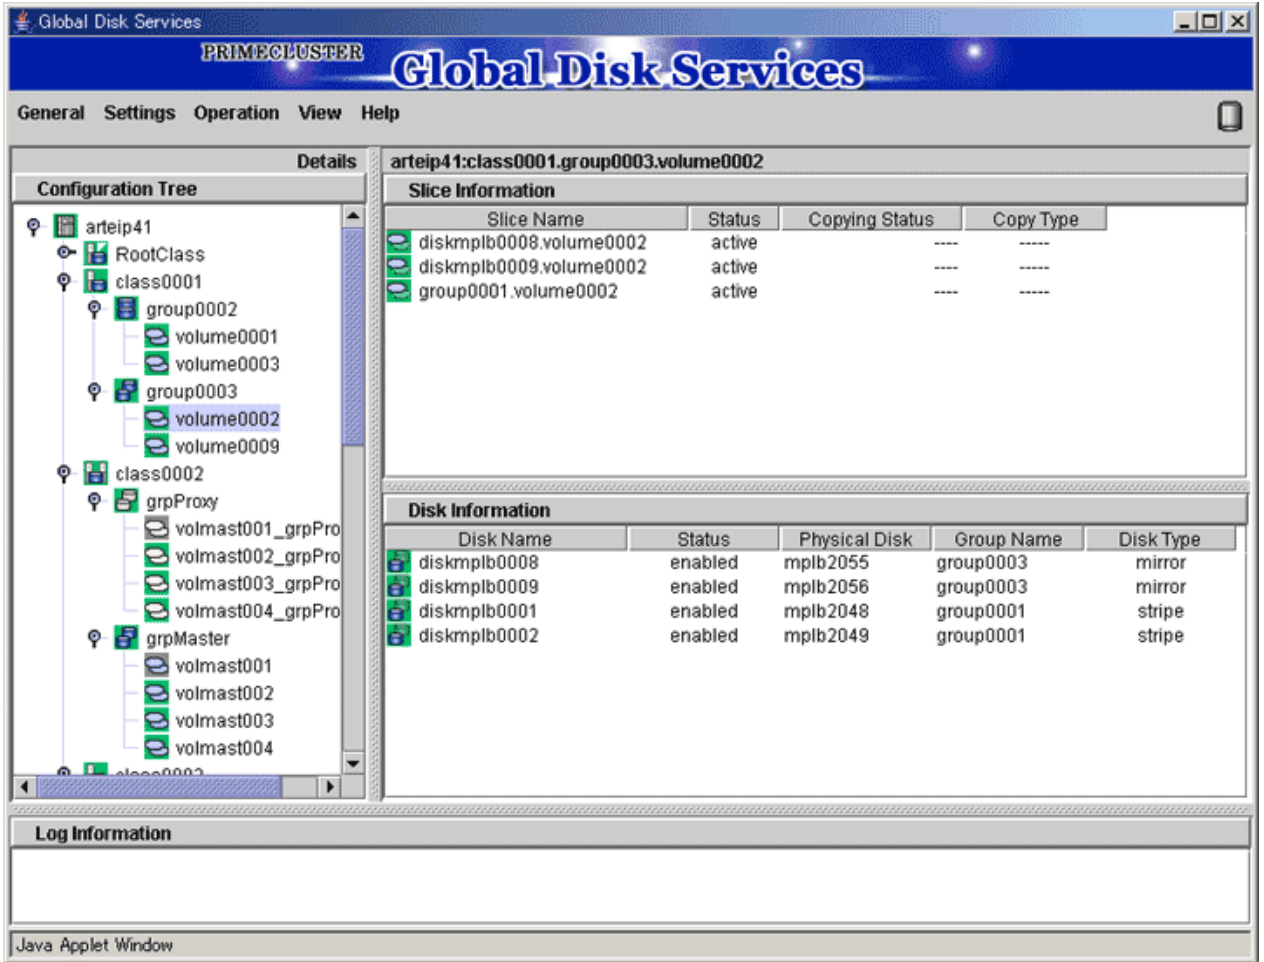
1. **Confirming the volume status**

A volume containing a temporarily detached slice cannot be removed.

When there is a temporarily detached slice, you must attach the slice before removing the volume.

Volume status can be confirmed in the Main screen.

Figure 5.110 Displaying Volume Information



2. Attaching a temporarily detached slice

When there is a temporarily detached slice, go to [Operation] menu and select [Attach Slice]. The detached slice will be attached again.

For information on Attach Slice operation, see "[5.3.2.1 Backup \(by Slice Detachment\)](#)."

3. Removing volume

The procedures for removing a volume are explained below.

1. Displaying the Volume Configuration screen

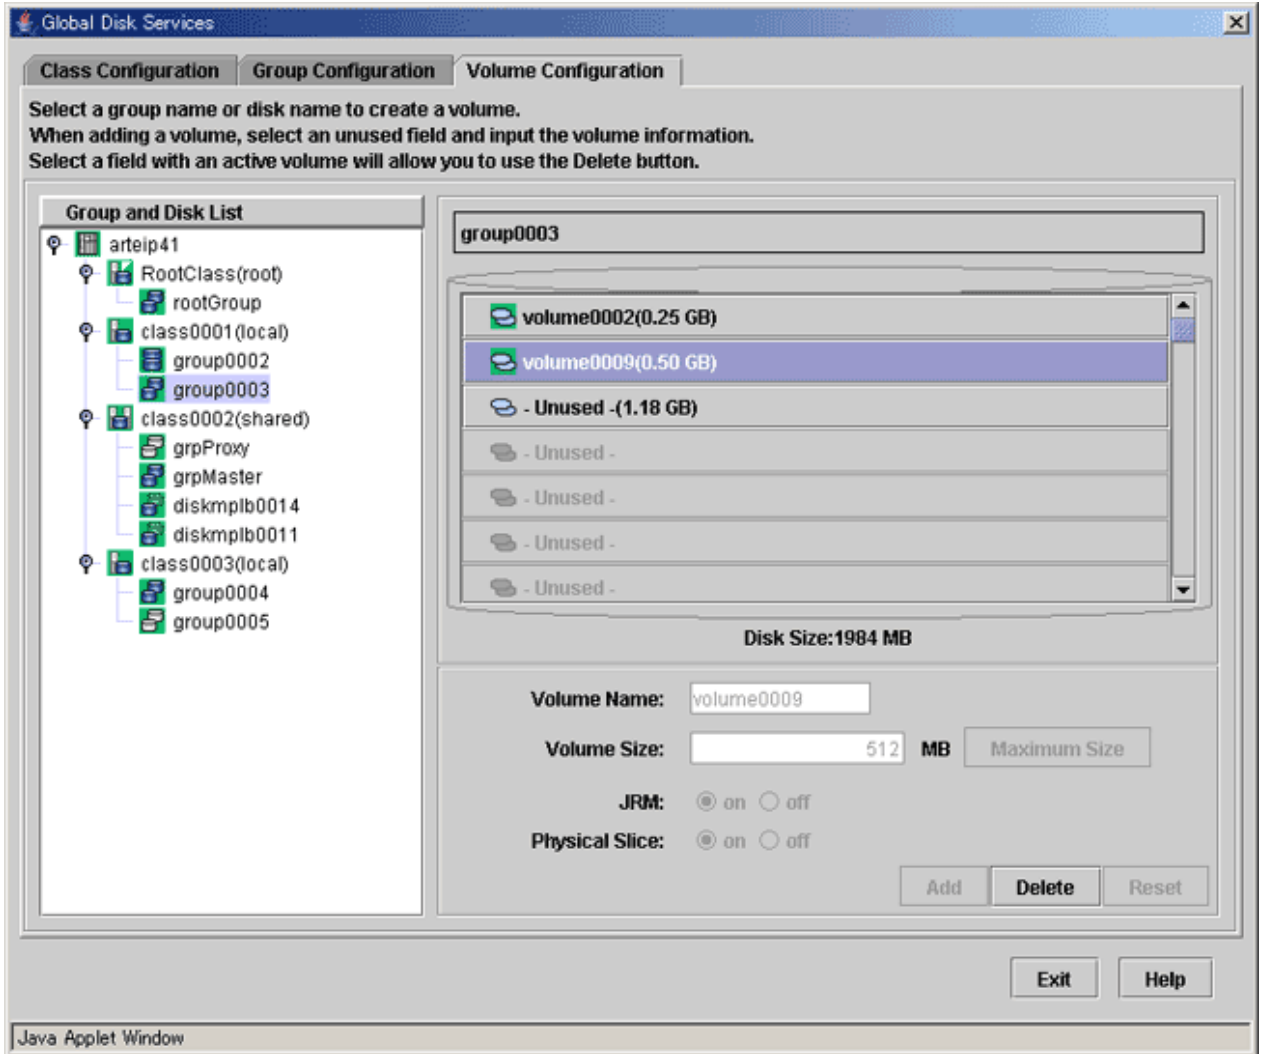
In the Main screen [Settings] menu, select [Volume Configuration]. Volume Configuration screen appears.

2. Removing a volume

Select the volume you want to remove, and click <Delete>.

When removing more than one volume, select the next volume, and repeat the same process.

Figure 5.111 Removing Volume



5.5.3 Removing a Group

Disconnecting all disks/lower level groups registered with a group will automatically remove the group.

The procedures are explained below.

1. Removing all volumes within the group

If there is even a single volume, you cannot remove the group.

Remove all volumes within the group by following the procedures described in "5.5.2 Removing a Volume."

2. Disconnecting disks/lower level groups from group

Disconnect all disks/lower level groups by following the procedures below.

1. Displaying Group Configuration screen

In the Main screen [Settings] menu, select [Group Configuration]. Group Configuration screen appears.

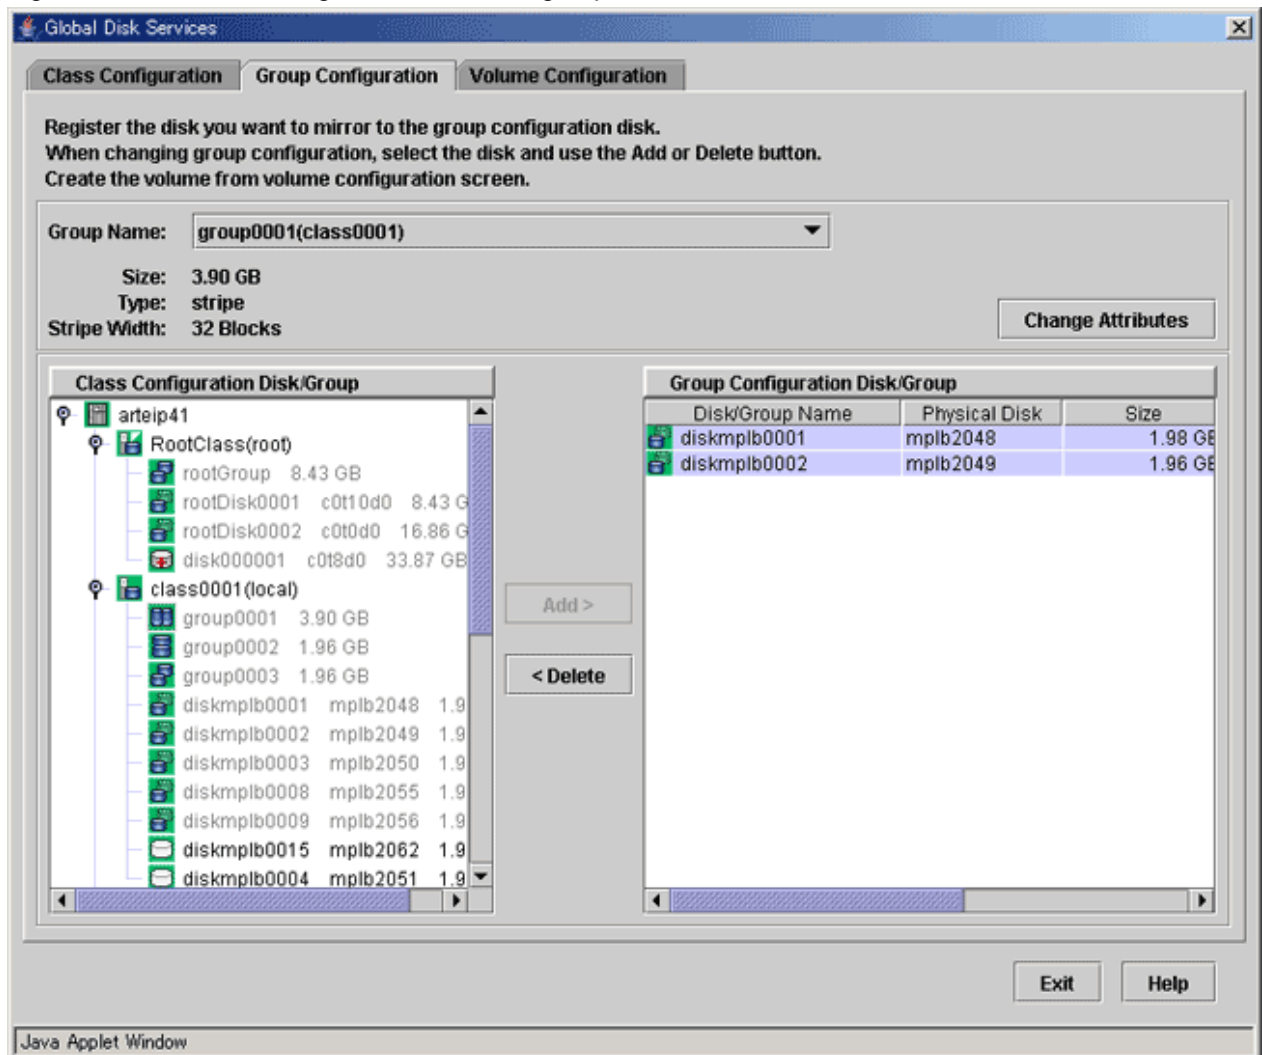
2. Selecting group to remove

Select the group you want to remove from Group Configuration screen [Group Name] field.

3. Disconnecting disk/lower level group

Select the disk/lower level group you want to remove from [Group Configuration Disk/Group] field, and click <Delete>.

Figure 5.112 Disconnecting Disk/Lower level group



5.5.4 Removing a Class

Removing all disks within a class will automatically remove the class.

The procedure is explained below.

1. Removing all groups within the class

Remove all groups within the class by following the procedure described in "5.5.3 Removing a Group."

2. Removing all disks from class

Disconnect all disks by following the procedure below.

1. Displaying Class Configuration screen

In the Main screen [Settings] menu, select [Class Configuration]. The Class Configuration screen appears.

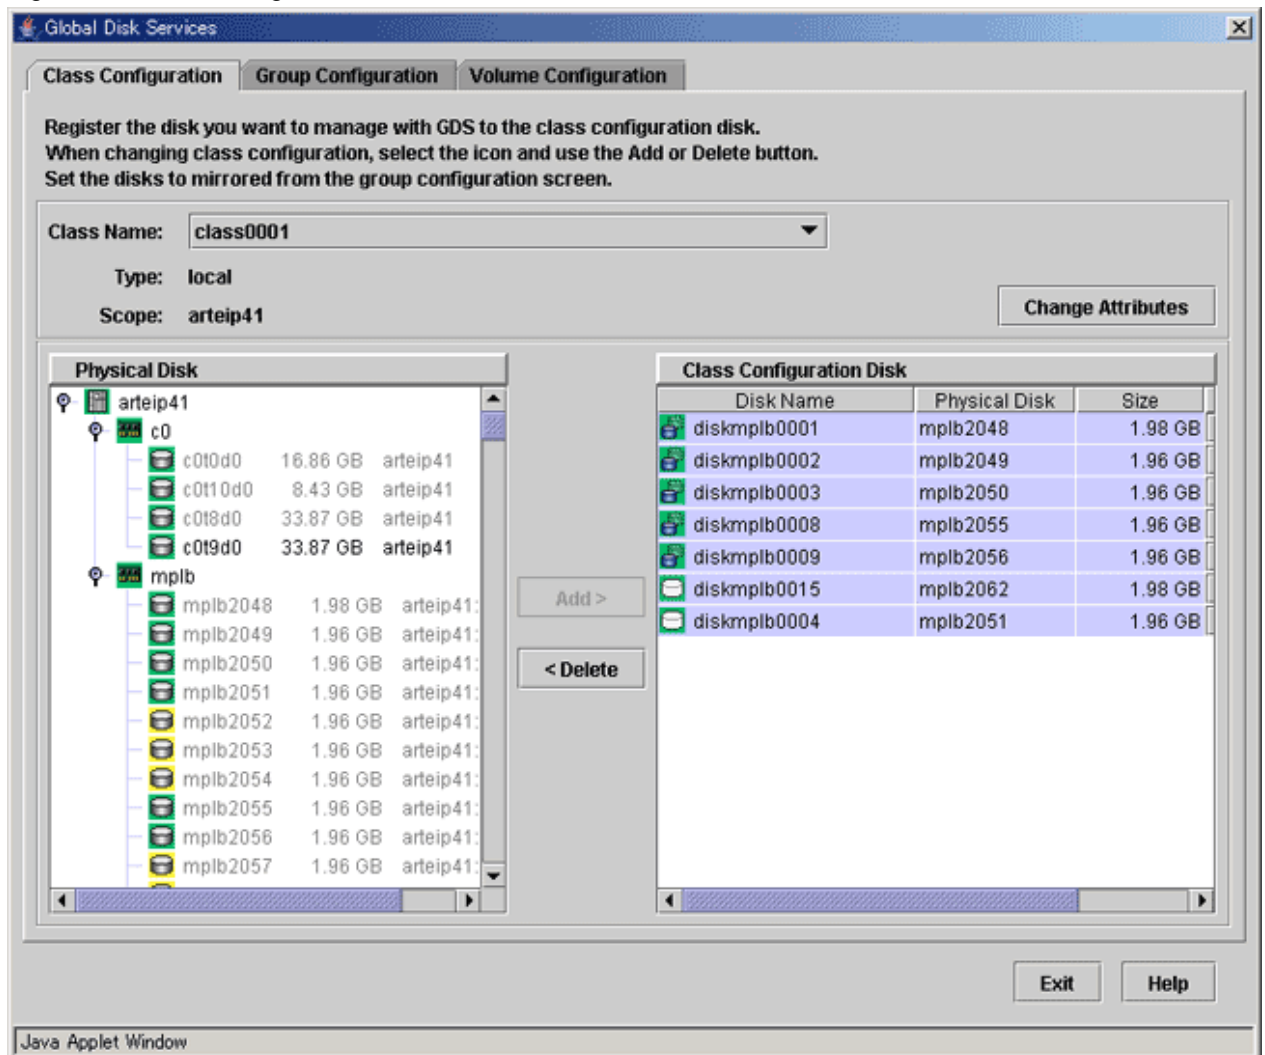
2. Selecting class to remove

Select the class you want to remove from Class Configuration screen [Class Name] field.

3. Removing disk

Select the disk you want to remove from [Class Configuration Disk] field, and click <Delete>.

Figure 5.113 Removing Disk



5.5.5 Unmirroring the System Disk

In this section, how to unmirror a system disk is explained.

In the [Settings] menu, select [Unmirror System Disk].



Note

When You Cannot Unmirror the System Disk

If system disks are under the following conditions, the system disks cannot be unmirrored. In these situations, restore the disk status first and unmirror the system disks.

- A slice has been detached temporarily with [Detach Slice].
- A disk has been disconnected with [Swap Physical Disk].
- A disk is in disabled status.
- All disks in each group contain a slice that is not active.
- There is a proxy volume.

- There is a volume to which a proxy volume is joined.

Note

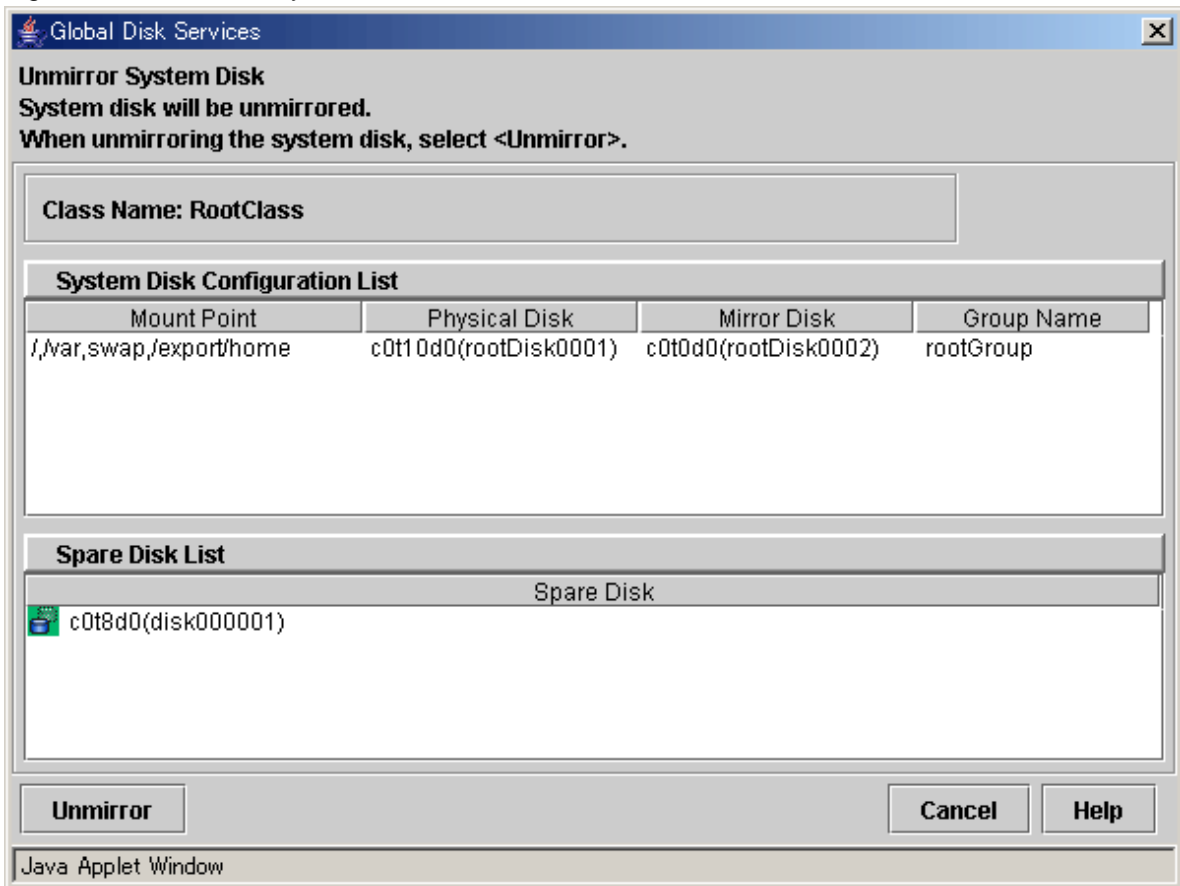
For Safe Unmirroring of System Disk

To safely unmirror the system disk, exit all active applications before proceeding to cancel system disk settings. After unmirroring the system disk is complete, promptly reboot the system in multi-user mode.

1. Confirming system disk configuration

System disk mirroring configurations are displayed.

Figure 5.114 Unmirror System Disk



If unmirroring is performed, the disk displayed in the [Mirror Disk] field is disconnected and the disk displayed in the [Physical Disk] field will be used as the system disk.

If the root class includes only system disks with their mirror disks and spare disks, the entire root class is removed. Here, the spare disks displayed in the spare disk field are also removed.

If the root class includes objects other than system disks, such as single disks and mirror groups, those settings are retained and only the unmirroring of system disks is executed. Here, the spare disk field does not display any information and spare disks are not removed.

To unmirror the system disk, click <Unmirror>. To cancel the unmirroring operation, click <Cancel>.

2. Confirming the unmirroring of system disk

If you click <Unmirror > in the Unmirror System Disk screen, the screen below appears.

Figure 5.115 Unmirroring System Disk Confirmation Screen

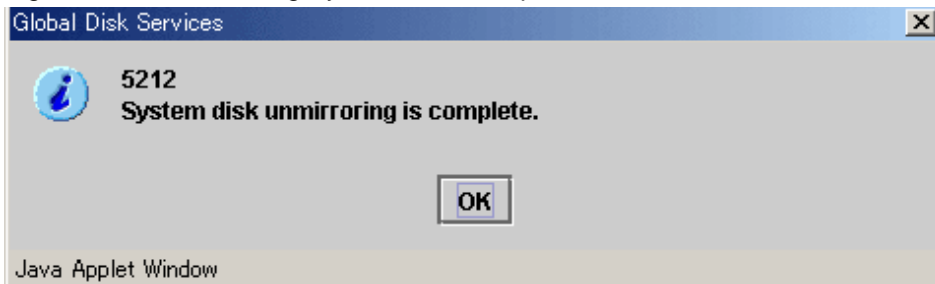


To continue the process, click <Yes>. Clicking <No> cancels the system disk unmirroring process.

3. Unmirroring System Disk Completion screen


If you click <Yes> on the Unmirroring System Disk Confirmation screen, the screen below appears.

Figure 5.116 Unmirroring System Disk Completion Screen



Confirm that unmirroring of system disk is complete and click <OK>.

Figure 5.117 System Reboot Notification Screen



The unmirroring of system disks will take place when the system is rebooted in multi-user mode.

Click <OK> and reboot the system in multi-user mode immediately.

Note

To stop or reboot the system immediately after rebooting the system, check the `sdxinfo(1)` command output to see whether the system disk has already been unmirrored in advance.

5.5.6 Breaking a Proxy

This subsection describes the procedures for breaking the relationships between masters and proxies in GDS Snapshot installed systems.

1. Selecting a master for break

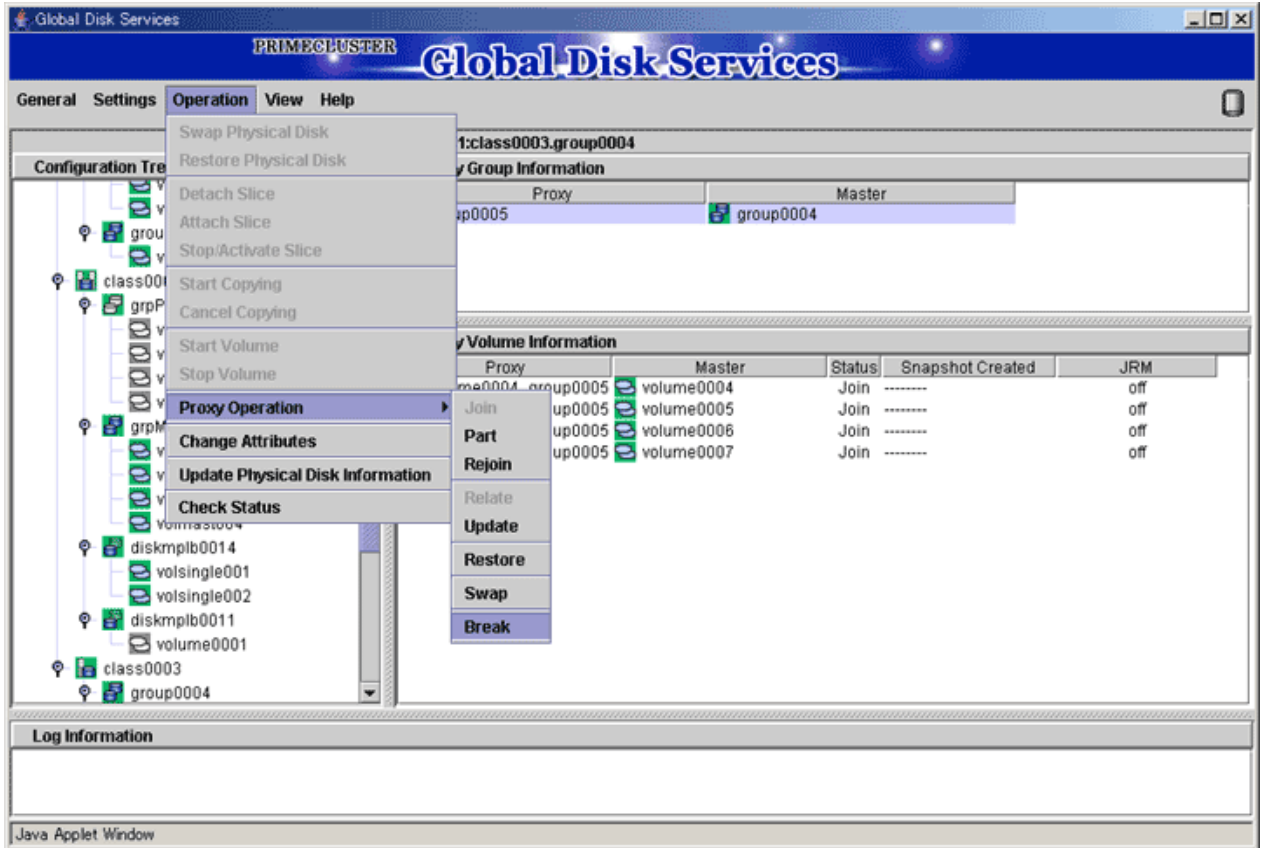
Click an icon of a master object for break on the GDS Configuration Tree in the Main screen.

Select [View]:[Details]:[Proxy Object] to view all the proxy objects related to the master selected on the GDS Configuration Tree, in the Object Information Field. In this field, click an icon of a proxy object for break.

2. Selecting the [Break] menu

Select [Operation]:[Proxy Operation]:[Break] in the Main screen.


Figure 5.118 Break



3. Confirming

A confirmation dialog box appears asking you whether to break the proxy.

Figure 5.119 Confirming "Break"




To continue the process, click <Yes>. Clicking <No> cancels the proxy break process.

4. Information message of the completion

A message window appears informing you that the break process is complete.

Figure 5.120 Information Message of "Break" Completion



Click <OK> to close the information message window.

Chapter 6 Backing Up and Restoring

The GDS mirroring function secures data from disk failures by preserving replicas of data on multiple disks. However, the mirroring function is not capable of protecting data against accidental erasure by the user or data crash due to an application malfunction. In addition, data can be lost due to multiple disk breakdown or large-scale disaster. To recover data when these troubles occur, data backup is mandatory. Be sure to back up data on a regular basis.

This chapter discusses the operating procedures for backing up and restoring GDS volumes to provide useful information concerning operational design for backup and restore.

6.1 Backing Up and Restoring a System Disk

Mirroring system disks will protect the data when a physical disk on one side fails. However, data must be restored from backup data created in advance if the data is damaged due to a critical failure caused by multiple breakdown and so on or by an operation mistake.

This section discusses the method of backing up data on a system disk to tape and the method of restoring data back to the system disk from tape. You must follow different restore procedures depending on whether or not the system can be booted.



Note

Data backed up before system disk mirroring cannot be restored back to the mirrored system disk. If system disk mirroring is configured, perform system disk backup using this procedure.



See

For backing up and restoring volumes (e.g. /opt, /home) other than system volumes (/, /usr, /var, swap area) among volumes in the root class, see "6.3 Backing Up and Restoring Local Disks and Shared Disks."

6.1.1 Backing Up

For securing consistency of backup target files, boot the system from a CD-ROM device (or network) or boot the system in single user mode to create backups. To ensure consistency, it is recommended to boot from a CD-ROM device (or network).

a) When booting the system from a CD-ROM device (or the network) and creating backups

a1) If there is a possibility of write access to backup target slices during backup, temporarily unmirror backup target disks.

For example, if the fsck(1M) command is used for a backup target file system, the command may write data to the backup target slices. Additionally, if a file system where ufs logging is on is backed up with the ufsdump(1M) command, the command writes data to the backup target slices. In these situations, it is necessary to unmirror the disks temporarily using this procedure in advance.

The command line shown below is an example of disconnecting disk Root2 from group Group1 for backing up Root1 in a situation where disks Root1 and Root2 are connected to Group1 and mirrored.

```
# sdxdisk -D -c System -g Group1 -d Root2
```

Confirm that only one disk is connected to group Group1 (Group1 is indicated in the GROUP field of only one disk).

```
# sdxinfo -D -c System
OBJ  NAME  TYPE  CLASS  GROUP  DEVNAM  DEVBLKS  DEVCONNECT  STATUS
-----
disk  Root1  mirror  System  Group1  c0t0d0  35368272  node1  ENABLE
disk  Root2  undef  System  *      c1t0d0  35368272  node1  ENABLE
```

Information

If disk Root1 has an INVALID slice, disconnect the Root1.

If a keep disk is disconnected (if the disconnected disk's TYPE field value is keep in the `sdxinfo -D` command output), change the type attribute to undef to keep its connectivity to a group (or remove the disk from the class once and register it again as an undef disk). For the disk type attribute setting method, see "[Changing the Disk Attributes](#)" of "[5.4.1 Class Configuration](#)" or "[D.7 sdxattr - Set objects attributes](#)."

Example) Changing the type attribute of Root1 to undef after keep disk Root1 is disconnected from Group1

```
# sdxattr -D -c System -d Root1 -a type=undef
```

See

- When using GDS Management View, see "[5.4.2 Group Configuration](#)" of "[5.4 Changes](#)."
- For the `fsck(1M)` command, `ufs logging`, and the `ufsdump(1M)` command, see the Solaris manual.

Note

If this procedure is not performed and data is written in backup target slices in the following procedures, the synchronized status of backup target volumes is not ensured. In this situation, using the procedure described in "[6.1.3 Restoring \(When the System Cannot Be Booted\)](#)," restore the backup target volumes.

a2) Shut down the system.

```
# shutdown -g0 -i0 -y
```

a3) Boot the system from a CD-ROM device (or the network).

- From a CD-ROM device

```
ok boot cdrom -s
```

- From the network

```
ok boot net -s
```

Note

- The version of an OS booted from a CD-ROM device (or network) must correspond to the version of the OS you are backing up.
- In a system where the disks are mirrored by GDS, if the system is booted from a CD-ROM device (or the network) and operations are performed improperly, the mirroring status may collapse. Be sure to perform operations according to the procedures described here. Operations other than those described here should not be performed.
- When mounting the file system on a physical slice at the time of recovery, mounting as read-only is required. If the slice is mounted without the read-only option on, it can be written in and thus the mirrored state may be ruined.

a4) Check the physical slice names of the root (/), /usr, and /var file systems you want to back up and make a note of them.

```
# eeprom nvramrc | grep sdx-root-slice
devalias sdx-root-slice-1 /devices/pci@1f,4000/scsi@3/sd@0,0:a,raw
devalias sdx-root-slice-2 /devices/pci@1e,4000/scsi@3/sd@0,0:a,raw

# eeprom nvramrc | grep sdx-usr-slice
devalias sdx-usr-slice-1 /devices/pci@1f,4000/scsi@3/sd@0,0:g,raw
devalias sdx-usr-slice-2 /devices/pci@1e,4000/scsi@3/sd@0,0:g,raw

# eeprom nvramrc | grep sdx-var-slice
devalias sdx-var-slice-1 /devices/pci@1f,4000/scsi@3/sd@0,0:b,raw
devalias sdx-var-slice-2 /devices/pci@1e,4000/scsi@3/sd@0,0:b,raw
```

In this example, the original physical slice names are:

```
/: /devices/pci@1f,4000/scsi@3/sd@0,0:a,raw
/usr : /devices/pci@1f,4000/scsi@3/sd@0,0:g,raw
/var : /devices/pci@1f,4000/scsi@3/sd@0,0:b,raw
```

The mirror physical slice names are:

```
/: /devices/pci@1e,4000/scsi@3/sd@0,0:a,raw
/usr : /devices/pci@1e,4000/scsi@3/sd@0,0:g,raw
/var : /devices/pci@1e,4000/scsi@3/sd@0,0:b,raw
```



Information

- If step a1) is performed, the mirror physical slice names are not given.
- If /usr and /var file systems do not exist, the physical slice names are not displayed as shown below.

```
# eeprom nvramrc | grep sdx-usr-slice
# eeprom nvramrc | grep sdx-var-slice
```

- For MPLB disks, check the physical slice names of paths after the first one as follows. In an environment booted from the network, if paths after the first one are defined, it is also possible to create backups from the physical slices of the paths.

```
# eeprom nvramrc | grep sdx-root-slice (First path)
devalias sdx-root-slice-1 /devices/pci@1f,4000/fibre-channel@4/
sd@10,0:a,raw
...

# eeprom nvramrc | grep sdx-root-alt-slice (Second path)
devalias sdx-root-alt-slice-1b /devices/pci@1f,4000/fibre-
channel@5/sd@10,0:a,raw
...
```

a5) From the physical slice device names confirmed in step a4), find the physical slice names.

The following procedure is an example of checking the original physical slice name of root (/).

```
# ls -l /dev/rdisk | grep /devices/pci@1f,4000/scsi@3/sd@0,0:a,raw
lrwxrwxrwx 1 root root 45 Oct 19 23:36 c0t0d0s0 -> ../../
devices/pci@1f,4000/scsi@3/sd@0,0:a,raw
```

For the grep command argument, specify the device name confirmed in step a4).

In this example, the original physical slice name of root (/) is c0t0d0s0.

a6) Create a backup of file system data in a tape medium.

The command line shown below is an example of one used to back up root file system data to a tape medium of tape device /dev/rmt/0 using the ufsdump(1M) command when the root file system type is ufs.

```
# ufsdump 0cf /dev/rmt/0 /dev/rdisk/c0t0d0s0
```

 See

For details on the ufsdump(1M) command, see the Solaris manual. After understanding the usage and the meaning of options, specify options properly.

a7) Reboot the system.

```
# shutdown -g0 -i6 -y
```

 See

If step a1) is not performed and a file system where ufs logging is on is backed up with the ufsdump(1M) command, the synchronized status of mirrors may collapse and it may be impossible to start the OS. After this kind of operation, be sure to restore the relevant file system. For the restore method, see "[6.1.2 Restoring \(When the System Can Be Booted\)](#)" and "[6.1.3 Restoring \(When the System Cannot Be Booted\)](#)."

a8) When step a1) was performed, reconnect the disk disconnected in that step to the group.

```
# sdxdisk -C -c System -g Group1 -d Root2
```

Confirm that disk Root2 is connected to group Group1 (Group1 is indicated in the GROUP field of the Root2 line).

```
# sdxinfo -D -c System
OBJ  NAME  TYPE  CLASS  GROUP  DEVNAM  DEVBLKS  DEVCONNECT  STATUS
-----
disk  Root1  mirror System  Group1 c0t0d0  35368272 node1      ENABLE
disk  Root2  mirror System  Group1 c1t0d0  35368272 node1      ENABLE
```

Synchronization copying will automatically take place, and when it is completed, the mirroring status is restored.

 See

When using GDS Management View, see "[5.4.2 Group Configuration](#)" of "[5.4 Changes](#)."

 Information

If a keep disk is disconnected from the group in step a1) and if the type attribute is not changed to undef, step a8) will result in an error and the error message "keep disk cannot be connected to existing group" will be output. In this event, change the disk's type attribute to undef first and retry step a8).

For the disk type attribute setting method, see "Information" in step a1).

b) When booting the system in single user mode and creating backups

b1) Exit all running application programs.

b2) Shut down the system.

```
# shutdown -g0 i0 -y
```

b3) Boot the system in single-user mode.

```
ok boot -s
```

b4) Check the volume of the file system to be backed up.

The following example shows the procedure for backing up the root (/) file system.

```
# mount
/ on /dev/sfdsk/System/dsk/root
read/write/setuid/largefiles on Mon Aug  9 04:02:38 1999
/usr on /dev/sfdsk/System/dsk/usr
read/write/setuid/largefiles on Mon Aug  9 04:02:38 1999
/var on /dev/sfdsk/System/dsk/var
read/write/setuid/largefiles on Mon Aug  9 04:02:38 1999
```

In this example, the device special file for the volume of the root (/) file system is /dev/sfdsk/System/dsk/root Volume.

b5) Back up data of the file system to a tape medium.

The following example shows the procedure for backing up data to a tape medium of tape device /dev/rmt/0 with the ufsdump(1M) command.

```
# ufsdump 0ucf /dev/rmt/0 /dev/sfdsk/System/rdsk/root
```

For the ufsdump command's argument, specify the device special file of the volume displayed in step b4).



For details on the ufsdump(1M) command, see the Solaris manual. After understanding the usage and the meaning of options, specify options properly.

b6) Reboot the system.

```
# shutdown -g0 i6 -y
```

6.1.2 Restoring (When the System Can Be Booted)

1) Exit all running application programs. If higher security is required, you should make a backup of the system disk in advance. For details on the backup procedure, see "6.1.1 Backing Up."

2) Disconnect disks other than the disk that will be the restore destination from the group to have only one disk connected to the group.

The command line shown below is an example of one used to disconnect disk Root2 from group Group1 to use disk Root1 as the restore destination when Root1 and Root2 are connected to Group1 and mirrored.

```
# sdxdisk -D -c System -g Group1 -d Root2
```

Confirm that only one disk is connected to group Group1 (Group1 is indicated in the GROUP field of only one disk).

```
# sdxinfo -D -c System
OBJ  NAME  TYPE  CLASS  GROUP  DEVNAM  DEVBLKS  DEVCONNECT  STATUS
```

```

-----
disk Root1 mirror System Group1 c0t0d0 35368272 node1 ENABLE
disk Root2 undef System * c1t0d0 35368272 node1 ENABLE
-----

```

Information

If disk Root1 has an INVALID slice, disconnect the Root1.

If a keep disk is disconnected (if the disconnected disk's TYPE field value is keep in the `sdxinfo -D` command output), change the type attribute to undef to keep its connectivity to a group (or remove the disk from the class once and register it again as an undef disk). For the disk type attribute setting method, see "[Changing the Disk Attributes](#)" of "[5.4.1 Class Configuration](#)" or "[D.7 sdxattr - Set objects attributes](#)."

Example) Changing the type attribute of Root1 to undef after keep disk Root1 is disconnected from Group1

```
# sdxattr -D -c System -d Root1 -a type=undef
```

See

When using GDS Management View, see "[5.4.2 Group Configuration](#)" of "[5.4 Changes](#)."

3) Shutdown the system.

```
# shutdown -g30 -i0 -y
```

4) Since file systems such as root (/), /usr, and /var cannot be unmounted during system operation, boot the system from a CD-ROM device (or network).

- From a CD-ROM device

```
ok boot cdrom -s
```

- From the network

```
ok boot net -s
```

Note

- In a system where mirroring is conducted by GDS, booting from a CD-ROM device and performing an illegal operation may disrupt the mirroring environment. Be sure to follow the procedure described here and to avoid operations not provided in this procedure.
- When mounting the file system on a physical slice at the time of recovery, mounting as read-only is recommended to write-lock.
- To restore Solaris 9, the version of an OS booted from a CD-ROM device (also network or emergency startup disk for restoration) must correspond to the version of the OS you are restoring.
- To restore Solaris 10, to which the kernel patch 137111-08 or earlier has been applied, the version of an OS booted from a CD-ROM device (also network or emergency startup disk for restoration) must correspond to the version of the OS, to which the kernel patch 137111-08 or earlier has been applied, you are restoring.
- To restore Solaris 10, to which the kernel patch 137137-09 or later has been applied, the version of an OS booted from a CD-ROM device (also network or emergency startup disk for restoration) must correspond to the version of the OS, to which the kernel patch 137137-09 or later has been applied, you are restoring.

5) Check the device names of physical slices that will be the restore destinations for the root (/), /usr, and /var file systems.

```
# eeprom nvramrc | grep sdx-root-slice
devalias sdx-root-slice-1 /devices/pci@1f,4000/scsi@3/sd@0,0:a,raw

# eeprom nvramrc | grep sdx-usr-slice
devalias sdx-usr-slice-1 /devices/pci@1f,4000/scsi@3/sd@0,0:g,raw

# eeprom nvramrc | grep sdx-var-slice
devalias sdx-var-slice-1 /devices/pci@1f,4000/scsi@3/sd@0,0:b,raw
```

In this example, the physical slice device names are:

```
/: /devices/pci@1f,4000/scsi@3/sd@0,0:a,raw
/usr : /devices/pci@1f,4000/scsi@3/sd@0,0:g,raw
/var : /devices/pci@1f,4000/scsi@3/sd@0,0:b,raw
```

Information

- If /usr and /var file systems do not exist, the physical slice names are not displayed as shown below.

```
# eeprom nvramrc | grep sdx-usr-slice
# eeprom nvramrc | grep sdx-var-slice
```

- For MPLB disks, check the physical slice names of paths after the first one as follows. In an environment booted from the network, if paths after the first one are defined, it is also possible to create backups from the physical slices of the paths.

```
# eeprom nvramrc | grep sdx-root-slice (First path)
devalias sdx-root-slice-1 /devices/pci@1f,4000/fibre-channel@4/
sd@10,0:a,raw
...

# eeprom nvramrc | grep sdx-root-alt-slice (Second path)
devalias sdx-root-alt-slice-1b /devices/pci@1f,4000/fibre-
channel@5/sd@10,0:a,raw
...
```

6) Identify the physical slice names by the physical slice device names confirmed in step 5).

```
# ls -l /dev/rdisk | grep /devices/pci@1f,4000/scsi@3/sd@0,0:a,raw
lrwxrwxrwx 1 root root 45 Oct 19 23:36 c0t0d0s0 -> ../../
devices/pci@1f,4000/scsi@3/sd@0,0:a,raw

# ls -l /dev/rdisk | grep /devices/pci@1f,4000/scsi@3/sd@0,0:g,raw
lrwxrwxrwx 1 root root 45 Oct 19 23:36 c0t0d0s6 -> ../../
devices/pci@1f,4000/scsi@3/sd@0,0:g,raw

# ls -l /dev/rdisk | grep /devices/pci@1f,4000/scsi@3/sd@0,0:b,raw
lrwxrwxrwx 1 root root 45 Oct 19 23:36 c0t0d0s1 -> ../../
devices/pci@1f,4000/scsi@3/sd@0,0:b,raw
```

For the grep command arguments, specify the device names confirmed in step 6).

In this example, the physical slice names of root (/), /usr, and /var are c0t0d0s0, c0t0d0s6, and c0t0d0s1 respectively.

7) Restore the backup data on a tape medium back to the root (/), /usr, and /var file systems.

The following example shows the procedure for restoring data backed up with the `ufsdump(1M)` command.

```
# newfs /dev/rdisk/c0t0d0s0
# mount -F ufs /dev/dsk/c0t0d0s0 /mnt
# cd /mnt
# ufsrestore rvf /dev/rmt/0
# cd /
# umount /mnt

# newfs /dev/rdisk/c0t0d0s6
# mount -F ufs /dev/dsk/c0t0d0s6 /mnt
# cd /mnt
# ufsrestore rvf /dev/rmt/0
# cd /
# umount /mnt

# newfs /dev/rdisk/c0t0d0s1
# mount -F ufs /dev/dsk/c0t0d0s1 /mnt
# cd /mnt
# ufsrestore rvf /dev/rmt/0
# cd /
# umount /mnt
```

Note

Data should not be restored from backup data created before the system disk was mirrored.

See

For details on the `ufsrestore(1M)` command, see the Solaris manual. After understanding the usage and the meaning of options, specify options properly.

8) Install the boot block to the root partition of system disk.

```
# installboot /usr/platform/`uname -i`/lib/fs/ufs/bootblk /dev/rdisk/
c0t0d0s0
```

9) Reboot the system.

```
# shutdown -g0 -i6 -y
```

10) Reconnect the disks that were disconnected in step 2).

```
# sdxdisk -C -c System -g Group1 -d Root2
```

Confirm that disk Root2 is connected to group Group1 (Group1 is indicated in the GROUP field of the Root2 line).

```
# sdxinfo -D -c System
```

OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	DEVCONNECT	STATUS
disk	Root1	mirror	System	Group1	c0t0d0	35368272	node1	ENABLE
disk	Root2	mirror	System	Group1	c1t0d0	35368272	node1	ENABLE

Synchronization copying will automatically take place, and when it is completed, the mirroring status is restored.



See

When using GDS Management View, see "5.4.2 Group Configuration" of "5.4 Changes."



Information

If a keep disk is disconnected from the group in step 2) and if the type attribute is not changed to undef, step 10) will result in an error and the error message "keep disk cannot be connected to existing group" will be output. In this event, change the disk's type attribute to undef first and retry step 10).

For the disk type attribute setting method, see "Information" in step 2).

6.1.3 Restoring (When the System Cannot Be Booted)

1) Boot the system from the CD-ROM unit (or network).

Perform the boot operation following the procedure of step 1a) if the operating system is Solaris 10, or the procedure of step 1b) if the operating system is Solaris 9.

1a) If the operating system is Solaris 10:

- From a CD-ROM device

```
ok boot cdrom -s
```

- From the network

```
ok boot net -s
```

1b) If the operating system is Solaris 9:

- From a CD-ROM device

```
ok boot net -as
```

Perform the following interactive procedure.

```
Fujitsu PRIMEPOWER 200 1x SPARC64-III 296MHz, No Keyboard
OpenBoot 3.11 A.01, 256 MB memory installed, Serial #15809255.
Ethernet address 0:80:17:28:3a:e7, Host ID: 80f13ae7.
SCF Version: 2.1.E , RCI Version: 0001

Rebooting with command: boot cdrom -as
Boot device: /pci@17,4000/scsi@3/disk@4,0:f File and args: -as
Enter filename [kernel/unix]:
Enter filename [kernel/sparcv9/unix]:
Enter default directory for modules [/platform/FJSV,GPUS/kernel /
```

```
platform/sun4us/kernel /kernel usr/
kernel]:

Name of system file [etc/system]: /dev/null
~~~~~
SunOS Release 5.8 Version Generic_108528-09 64-bit
Copyright 1983-2001 Sun Microsystems, Inc. All rights reserved.
root filesystem type [ufs]:
Enter physical name of root device
[/pci@17,4000/scsi@3/disk@4,0:b]:
Configuring /dev and /devices
...
```

- From the network

```
ok boot net0 -as
```

Perform the following interactive procedure.

```
Resetting ...
...
Executing last command: boot net0 -as
Boot device: /pci@83,4000/network@1,1 File and args: -as
2ac00 Enter filename [kernel/sparcv9/unix]:
Enter default directory for modules [/platform/FJSV,GPUSK/kernel /
platform/sun4
us/kernel /kernel /usr/kernel]:
Name of system file [etc/system]: /dev/null
~~~~~
SunOS Release 5.8 Version Generic_108528-01 64-bit
Copyright 1983-2000 Sun Microsystems, Inc. All rights reserved.
root filesystem type [nfs]:
Enter physical name of root device
[/pci@83,4000/network@1,1]:
root name:
whoami: no domain name
root name:
Configuring /dev and /devices
...
```

Note

- In a system where mirroring is conducted by GDS, booting from a CD-ROM device (or network) and performing an illegal operation may disrupt the mirroring environment. Be sure to follow the procedure described here and to avoid operations not provided in this procedure.
- When mounting the file system on a physical slice at the time of recovery, mounting as read-only is recommended to write-lock.
- To restore Solaris 9, the version of an OS booted from a CD-ROM device (also network or emergency startup disk for restoration) must correspond to the version of the OS you are restoring.
- To restore Solaris 10, to which the kernel patch 137111-08 or earlier has been applied, the version of an OS booted from a CD-ROM device (also network or emergency startup disk for restoration) must correspond to the version of the OS, to which the kernel patch 137111-08 or earlier has been applied, you are restoring.
- To restore Solaris 10, to which the kernel patch 137137-09 or later has been applied, the version of an OS booted from a CD-ROM device (also network or emergency startup disk for restoration) must correspond to the version of the OS, to which the kernel patch 137137-09 or later has been applied, you are restoring.

2) Check the device names of physical slices that will be the restore destinations for the root (/), /usr, and /var file systems.

```
# eeprom nvramrc | grep sdx-root-slice
devalias sdx-root-slice-1 /devices/pci@1f,4000/scsi@3/sd@0,0:a,raw
devalias sdx-root-slice-2 /devices/pci@1e,4000/scsi@3/sd@0,0:a,raw

# eeprom nvramrc | grep sdx-usr-slice
devalias sdx-usr-slice-1 /devices/pci@1f,4000/scsi@3/sd@0,0:g,raw
devalias sdx-usr-slice-2 /devices/pci@1e,4000/scsi@3/sd@0,0:g,raw

# eeprom nvramrc | grep sdx-var-slice
devalias sdx-var-slice-1 /devices/pci@1f,4000/scsi@3/sd@0,0:b,raw
devalias sdx-var-slice-2 /devices/pci@1e,4000/scsi@3/sd@0,0:b,raw
```

In this example, the physical slice device names are:

```
/: /devices/pci@1f,4000/scsi@3/sd@0,0:a,raw
/usr : /devices/pci@1f,4000/scsi@3/sd@0,0:g,raw
/var : /devices/pci@1f,4000/scsi@3/sd@0,0:b,raw
```

The mirror physical slice device names are:

```
/: /devices/pci@1e,4000/scsi@3/sd@0,0:a,raw
/usr : /devices/pci@1e,4000/scsi@3/sd@0,0:g,raw
/var : /devices/pci@1e,4000/scsi@3/sd@0,0:b,raw
```

Note

If multiplicity of the mirroring configuration is 3 or more, there are multiple mirror physical slices.

Information

- If /usr and /var file systems do not exist, the physical slice names are not displayed as shown below.

```
# eeprom nvramrc | grep sdx-usr-slice
# eeprom nvramrc | grep sdx-var-slice
```

- For MPLB disks, check the physical slice names of paths after the first one as follows. In an environment booted from the network, if paths after the first one are defined, it is also possible to create backups from the physical slices of the paths.

```
# eeprom nvramrc | grep sdx-root-slice (First path)
devalias sdx-root-slice-1 /devices/pci@1f,4000/fibre-channel@4/
sd@10,0:a,raw
...

# eeprom nvramrc | grep sdx-root-alt-slice (Second path)
devalias sdx-root-alt-slice-1b /devices/pci@1f,4000/fibre-
channel@5/sd@10,0:a,raw
...
```

3) Identify the physical slice names by the physical slice device names confirmed in step 2).

```

# ls -l /dev/rdisk | grep /devices/pci@1f,4000/scsi@3/sd@0,0:a,raw
lrwxrwxrwx  1 root      root          45 Oct 19 23:36 c0t0d0s0 -> ../../
devices/pci@1f,4000/scsi@3/sd@0,0:a,raw

# ls -l /dev/rdisk | grep /devices/pci@1e,4000/scsi@3/sd@0,0:a,raw
lrwxrwxrwx  1 root      root          45 Oct 19 23:36 c1t0d0s0 -> ../../
devices/pci@1e,4000/scsi@3/sd@0,0:a,raw

# ls -l /dev/rdisk | grep /devices/pci@1f,4000/scsi@3/sd@0,0:g,raw
lrwxrwxrwx  1 root      root          45 Oct 19 23:36 c0t0d0s6 -> ../../
devices/pci@1f,4000/scsi@3/sd@0,0:g,raw

# ls -l /dev/rdisk | grep /devices/pci@1e,4000/scsi@3/sd@0,0:g,raw
lrwxrwxrwx  1 root      root          45 Oct 19 23:36 c1t0d0s6 -> ../../
devices/pci@1e,4000/scsi@3/sd@0,0:g,raw

# ls -l /dev/rdisk | grep /devices/pci@1f,4000/scsi@3/sd@0,0:b,raw
lrwxrwxrwx  1 root      root          45 Oct 19 23:36 c0t0d0s1 -> ../../
devices/pci@1f,4000/scsi@3/sd@0,0:b,raw

# ls -l /dev/rdisk | grep /devices/pci@1f,4000/scsi@3/sd@0,0:b,raw
lrwxrwxrwx  1 root      root          45 Oct 19 23:36 c1t0d0s1 -> ../../
devices/pci@1f,4000/scsi@3/sd@0,0:b,raw

```

For the grep command arguments, specify the device names confirmed in step 3).

In this example, the physical slice names are:

```

/: c0t0d0s0
/usr : c0t0d0s6
/var : c0t0d0s1

```

The mirror physical slice names are:

```

/: c1t0d0s0
/usr : c1t0d0s6
/var : c1t0d0s1

```

4) Restore the backup data on a tape medium back to the original physical slices respectively.

The following example shows the procedure for restoring backup data created with the ufsdump(1M) command.

```

# newfs /dev/rdisk/c0t0d0s0
# mount -F ufs /dev/dsk/c0t0d0s0 /mnt
# cd /mnt
# ufsrestore rvf /dev/rmt/0
# cd /
# umount /mnt

# newfs /dev/rdisk/c0t0d0s6
# mount -F ufs /dev/dsk/c0t0d0s6 /mnt
# cd /mnt
# ufsrestore rvf /dev/rmt/0
# cd /
# umount /mnt

# newfs /dev/rdisk/c0t0d0s1

```



```
# mount -F ufs /dev/dsk/c0t0d0s1 /mnt
# cd /mnt
# ufsrestore rvf /dev/rmt/0
# cd /
# umount /mnt
```

Note

- If restoring data to an original physical slice fails due to an I/O error and so on, restore data to the mirror physical slice.
- Data should not be restored from backup data created before the system disk was mirrored.

See

For details on the ufsrestore(1M) command, see the Solaris manual. After understanding the usage and the meaning of options, specify options properly.

5) Install the boot block to the root partition of system disk.

```
# installboot /usr/platform/`uname -i`/lib/fs/ufs/bootblk /dev/rdisk/
c0t0d0s0
```

6) Check the sector size and cylinder size of disk and make a note of them.

```
# prtvtoc /dev/rdisk/c0t0d0s0
...
* Dimensions:
*   512 bytes/sector
*   248 sectors/track
*   19 tracks/cylinder
*   4712 sectors/cylinder
...
```

In this example, the sector size is 512 (bytes) = 1 (block) and the cylinder size is 512 x 4712 (bytes) = 4712 (blocks).

7) Check the physical slice that contains the first sector (VTOC area) of the disk.

In the prtvtoc(1M) command output, a physical slice with 0 in the First Sector field contains the disk's first sector.

```
# prtvtoc /dev/rdisk/c0t0d0s0
...
*
* Partition  Tag  Flags      First   Sector   Last
*           Tag  Flags      Sector Count   Sector  Mount Directory
*   0         2   00         0    409656  409655
*   1         7   00    409656  307242  716897
*   2         5   00         0    4117568  4117567
*   3         3   01    737906  388648  1126553
*   4        33   00    716898   21008   737905
*   5         4   00   1126554   942734  2069287
*   6         4   00   2069288  2048280  4117567
...
```

In this example, the physical slice (Partition 0) contains the first sector.

Information

Partition 2 is a special partition that corresponds to the entire disk.

8) Copy data from the restored original physical slices to the mirror physical slices.

Copy data by following step 8a) if the physical slice does not contain the first sector or by following step 8b) if the physical slice contains the first sector.

Note

- If the multiplicity of the mirroring configuration is 3 or more, there are multiple mirror physical slices. Be sure to copy data to all of the mirror physical slices.
- There is no need to copy data on an original slice that was not restored in step 4).

8a) If the physical slice does not contain the first sector:

To reduce copying time, specify the cylinder size found in step 7) to the bs option (I/O block size) of the dd(1M) command.

```
# dd if=/dev/rdisk/c0t0d0s6 of=/dev/rdisk/c1t0d0s6 bs=4712b
# dd if=/dev/rdisk/c0t0d0s1 of=/dev/rdisk/c1t0d0s1 bs=4712b
```

8b) If the physical slice contains the first sector:

Take the following steps.

8b-1) Copy only the first cylinder of the physical slice first. Skip the first sector not to overwrite the first sector (VTOC area) of the mirror slice and copy by sector.

To skip the first sector, specify 1 to the iseek option and the oseek option of the dd(1M) command. Specify the sector size found in step 6) to the bs option. Specify the cylinder size (in blocks) found in step 6) subtracting 1 to the count option.

```
# dd if=/dev/rdisk/c0t0d0s0 of=/dev/rdisk/c1t0d0s0 bs=512 iseek=1 oseek=1
count=4711
```

8b-2) Copy the rest of cylinders of the physical slice next.

To reduce copying time, specify the cylinder size found in step 6) to the bs option of the dd(1M) command. To skip the first sector copied in step 8b-1), specify 1 to the iseek option and the oseek option of the dd(1M) command.

```
# dd if=/dev/rdisk/c0t0d0s0 of=/dev/rdisk/c1t0d0s0 bs=4712b iseek=1 oseek=1
```

9) Reboot the system.

```
# shutdown -g0 -i6 -y
```

6.2 Backing Up and Restoring a System Disk through an Alternative Boot Environment

This section discusses the method of backing up system disks with the GDS Snapshot function and creating an alternative boot environment by use of the backup copies, and the method of restoring the system disks through the alternative boot environment.

GDS Snapshot can collect snapshots of system disks (replications at a certain point) in the disk areas for backup (proxy volumes) during the system operation. You can configure an alternative boot environment to allow the system to be booted through such a proxy volume in order that the system operates in the alternative boot environment even if the system cannot be booted normally due to a failed system disk or damaged data. After switching to the alternative boot environment, you can also restore the original boot environment simply by restoring backup disk data to the original system disks and rebooting the system.

In this manner, operation down time while backing up and restoring system disks and time required for recovery from a system disk failure can be reduced drastically.

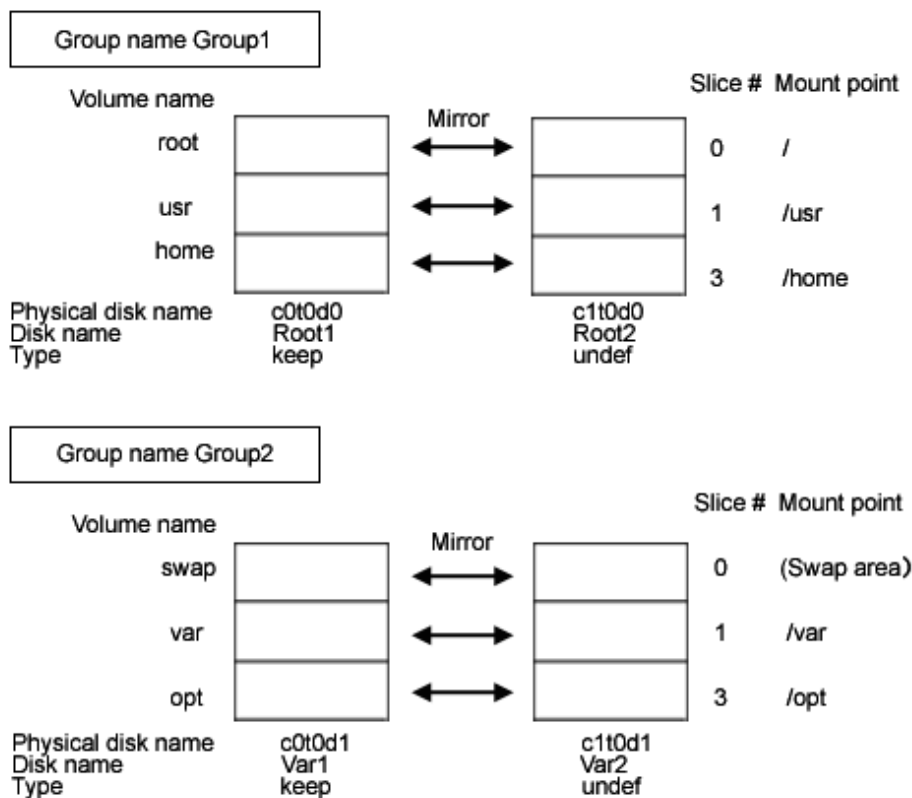
6.2.1 System Configuration

In preparation for alternative boot environment creation, mirror system disks. This sub-section provides an example of mirroring system disks in the configuration as shown in the figure below.

Information

System disk mirroring is not required. A configuration without c1t0d0 and c1t0d1 in the following figure is also available. However, the mirroring configuration as shown in the following figure is recommended for systems that require high availability.

Figure 6.1 System Disk Configuration



The disk areas (proxy groups) for backing up system disks are necessary.

Information

Mirroring of the disk areas for backup is not required. A configuration without c3t0d0 and c3t0d1 in the following figure is also available. However, the mirroring configuration as shown in the following figure is recommended for systems that require high availability.

Figure 6.2 Proxy Group Configuration

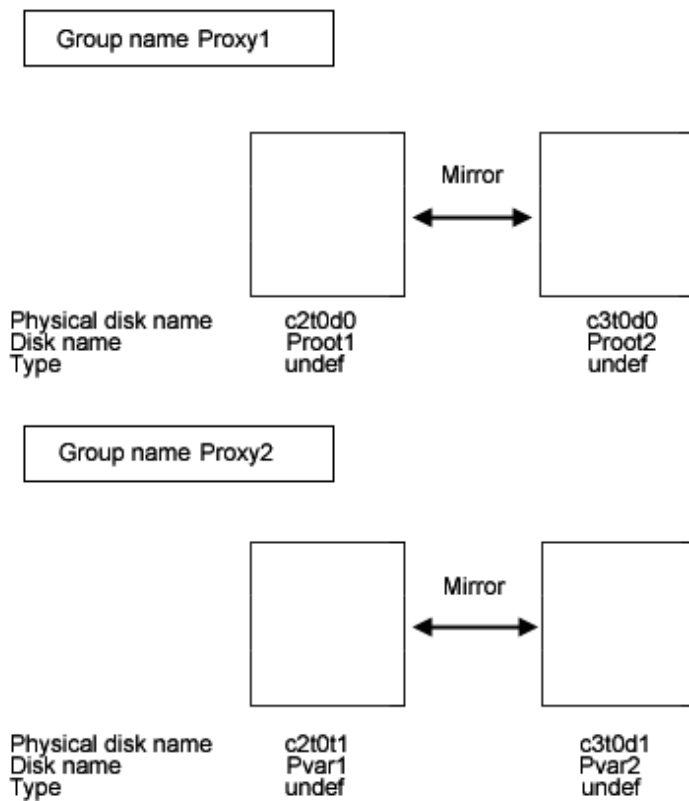
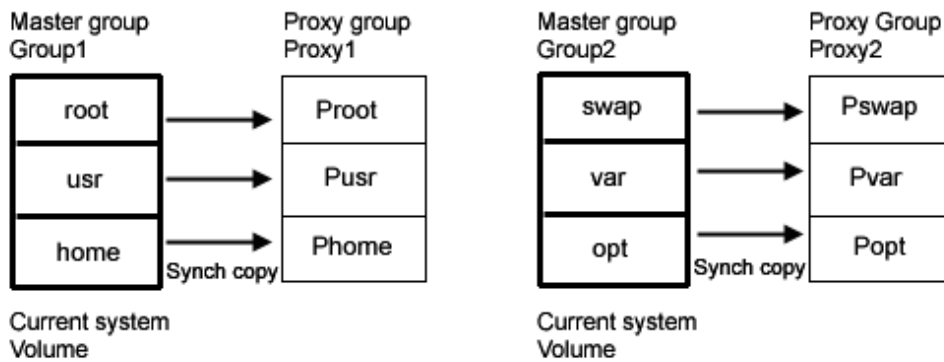


Figure 6.3 Joining Proxy Groups



 Note

Fujitsu Storage Systems ETERNUS SAN Boot Environment

When creating an alternative boot environment with the Advanced Copy function, disk areas for backup must be in a mirroring configuration with the multiplicity of one or be a single volume. If disk areas for backup are in a mirroring configuration with the multiplicity of two or more, the soft copy function is used instead of the Advanced Copy function. For details, see "[A.2.18 Using the Advanced Copy Function in a Proxy Configuration.](#)"

6.2.2 Summary of Backup

Snapshots of data in system disks must be collected in proxy groups during the system operation. Additionally, the environment allowing for booting through the proxy groups should be configured for errors due to a failed system disk and damaged data.

Figure 6.4 Backup

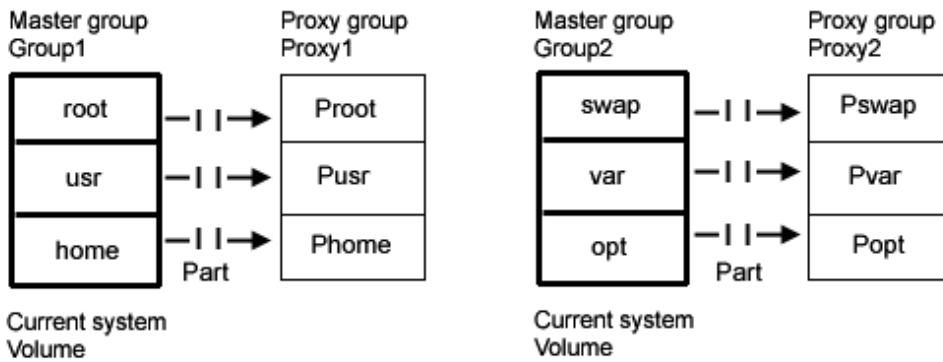



Figure 6.5 Backup Schedule



6.2.3 Summary of Restore

If booting the system becomes unavailable due to a failed system disk or damaged data, you can switch the environment to the alternative boot environment created on proxy groups to enable the system to continue operating. You can restore data back to the original system disks by joining the disk areas for backup as the master and the original system disks as the proxy.

Figure 6.6 Restore

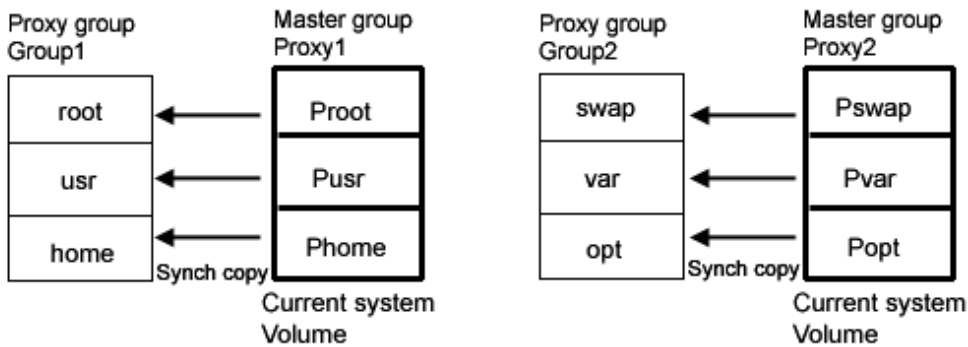
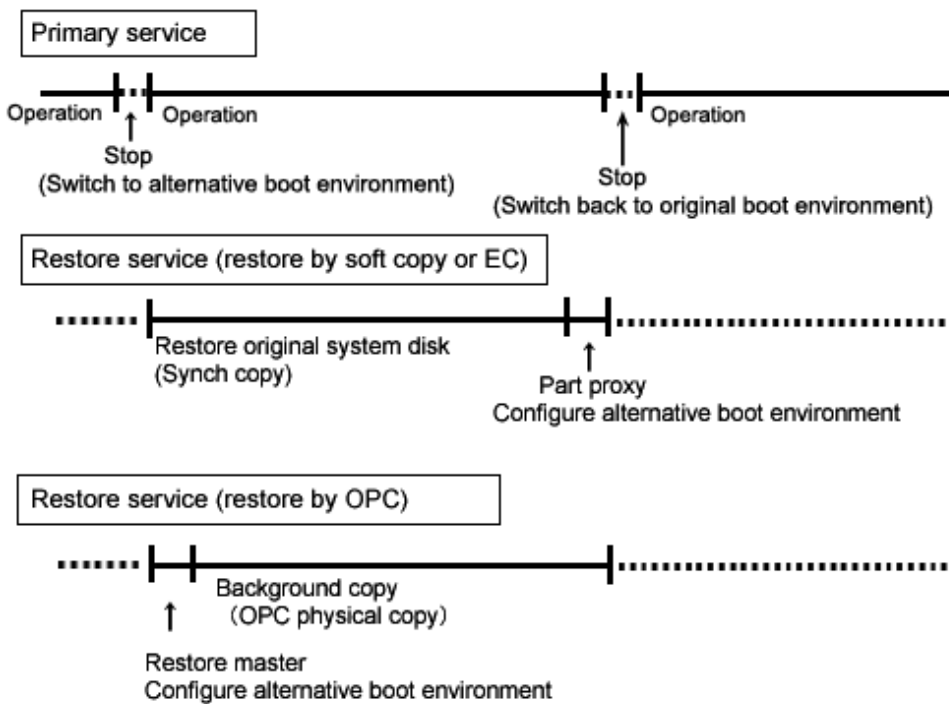


Figure 6.7 Restore Schedule



Information

Fujitsu Storage Systems ETERNUS SAN Boot Environment

When the OPC function is available, after switching to the alternative boot environment, it is possible to restore data from the proxy to the master without breaking the relationship between the master (original system disks) and the proxy (disk areas for backup). For conditions under which the OPC function is used, see "A.2.19 Instant Snapshot by OPC."

6.2.4 Summary of Procedure

Figure 6.8 Outline of the Configuration Procedure




Figure 6.9 Outline of the Backup Procedure

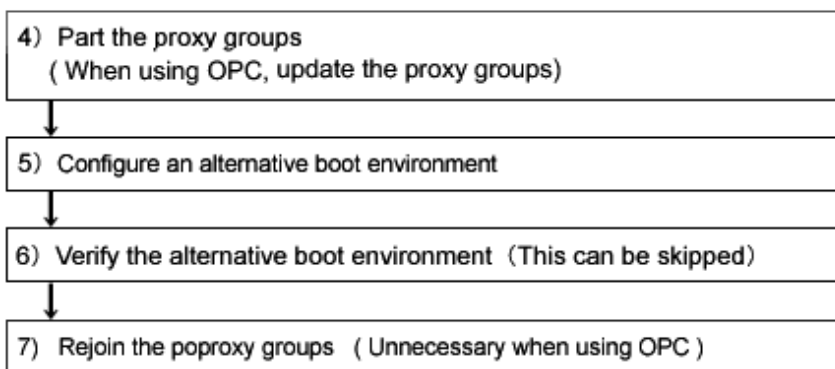



Figure 6.10 Outline of the Restore Procedure



6.2.5 Configuring an Environment

1) Mirroring system disks

In preparatory for alternative boot environment creation, mirror system disks. This sub-section describes the procedures for mirroring system disks in the configuration as shown in "[6.2.1 System Configuration](#)."



See

.....
For details on GDS Management View, see "[5.2.1 System Disk Settings](#)."
.....

1-1) Exit all active application programs.

To ensure safe mirroring, exit all running application programs. If higher security is required, you should make backups of system disks.

1-2) Register system disks with the root class.

```
# sdxdisk -M -c System -a type=root -d  
c0t0d0=Root1:keep,c1t0d0=Root2:undef,c0t0d1=Var1:keep,c1t0d1=Var2:undef
```

1-3) Connect system disks to groups respectively.

```
# sdxdisk -C -c System -g Group1 -d Root1,Root2 -v 0=root:on,1=usr:on,  
3=home:on  
# sdxdisk -C -c System -g Group2 -d Var1,Var2 -v 0=swap:on,1=var:on,3=opt:on
```

1-4) Confirm that the mirror definition is complete.

```
# sdxroot -M -c System -d Root1,Var1
```

1-5) Reboot the system.

```
# shutdown -y -g0 -i6
```

1-6) Confirm that mirroring is complete.

Use the mount(1M) command or the sdxinfo command to verify that the system disks have been mirrored properly.

2) Creating proxy groups

Create the disk areas (proxy groups) for backing up system disks. The following describes the procedure for creating proxy groups in the configuration as shown in ["6.2.1 System Configuration."](#)

2-1) Register the disks with the root class.

```
# sdxdisk -M -c System -d  
c2t0d0=Proot1,c3t0d0=Proot2,c2t0d1=Pvar1,c3t0d1=Pvar2
```



See

For details on GDS Management View, see ["5.4.1 Class Configuration."](#)

2-2) Connect the disks to groups respectively.

```
# sdxdisk -C -c System -g Proxy1 -d Proot1,Proot2  
# sdxdisk -C -c System -g Proxy2 -d Pvar1,Pvar2
```



See

For details on GDS Management View, see ["5.2.3.3 Group Configuration."](#)

3) Joining the proxy groups

Copy data in system disks into the backup disks by joining a group of the backup disks (proxy group) to a group of the system disks (master group). The following describes the procedure for joining proxy groups in the configuration as shown in ["6.2.1 System Configuration."](#)

3-1) Join the proxy groups.

```
# sdxproxy Join -c System -m Group1 -p Proxy1 -a  
root=Proot:on,usr=Pusr:on,home=Phome:on  
# sdxproxy Join -c System -m Group2 -p Proxy2 -a  
swap=Pswap:on,var=Pvar:on,opt=Popt:on
```




See

For details on GDS Management View, see "5.2.5.1 Join."

3-2) Confirm that synchronization copying is complete.

When using the OPC function, go on to step 3-3) without waiting until the synchronization copying is completed.

```
# sdxinfo -s -c System
OBJ    CLASS  GROUP  DISK    VOLUME  STATUS
-----
slice  System Group1 Root1    root    ACTIVE
slice  System Group1 Root2    root    ACTIVE
slice  System Group1 Root1    usr     ACTIVE
slice  System Group1 Root2    usr     ACTIVE
slice  System Group1 Root1    home   ACTIVE
slice  System Group1 Root2    home   ACTIVE
slice  System Group2 Var1     swap   ACTIVE
slice  System Group2 Var2     swap   ACTIVE
slice  System Group2 Var1     var    ACTIVE
slice  System Group2 Var2     var    ACTIVE
slice  System Group2 Var1     opt    ACTIVE
slice  System Group2 Var2     opt    ACTIVE
slice  System Proxy1 Proot1   Proot   STOP
slice  System Proxy1 Proot2   Proot   STOP
slice  System Proxy1 Proot1   Pusr    STOP
slice  System Proxy1 Proot2   Pusr    STOP
slice  System Proxy1 Proot1   Phome   COPY
slice  System Proxy1 Proot2   Phome   COPY
slice  System Proxy2 Pvar1   Pswap   STOP
slice  System Proxy2 Pvar2   Pswap   STOP
slice  System Proxy2 Pvar1   Pvar    COPY
slice  System Proxy2 Pvar2   Pvar    COPY
slice  System Proxy2 Pvar1   Popt    COPY
slice  System Proxy2 Pvar2   Popt    COPY
```

If synchronization copying is in process, COPY is displayed in the STATUS field for slice information of the copy destination proxy group. If the STATUS of every slice in the proxy group is STOP, the synchronization copying is completed.



Information

On the GDS Management View main screen, slices composing proxy volumes are indicated as below.

- If synchronization copying is in process, the status is "copy" and the icon color is blue.
- After synchronization copying is completed, the status is "stop" and the icon color is black.



Information

Fujitsu Storage Systems ETERNUS SAN Boot Environment

When using the OPC function, perform the following steps 3-3) through to 3-5), 4), and 5) without waiting until the synchronization copying is completed. For conditions under which the OPC function is used, see "A.2.19 Instant Snapshot by OPC."

3-3) Break the relationship between the master and the proxy once.

```
# sdxproxy Break -c System -p Proxy1 -e force
# sdxproxy Break -c System -p Proxy2 -e force
```

3-4) Set proxy volume access mode attributes to rw (read and write).

```
# sdxattr -V -c System -v Proot -a mode=rw
# sdxattr -V -c System -v Pusr -a mode=rw
# sdxattr -V -c System -v Phome -a mode=rw
# sdxattr -V -c System -v Pswap -a mode=rw
# sdxattr -V -c System -v Pvar -a mode=rw
# sdxattr -V -c System -v Popt -a mode=rw
```

3-5) Relate the master groups and the proxy groups again.

```
# sdxproxy Relate -c System -m Group1 -p Proxy1
# sdxproxy Relate -c System -m Group2 -p Proxy2
```

6.2.6 Backing Up

4) Parting the proxy groups

Once synchronization copying is completed, the master group and the proxy group become synchronized. Snapshots of a master group can be collected in a proxy group by parting the synchronized master group and proxy group.

4-1) Secure consistency of the file systems.

To secure consistency of snapshot file systems, restrain update of the file systems. However, file systems such as /, /usr, and /var are necessary for the system to operate, they cannot be unmounted during system operation. You can follow the procedure as below in order to reduce write occurrences to system disks and write occurrences to system disks that are not updated yet.

- a. Boot the system in single user mode. (This can be skipped.)
When creating backups using the OPC function in a Fujitsu Storage Systems ETERNUS SAN boot environment, do not perform this procedure.
- b. Exit all active application programs writing in the system disks. (This can be skipped.)
- c. Execute the sync(1M) command to write file system data updated in memory but not yet written to the disks.



Note

The system may hang up if a / file system is locked not to be written in with the lockfs(1M) command. In addition, GDS and GDS Snapshot commands hang up if a /var file system is locked not to be written in. Therefore, / and /var should not be applied the write lock.

Even if all the steps a., b., and c. are enforced, it is impossible to completely restrain update on the file system. As a result, inconsistency similar to that after a system panic might occur in a snapshot file system.

If a., b., and c. are all enforced, a snapshot file system will be similar to a file system after a panic occurs in single user mode.

If only c. is enforced skipping a. and b., a snapshot file system will be similar to a file system after a panic occurs during the system operation.

In any of these situations, a file system may have inconsistency, and the file system should be checked for consistency and repaired as described in step 5-2).

4-2) Part the proxy groups.

```
# sdxproxy Part -c System -p Proxy1,Proxy2
```



See

When using GDS Management View, see "Part" in "5.3.2.2 Backup (by Synchronization)."



Information

Fujitsu Storage Systems ETERNUS SAN Boot Environment

When the OPC function is available, after proxy groups are parted, use the following command to copy data from the master groups to the proxy groups. When using GDS Management View, see "Update" in "5.3.2.3 Backup (by OPC)." In the Update Proxy dialog box, select "Yes" to "Instant Snapshot."

```
# sdxproxy Update -c System -p Proxy1,Proxy2 -e instant
```

4-3) When the system was booted in single user mode in a. of step 4-1), reboot it in multi-user mode.

4-4) When application programs were exited in b. of step 4-1), launch the application programs.

5) Configuring an alternative boot environment

Make booting available through the proxy volumes in preparation for an error due to a failed system disk or damaged data.

5-1) Set the access mode attribute of the proxy volumes to rw (read and write).

If the access mode attribute of proxy volumes created in the proxy group is ro (read only), it must be changed to rw (read and write). The access mode attribute can be viewed in the MODE field output by the `sdxinfo -V -e long` command. If the access mode attribute is already set to rw (read and write), executing the following commands is not necessary.

```
# sdxvolume -F -c System -v Proot,Pusr,Phome,Pswap,Pvar,Popt
# sdxattr -V -c System -v Proot -a mode=rw
# sdxattr -V -c System -v Pusr -a mode=rw
# sdxattr -V -c System -v Phome -a mode=rw
# sdxattr -V -c System -v Pswap -a mode=rw
# sdxattr -V -c System -v Pvar -a mode=rw
# sdxattr -V -c System -v Popt -a mode=rw
```

5-2) Verify and repair the file systems on the proxy volumes.

There may be inconsistency in file systems on proxy volumes, and so verify and repair them using the `fsck(1M)` command.

```
# sdxvolume -N -c System -v Proot,Pusr,Phome,Pswap,Pvar,Popt
# fsck -y /dev/sfdsk/System/rdisk/Proot
# fsck -y /dev/sfdsk/System/rdisk/Pusr
# fsck -y /dev/sfdsk/System/rdisk/Phome
# fsck -y /dev/sfdsk/System/rdisk/Pvar
# fsck -y /dev/sfdsk/System/rdisk/Popt
```

5-3) Configure the alternative boot environment.

```
# sdxproxy Root -c System -p Proxy1,Proxy2
```

Once the alternative boot environment is configured, the following message is output.

```
SDX:sdxproxy: INFO: completed definitions of alternative boot environment:
current-boot-device=/pci@1f,4000/scsi@3/disk@0,0:a pci@1f,4000/scsi@3/
disk@1,0:a
```

```
alternative-boot-device=/pci@1f,4000/scsi@3/disk@2,0:a /pci@1f,4000/scsi@3/disk@3,0:a
```

Be sure to keep a copy of the output boot device names in the current boot environment (current-boot-device value) and in the alternative boot environment (alternative-boot-device value).

Information

Fujitsu Storage Systems ETERNUS SAN Boot Environment

When using the OPC function, confirm that the OPC physical copying is completed.

```
# sdxinfo -S -c System
OBJ    CLASS   GROUP   DISK    VOLUME  STATUS
-----
slice  System  Group1  Root1   root    ACTIVE
slice  System  Group1  Root2   root    ACTIVE
slice  System  Group1  Root1   usr     ACTIVE
slice  System  Group1  Root2   usr     ACTIVE
slice  System  Group1  Root1   home    ACTIVE
slice  System  Group1  Root2   home    ACTIVE
slice  System  Group2  Var1    swap    ACTIVE
slice  System  Group2  Var2    swap    ACTIVE
slice  System  Group2  Var1    var     ACTIVE
slice  System  Group2  Var2    var     ACTIVE
slice  System  Group2  Var1    opt     ACTIVE
slice  System  Group2  Var2    opt     ACTIVE
slice  System  Proxy1  Proot1  Proot   STOP
slice  System  Proxy1  Proot2  Proot   STOP
slice  System  Proxy1  Proot1  Pusr    STOP
slice  System  Proxy1  Proot2  Pusr    STOP
slice  System  Proxy1  Proot1  Phome   COPY
slice  System  Proxy1  Proot2  Phome   COPY
slice  System  Proxy2  Pvar1   Pswap   STOP
slice  System  Proxy2  Pvar2   Pswap   STOP
slice  System  Proxy2  Pvar1   Pvar    COPY
slice  System  Proxy2  Pvar2   Pvar    COPY
slice  System  Proxy2  Pvar1   Popt    COPY
slice  System  Proxy2  Pvar2   Popt    COPY
```

If the copying is in process, COPY is displayed in the STATUS field for slices of the copy destination proxy group. If the STATUS of every slice in the proxy group is STOP, the copying is completed.

5-4) Stop the proxy volumes.

To protect data in the alternative boot environment from illegal write access, the proxy volumes should be inactivated.

```
# sdxvolume -F -c System -v Proot,Pusr,Phome,Pswap,Pvar,Popt
```

Information

When using GDS Management View, select a proxy volume and execute [Operation]:[Stop Volume] in the Main Screen.

6) Verifying the alternative boot environment. (This can be skipped.)

Confirm that the system can be booted in the alternative boot environment.

6-1) Boot the system through the alternative boot environment.

Execute the boot command in the OpenBoot environment. For the boot command, you can specify one of the boot devices in the alternative boot environment output in the message as shown in step 5-3).

```
# shutdown -y -g0 -i0
ok boot /pci@1f,4000/scsi@3/disk@2,0:a
```

Information

The boot command may fail as below if the number of the boot device names in the alternative boot environment output in step 5-3) exceeds the number of boot devices in the alternative boot environment.

```
ok boot /pci@1f,4000/scsi@3/disk@2,0:a
Rebooting with command: boot /pci@1f,4000/scsi@3/disk@2,0:a
Boot device: /pci@1f,4000/scsi@3/disk@2,0:a File and args:
Evaluating: boot /pci@1f,4000/scsi@3/disk@2,0:a

Can't open boot device

ok
```

If this happens, you should specify another boot device name in the alternative boot environment output in step 5-3) and try the boot command again.

6-2) Confirm that it was booted normally.

Using the mount(IM) command or the sdxinfo command, make sure that it was booted normally in the alternative boot environment and that GDS objects do not contain errors. Additionally, according to need, you should confirm that file system contents in the alternative boot environment are correct and that applications can normally run.

6-3) Return to the original boot environment.

Execute the boot command in the OpenBoot environment. For the boot command, you can specify one of the boot devices in the current boot environment output in the message as shown in step 5-3).

```
# shutdown -y -g0 -i0
ok boot /pci@1f,4000/scsi@3/disk@0,0:a
```

Information

The boot command may fail as below if the number of the boot device names in the current boot environment output in step 5-3) exceeds the number of boot devices in the current boot environment.

```
ok boot /pci@1f,4000/scsi@3/disk@0,0:a
Rebooting with command: boot /pci@1f,4000/scsi@3/disk@0,0:a
Boot device: /pci@1f,4000/scsi@3/disk@0,0:a File and args:
Evaluating: boot /pci@1f,4000/scsi@3/disk@0,0:a

Can't open boot device

ok
```

If this happens, you should specify another boot device name in the current boot environment output in step 5-3) and try the boot command again.

6-4) Stop the proxy volumes.

To protect data in the alternative boot environment from illegal write access, the proxy volumes should be inactivated.

```
# sdxvolume -F -c System -v Proot,Pusr,Phome,Pswap,Pvar,Popt
```

Information

When using GDS Management View, select a proxy volume and execute [Operation]:[Stop Volume] in the Main Screen.

7) Rejoining the proxy groups

To back up the system disks again, copy contents in the system disks to the backup disks again by rejoining a group of the backup disks (proxy group) to a group of the system disks (master group).

Note

Copying by rejoining is finished quickly since only the updated blocks in the master and the proxy are copied through the just resynchronization mechanism (JRM). However, if the system is rebooted after the proxy is parted, the JRM is disabled and the entire volumes will be copied when the proxy is rejoined. Therefore, copying of the entire volumes is conducted instead of just resynchronization when the proxy groups are rejoined in step 7-1) if the system is rebooted after step 4-3) or step 6).

Information

Fujitsu Storage Systems ETERNUS SAN Boot Environment

When using the OPC function, do not rejoin the proxy groups.

7-1) Rejoin the proxy groups.

```
# sdxproxy Rejoin -c System -p Proxy1,Proxy2
```

See

When using GDS Management View, see "Rejoin" in "5.3.2.2 Backup (by Synchronization)."

7-2) Confirm that synchronization copying is complete.

```
# sdxinfo -S -c System
OBJ    CLASS  GROUP  DISK    VOLUME  STATUS
-----
slice  System Group1  Root1   root    ACTIVE
slice  System Group1  Root2   root    ACTIVE
slice  System Group1  Root1   usr     ACTIVE
slice  System Group1  Root2   usr     ACTIVE
slice  System Group1  Root1   home    ACTIVE
slice  System Group1  Root2   home    ACTIVE
slice  System Group2  Var1    swap    ACTIVE
slice  System Group2  Var2    swap    ACTIVE
slice  System Group2  Var1    var     ACTIVE
slice  System Group2  Var2    var     ACTIVE
slice  System Group2  Var1    opt     ACTIVE
slice  System Group2  Var2    opt     ACTIVE
slice  System Proxy1  Proot1  Proot   STOP
slice  System Proxy1  Proot2  Proot   STOP
slice  System Proxy1  Proot1  Pusr    STOP
```

slice	System	Proxy1	Proot2	Pusr	STOP
slice	System	Proxy1	Proot1	Phome	COPY
slice	System	Proxy1	Proot2	Phome	COPY
slice	System	Proxy2	Pvar1	Pswap	STOP
slice	System	Proxy2	Pvar2	Pswap	STOP
slice	System	Proxy2	Pvar1	Pvar	COPY
slice	System	Proxy2	Pvar2	Pvar	COPY
slice	System	Proxy2	Pvar1	Popt	COPY
slice	System	Proxy2	Pvar2	Popt	COPY

If synchronization copying is in process, COPY is displayed in the STATUS field for slices of the copy destination proxy group. If the STATUS of every slice in the proxy group is STOP, the synchronization copying is completed.

 **Information**

On the GDS Management View main screen, slices composing proxy volumes are indicated as below.

- If synchronization copying is in process, the status is "copy" and the icon color is blue.
- After synchronization copying is completed, the status is "stop" and the icon color is black.

7-3) Part the proxy groups, configure the alternative boot environment, and confirm that the alternative boot environment is valid following the step from 4) to 6).

6.2.7 Restoring

8) Switching to the alternative boot environment

If the system cannot be booted due to a failed system disk or damaged data, switch the environment to the alternative boot environment created in the proxy volume to allow the system to continue operating.

8-1) Boot the system through the alternative boot environment.

Execute the boot command in the OpenBoot environment. For the boot command, you can specify one of the boot devices in the alternative boot environment output in the message as shown in step 5-3).

```
# shutdown -y -g0 -i0
ok boot /pci@1f,4000/scsi@3/disk@2,0:a
```

 **Information**

The boot command may fail as below if the number of the boot device names in the alternative boot environment output in step 5-3) exceeds the number of boot devices in the alternative boot environment.

```
ok boot /pci@1f,4000/scsi@3/disk@2,0:a
Rebooting with command: boot /pci@1f,4000/scsi@3/disk@2,0:a
Boot device: /pci@1f,4000/scsi@3/disk@2,0:a File and args:
Evaluating: boot /pci@1f,4000/scsi@3/disk@2,0:a

Can't open boot device

ok
```

If this happens, you should specify another boot device name in the alternative boot environment output in step 5-3) and try the boot command again.

8-2) Confirm that it was booted normally.

Using the mount(IM) command or the sdxinfo command, make sure that it was booted normally in the alternative boot environment and that GDS objects do not contain errors. Additionally, according to need, you should confirm that file system contents in the alternative boot environment are correct and that applications can normally run.

8-3) Break up the former boot environment according to need.

To break up the former boot environment, break the master and proxy relationships, remove the master volumes, and remove groups and disks from the master groups as follows. You may not perform this procedure when restoring system disks in step 9).

```
# sdxproxy Break -c System -p Proxy1
# sdxproxy Break -c System -p Proxy2
# sdxvolume -F -c System -v root,usr,home,swap,var,opt
# sdxvolume -R -c System -v root
# sdxvolume -R -c System -v usr
# sdxvolume -R -c System -v home
# sdxvolume -R -c System -v swap
# sdxvolume -R -c System -v var
# sdxvolume -R -c System -v opt
# sdxgroup -R -c System -g Group1
# sdxgroup -R -c System -g Group2
# sdxdisk -R -c System -d Root1
# sdxdisk -R -c System -d Root2
# sdxdisk -R -c System -d Var1
# sdxdisk -R -c System -d Var2
```



When using GDS Management View, see "[5.5.6 Breaking a Proxy](#)," "[5.5.2 Removing a Volume](#)," "[5.5.3 Removing a Group](#)," and "[5.4.1 Class Configuration](#)."

9) Restoring system disks

After rebooting the system in the alternative boot environment, restore backup disk data back to the original system disks.

9-1) Cancel the master and proxy relationship.

```
# sdxproxy Break -c System -p Proxy1
# sdxproxy Break -c System -p Proxy2
```



When using GDS Management View, see "[5.5.6 Breaking a Proxy](#)."



Fujitsu Storage Systems ETERNUS SAN Boot Environment

When using the OPC function, do not break the master and proxy relationship. However, the original system disks are not in a mirroring configuration and need replacement, break the relationship.

9-2) Remove the master volumes.

```
# sdxvolume -F -c System -v root,usr,home,swap,var,opt
# sdxvolume -R -c System -v root
```



```
# sdxvolume -R -c System -v usr
# sdxvolume -R -c System -v home
# sdxvolume -R -c System -v swap
# sdxvolume -R -c System -v var
# sdxvolume -R -c System -v opt
```



See

When using GDS Management View, see "5.5.2 Removing a Volume."



Information

Fujitsu Storage Systems ETERNUS SAN Boot Environment

When using the OPC function, only stop master volumes (with sdxvolume -F) and do not remove the master volumes (with sdxvolume -R). However, the original system disks are not in a mirroring configuration and need replacement, remove the master volumes.

9-3) If an original system disk crashed, swap the failed disk.

The following is an example of swapping disk Root1 (physical disk c0t0d0).



See

When using GDS Management View, see "5.3.4 Disk Swap."

9-3-1) Exclude the disk to be swapped from the GDS management to make it exchangeable.

```
# sdxswap -O -c System -d Root1
```

9-3-2) Swap physical disk c0t0d0.

9-3-3) Include the swapped disk into the GDS management.

```
# sdxswap -I -c System -d Root1
```

9-4) Join a group of the backup disks as the master and a group of the original system disks as the proxy.

```
# sdxproxy Join -c System -m Proxy1 -p Group1 -a
Proot=root:on,Pusr=usr:on,Phome=home:on
# sdxproxy Join -c System -m Proxy2 -p Group2 -a
Pswap=swap:on,Pvar=var:on,Popt=opt:on
```



See

When using GDS Management View, see "5.2.5.1 Join."



Information

Fujitsu Storage Systems ETERNUS SAN Boot Environment

When the relationship between the master and the proxy was not canceled, by using the following command, it is possible to restore data from the proxy groups (backup disk groups) to the master groups (original system disk groups) with the OPC function. When using GDS

Management View, see "Restore" in "5.3.3 Restore." In the Restore Master dialog box, select "No" to "Rejoin" and "Yes" to "Instant Restore."

```
# sdxproxy Restore -c System -p Proxy1,Proxy2 -e instant
```

9-5) Confirm that synchronization copying is complete.

```
# sdxinfo -S -c System
OBJ      CLASS  GROUP  DISK    VOLUME  STATUS
-----
slice    System Group1  Root1   root    STOP
slice    System Group1  Root2   root    STOP
slice    System Group1  Root1   usr     STOP
slice    System Group1  Root2   usr     STOP
slice    System Group1  Root1   home    COPY
slice    System Group1  Root2   home    COPY
slice    System Group2  Var1    swap    STOP
slice    System Group2  Var2    swap    STOP
slice    System Group2  Var1    var     COPY
slice    System Group2  Var2    var     COPY
slice    System Group2  Var1    opt     COPY
slice    System Group2  Var2    opt     COPY
slice    System Proxy1  Proot1  Proot   ACTIVE
slice    System Proxy1  Proot2  Proot   ACTIVE
slice    System Proxy1  Proot1  Pusr    ACTIVE
slice    System Proxy1  Proot2  Pusr    ACTIVE
slice    System Proxy1  Proot1  Phome   ACTIVE
slice    System Proxy1  Proot2  Phome   ACTIVE
slice    System Proxy2  Pvar1   Pswap   ACTIVE
slice    System Proxy2  Pvar2   Pswap   ACTIVE
slice    System Proxy2  Pvar1   Pvar    ACTIVE
slice    System Proxy2  Pvar2   Pvar    ACTIVE
slice    System Proxy2  Pvar1   Popt    ACTIVE
slice    System Proxy2  Pvar2   Popt    ACTIVE
```

If synchronization copying is in process, COPY is displayed in the STATUS field for slices of the copy destination proxy group. If the STATUS of every slice in the proxy group is STOP, the synchronization copying is completed.

Information

On the GDS Management View main screen, slices composing copy destination volumes are indicated as below.

- If synchronization copying is in process, the status is "copy" and the icon color is blue.
- After synchronization copying is completed, the status is "stop" and the icon color is black.

Note

Fujitsu Storage Systems ETERNUS SAN Boot Environment

In a situation where the multiplicity of master volumes is two or more, if restore is performed with the OPC function, the master volume slices other than the OPC copy destinations will be invalid. To restore the mirror status of the master volumes, execute the resynchronization copying of the master volumes. If the boot environment is switched to the original boot environment in step 9-6), the master volumes are rebooted and the resynchronization copying is automatically executed.

11) Unmirroring system disks in an alternative boot environment



See

When using GDS Management View, see "5.5.5 Unmirroring the System Disk."

11-1) Exit all active application programs.

To ensure safe mirroring cancellation, exit all running application programs. If higher security is required, you should make backups of system disks.

11-2) Remove those disks not used as system disks after canceling the mirroring.

```
# sdxdisk -D -c System -g Proxy1 -d Proot2
# sdxdisk -D -c System -g Proxy2 -d Pvar2
```

11-3) Confirm that cancellation of the mirroring is complete.

```
# sdxroot -R -c System -d Proot1,Pvar1
```

11-4) Reboot the system.

```
# shutdown -y -g0 -i6
```

11-5) Verify that the mirroring was cancelled normally.

Using the mount(IM) command or the sdxinfo command, verify that the system disk mirroring was cancelled properly.

11-6) Cancel the system disk management.

```
# sdxvolume -F -c System -v Proot,Pusr,Phome,Pswap,Pvar,Popt
# sdxvolume -R -c System -v Proot
# sdxvolume -R -c System -v Pusr
# sdxvolume -R -c System -v Phome
# sdxvolume -R -c System -v Pswap
# sdxvolume -R -c System -v Pvar
# sdxvolume -R -c System -v Popt
# sdxgroup -R -c System -g Proxy1
# sdxgroup -R -c System -g Proxy2
# sdxdisk -R -c System -d Proot1
# sdxdisk -R -c System -d Proot2
# sdxdisk -R -c System -d Pvar1
# sdxdisk -R -c System -d Pvar2
```

6.3 Backing Up and Restoring Local Disks and Shared Disks

This section discusses the methods of backing up and restoring local disks and shared disks in a system where GDS Snapshot has not been installed. Among volumes in the root class, volumes (e.g. /opt, /home) other than system volumes (/ , /usr, /var, swap area) can also be backed up and restored following the procedures described here.

The following is an example of backing up and restoring volume Volume1 in class Class1.

6.3.1 Offline Backup

1) Stopping the services

1a) With a shared volume used in a cluster application

1a-1) Exit all cluster applications.

1a-2) Activate the volume on a node on which backup is conducted.

```
# sdxvolume -N -c Class1 -v Volume1
```

1b) With a volume not used by a cluster application

1b-1) Stop the services using the volume.

1b-2) When the volume is used as a file system, unmount the file system. In the following example, the mount point is /mnt1.

```
# cd /  
# umount /mnt1
```

2) Backing Up

Back up volume data. In the following examples, data is backed up to a tape medium of tape device /dev/rmt/0.



See

For details on the backup methods, see the manuals of file systems to be backed up and used commands.

- When backing up data held in a raw device with the dd(1M) command

```
# dd if=/dev/sfdsk/Class1/rdsk/Volume1 of=/dev/rmt/0 bs=32768
```

- When backing up a ufs file system with the ufsdump(1M) command

```
# ufsdump 0ucf /dev/rmt/0 /dev/sfdsk/Class1/rdsk/Volume1
```

- When backing up the GFS Local File System with the sfxdump(1M) command

```
# sfxdump 0ucf /dev/rmt/0 /dev/sfdsk/Class1/rdsk/Volume1
```

3) Resuming the services

3a) With a shared volume used in a cluster application

3a-1) Inactivate the volume on the node where backup was conducted.

```
# sdxvolume -F -c Class1 -v Volume1
```

3a-2) Launch cluster applications.

3b) With a volume not used by a cluster application

3b-1) When the volume is used as a file system, mount the file system. The following shows examples when the mount point is /mnt1.

- For a ufs file system

```
# mount -F ufs /dev/sfdsk/Class1/dsk/Volume1 /mnt1
```

- For the GFS Local File System

```
# mount -F sfxfs /dev/sfdsk/Class1/dsk/Volume1 /mnt1
```

3b-2) Resume the services.

6.3.2 Online Backup (by Slice Detachment)

For mirror volumes, data can be backed up through snapshot by slice detachment during the services operation.



See

For use conditions on snapshots by detaching a slice, see "[A.2.12 Creating a Snapshot by Slice Detachment](#)."

To secure consistency of data in a detached slice, the services must be stopped temporarily when detaching the slice.



See

For securing consistency of snapshot data, see "[A.2.22 Ensuring Consistency of Snapshot Data](#)."

1) Stop the services

1a) With a shared volume used by a cluster application

Exit the cluster application.

1b) With a volume not used by a cluster application

1b-1) Stop the services using the volume.

1b-2) When the volume is used as a file system, unmount the file system. In the following example, the mount point is /mnt1.

```
# cd /  
# umount /mnt1
```

2) Detaching the slice

Detach the slice from the mirror volume. The following shows an example of detaching the slice from disk Disk1.

```
# sdxslice -M -c Class1 -d Disk1 -v Volume1
```



See

When using GDS Management View, see "[Slice Detachment](#)" in "[5.3.2.1 Backup \(by Slice Detachment\)](#)."

3) Resuming the services

3a) With a shared volume used by a cluster application

Execute the cluster application.

3b) With a volume not used by a cluster application

3b-1) When the volume is used as a file system, mount the file system. In the following example, the mount point is /mnt1.

- For a ufs file system

```
# mount -F ufs /dev/sfdsk/Class1/dsk/Volume1 /mnt1
```

- For the GFS Local File System

```
# mount -F sfxfs /dev/sfdsk/Class1/dsk/Volume1 /mnt1
```

3b-2) Resume the services.

4) When the volume is used as the GFS Local File System, reconfigure the partition information for the GFS Local File System on the detached slice.

```
# sfxadm /dev/sfdsk/Class1/rdsk/Disk1.Volume1
```

5) When the volume is used as a file system, check and repair consistency of the file system. If the file system was unmounted when the slice was detached in step 1), this step can be skipped.

- For a ufs file system

```
# fsck -F ufs -y /dev/sfdsk/Class1/rdsk/Disk1.Volume1
```

- For the GFS Local File System

```
# fsck -F sfxfs -y /dev/sfdsk/Class1/rdsk/Disk1.Volume1
```

6) Backing Up

Back up data in the detached slice. In the following examples, data is backed up to a tape medium of tape device /dev/rmt/0.



See

For details on the backup methods, see the manuals of file systems to be backed up and used commands.

- When backing up data held in a raw device with the dd(1M) command

```
# dd if=/dev/sfdsk/Class1/rdsk/Disk1.Volume1 of=/dev/rmt/0 bs=32768
```

- When backing up a ufs file system with the ufsdump(1M) command

```
# ufsdump 0ucf /dev/rmt/0 /dev/sfdsk/Class1/rdsk/Disk1.Volume1
```

- When backing up the GFS Local File System with the sfxdump(1M) command

```
# sfxdump 0ucf /dev/rmt/0 /dev/sfdsk/Class1/rdsk/Disk1.Volume1
```

7) Reattaching the slice

Reattach the detached slice to the mirror volume.

```
# sdxslice -R -c Class1 -d Disk1 -v Volume1
```



See

When using GDS Management View, see "Attach Slice" of "5.3.2.1 Backup (by Slice Detachment)."

6.3.3 Restoring

1) Stopping the services

1a) With a shared volume used by a cluster application.

Exit the cluster application.

1b) With a volume not used by a cluster application

1b-1) Stop the services using the volume.

1b-2) When the volume is used as a file system, unmount the file system. In the following example, the mount point is /mnt1.

```
# cd /
# umount /mnt1
```

2) Restoring

Restore the volume data. The following shows an example of restoring data from a tape medium of tape device /dev/rmt/0.



See

For details on the restore methods, see the manuals of file systems to be restored and used commands.

- When restoring data held in a raw device with the dd(1M) command

```
# dd if=/dev/rmt/0 of=/dev/sfdsk/Class1/rdisk/Volume1 bs=32768
```

- When restoring a ufs file system with the ufsrestore(1M) command

```
# newfs /dev/sfdsk/Class1/rdisk/Volume1
# mount -F ufs /dev/sfdsk/Class1/dsk/Volume1 /mnt1
# cd /mnt1
# ufsrestore rvf /dev/rmt/0
# rm /mnt1/restoresymtable
# cd /
# umount /mnt1
```

- When restoring the GFS Local File System with the sfxrestore(1M) command

```
# sfxnewfs /dev/sfdsk/Class1/rdisk/Volume1
# mount -F sfxfs /dev/sfdsk/Class1/dsk/Volume1 /mnt1
# cd /mnt1
# sfxrestore rvf /dev/rmt/0
# rm /mnt1/restoresymtable
# cd /
# umount /mnt1
```

3) When the volume is used as the GFS Local File System, add the node configuration information of the GFS Local File System on a node where restore was not conducted.

3-1) Inactivate the volume on the node where restore was conducted.

```
# sdxvolume -F -c Class1 -v Volume1
```

3-2) Activate the volume on a node where restore was not conducted.

```
# sdxvolume -N -c Class1 -v Volume1
```

3-3) Add the node configuration information of the GFS Local File System on the node where restore was not conducted.

```
# sfxnode /dev/sfdsk/Class1/rdsk/Volume1
```

3-4) Inactivate the volume on the node where restore was not conducted.

```
# sdxvolume -F -c Class1 -v Volume1
```

3-5) Restart the volume on the node where the restore was performed.

```
# sdxvolume -N -c Class1 -v Volume1
```

4) Resuming the services

4a) With a shared volume used by a cluster application

Execute the cluster application.

4b) With a volume not used by a cluster application

4b-1) When the volume is used as a file system, mount the file system. In the following example, the mount point is /mnt1.

- For a ufs system

```
# mount -F ufs /dev/sfdsk/Class1/dsk/Volume1 /mnt1
```

- For a GFS Local File System

```
# mount -F sfxfs /dev/sfdsk/Class1/dsk/Volume1 /mnt1
```

4b-2) Resume the services.

6.4 Online Backup and Instant Restore through Proxy Volume

This subsection describes the procedures for online backup and instant restore of local volumes and shared volumes through use of snapshots by GDS Snapshot proxy volumes.

There are two online backup methods: one uses the function as "1.5.1 Snapshot by Synchronization" and the other uses the function as "1.5.3 Instant Snapshot by OPC." By cooperation with the Advanced Copy function (EC or OPC) of the Fujitsu Storage Systems ETERNUS, or a copy function (EMC TimeFinder or EMC SRDF) of EMC's Symmetrix storage unit, backup and restore processing that does not impose the load on the server or SAN becomes possible.

The following table shows backup and restore methods and their features.

Operation	Disk Unit's Copy Function	Master Configuration (*1)	Proxy Configuration (*1)	Feature
Backup(Sync)	-	single mirror stripe concat	single mirror stripe concat	<ul style="list-style-type: none"> - With synchronized masters and proxies, snapshots are created simply by parting the proxies. - By use of the JRM for proxies, high-speed resynchronization is executable for snapshot re-creation. - Striping and concatenating masters are possible. - Mirroring, striping and concatenating proxies are possible.
	EC REC TimeFinder SRDF	single mirror	single	<ul style="list-style-type: none"> - With synchronized masters and proxies, snapshots are created simply by parting the proxies. - By use of a disk unit's copy function, snapshot creation that does not impose the load on the server or SAN is possible. - By use of a disk unit's incremental copy function, high-speed resynchronization that does not impose the load on the server or the SAN is performed for snapshot re-creation.
	OPC ROPC	single mirror	single	<ul style="list-style-type: none"> - With synchronized masters and proxies, snapshots are created simply by parting the proxies. - By use of a disk unit's copy function, snapshot creation that does not impose the load on the server or SAN is possible.
Backup(OPC)	OPC ROPC	single mirror	single	<ul style="list-style-type: none"> - There is no need to synchronize masters and proxies before creating snapshots, and backups can be created any time without scheduling in advance. - By use of a disk unit's copy function, snapshot creation that does not impose the load on the server or SAN is possible.
Restore	-	single mirror stripe concat	single mirror stripe concat	<ul style="list-style-type: none"> - By use of the JRM for proxies, high-speed restore is possible. - Striping and concatenating masters are possible. - Striping and concatenating proxies are possible.
	OPC ROPC	single mirror	single mirror	<ul style="list-style-type: none"> - By use of a disk unit's copy function, restore that does not impose the server or SAN is possible.

(*1) In the table above, configurations indicate objects as follows.

Configuration	Description
single	This configuration means any one of the following objects. <ul style="list-style-type: none"> - A single volume created in a single disk (excepting when using TimeFinder or SRDF) - A mirror group not hierarchized, with the multiplicity of one - A volume created in a mirror group not hierarchized, with the multiplicity of one (excepting when using TimeFinder or SRDF)
mirror	This configuration means either of the following objects. <ul style="list-style-type: none"> - A mirror group not hierarchized, with the multiplicity of two and higher - A mirror volume created in a mirror group not hierarchized, with the multiplicity of two and higher (excepting when using TimeFinder or SRDF)
stripe	This configuration means either of the following objects. <ul style="list-style-type: none"> - A mirror group to which one or more stripe groups are connected - A volume created in a mirror group to which one or more stripe groups are connected
concat	This configuration means either of the following objects. <ul style="list-style-type: none"> - A mirror group to which one or more concatenation groups are connected - A volume created in a mirror group to which one or more concatenation groups are connected

For details on backup and restore by cooperation with copy functions of disk units (EC, OPC, TimeFinder and SRDF), see "[1.5.2 Snapshot Function with No Load to the Server/SAN](#)," "[A.2.18 Using the Advanced Copy Function in a Proxy Configuration](#)," "[A.2.19 Instant Snapshot by OPC](#)" and "[A.2.21 Using EMC TimeFinder or EMC SRDF in a Proxy Configuration](#)."

6.4.1 Online Backup (by Synchronization)

In a system employing GDS Snapshot, backups can be created using snapshots from proxy volumes during the services operation.


However, to secure consistency of snapshot data, the service must be stopped temporarily when creating a snapshot (when parting a proxy volume from a master volume).



See

.....
 For securing consistency of snapshot data, see "[A.2.22 Ensuring Consistency of Snapshot Data](#)."

Figure 6.11 Backup Schedule



Note

Instant Restore

If an error occurs in master volume data while the master volume and a proxy volume are joined, the same error occurs in data of the proxy volume and instant restore becomes unavailable. If this happens, you must restore data from tape. After you execute online backup by parting a proxy volume, you are recommended to leave the proxy volume parted until just before re-executing online backup.

Note

Snapshot through Disk Unit's Copy Functions

If groups are hierarchized, or if proxies are in mirroring configuration with the multiplicity of two and higher, copy functions of disk units are unavailable for copying data from masters to proxies. For details, see "[A.2.18 Using the Advanced Copy Function in a Proxy Configuration](#)" and "[A.2.21 Using EMC TimeFinder or EMC SRDF in a Proxy Configuration](#)."

Note

When Using EMC TimeFinder or EMC SRDF

- In the procedures described below, volumes are used for snapshot operations. However, when using TimeFinder or SRDF, use groups for similar operations.
- When using TimeFinder, if standard devices that compose the master group and BCV devices that compose the proxy group are established, cancel the BCV pairs prior to joining the master and the proxy.
- When using SRDF, have source (R1) devices that compose the master group and target (R2) devices that compose the proxy group split prior to joining the master and the proxy.
- After relating the master and the proxy, do not perform TimeFinder or SRDF operations for BCV pairs or SRDF pairs that compose the master and proxy using the SYMCLI command and so on before canceling the relation.

Procedure

- 1) Joining a proxy volume

To prepare for creating a snapshot, relate and join proxy volume Volume2 as the copy destination of master volume Volume1 to the master volume. This example shows the procedure when Volume1 and Volume2 belong to Class1. Execute the following commands on a node that belongs to the scope of class Class1.

1-1) Stop proxy volume Volume2. If Class1 is a shared class, specify the "-e allnodes" option to stop Volume2 on all nodes.

```
# sdxvolume -F -c Class1 -v Volume2
```

Information

When using GDS Management View, select a proxy volume and execute [Operation]:[Stop Volume] in the Main Screen.

1-2) Relate and join proxy volume Volume2 to master volume Volume1.

```
# sdxproxy Join -c Class1 -m Volume1 -p Volume2
```

After returning from the command, synchronization copying from Volume1 to Volume2 is executed.

Information

Relating and Joining a Pair of Groups

If the proxy group includes volumes, remove the volumes before executing the sdxproxy Join command, and also specify the -a option for this command.

Example)

Relate and join proxy group Group2 to master group Group1.

Assign volume names Proxy1 and Proxy2 to the proxy volumes automatically created in Group2 corresponding to volumes Volume1 and Volume2 in Group1.

```
# sdxproxy Join -c Class1 -m Group1 -p Group2 \  
-a Volume1=Proxy1:on,Volume2=Proxy2:on
```

See

When using GDS Management View, see "5.2.5.1 Join."

Information

GFS Local File System in Multi-partition Configuration

If multiple volumes constitute one GFS Local File System, relate and join the proxy volume to all of the volumes.

2) Confirming the completion of copying

Confirm that the synchronization copying is complete.

```
# sdxinfo -S -c Class1 -o Volume2  
OBJ    CLASS  GROUP  DISK    VOLUME  STATUS  
-----  
slice  Class1  Group2  Disk3   Volume2  STOP  
slice  Class1  Group2  Disk4   Volume2  STOP
```

If all the displayed slices' STATUS fields are "STOP", synchronization copying is complete.
If the synchronization copying is still in progress, "COPY" will be displayed in the STATUS field.

3) Stopping the services

To secure consistency of snapshot data, stop the services before creating a snapshot and restrain the master volume from being written in.

3a) When the master volume is being used for a cluster application

Inactivate the cluster application.

3b) When the master volume is not being used for a cluster application

3b-1) Stop the services for which the master volume is being used.

3b-2) When the master volume is used as a file system, use either the following methods to prevent write access to the file system. In the following examples, the mount point is /DATA.

- To unmount the file system

```
# cd /  
# umount /DATA
```

- To write-lock the file system

```
# lockfs -w /DATA
```

4) Parting the proxy volume

Create a snapshot of master volume Volume1 by parting proxy volume Volume2 from master volume Volume1. Execute the following commands on a node that belongs to the scope of class Class1.

```
# sdxproxy Part -c Class1 -p Volume2
```



See

When using GDS Management View, see "Part" in "5.3.2.2 Backup (by Synchronization)." For setting of the part environment, select "No" to "Instant Snapshot."



Information

GFS Local File System in Multi-partition Configuration

Part proxy volumes for all the volumes that constitute the file system.

5) Resuming the services

5a) When the master volume is used for a cluster application

Activate the cluster application.

5b) When the master volume is not used for a cluster application

5b-1) When the master volume is used as a file system, use either the following methods to permit write access from applications to the file system on master volume Volume1. In the following examples, the mount point is /DATA.

- When the file system was unmounted in step 3b-2), mount the file system.
- For the ufs file system

```
# mount -F ufs /dev/sfdsk/Class1/dsk/Volume1 /DATA
```

- For the GFS Local File System

```
# mount -F sfxfs /dev/sfdsk/Class1/dsk/Volume1 /DATA
```

Information

GFS Local File System in Multi-partition Configuration

For the mount(1M) command, specify a device special file name of the representative partition.

- When write access to the file system was write-locked in step 3b-2), unlock the file system.

```
# lockfs -u /DATA
```

5b-2) Start the services for which the master volume is used.

6) Backing up to tape

Back up the snapshot data on the proxy volume to tape. Execute the following commands on a node that belongs to the scope of class Class1.

See

For details on the backup methods, see the manuals of file systems to be backed up and used commands.

6a) When backing up data held in a raw device with the dd(1M) command

```
# dd if=/dev/sfdsk/Class1/rdsk/Volume2 of=/dev/rmt/0 bs=32768
```

6b) When backing up the ufs file system with the tar(1) command

6b-1) Before mounting

Check and repair consistency of the ufs file system on the proxy volume with the fsck(1M) command. When the file system on the master volume was unmounted in step 3b-2), skip this step.

```
# fsck -F ufs /dev/sfdsk/Class1/rdsk/Volume2
```

6b-2) Mounting the snapshot

Mount the ufs file system on proxy volume Volume2 on temporary mount point /DATA_backup.

```
# mkdir /DATA_backup
# mount -F ufs /dev/sfdsk/Class1/dsk/Volume2 /DATA_backup
```

6b-3) Backing up to tape

This example shows the procedure when data is backed up to a tape medium of tape device /dev/rmt/0 with the tar(1) command.

```
# cd /DATA_backup
# tar cvf /dev/rmt/0 .
```

6b-4) Unmounting the snapshot

Unmount the file system mounted in step 6b-2).

```
# cd /
# umount /DATA_backup
# rmdir /DATA_backup
```

6c) When backing up the ufs file system with the ufsdump(1M) command

6c-1) Before mounting

Check and repair consistency of the ufs file system on proxy volume Volume2 with the fsck(1M) command. When the file system on the master volume was unmounted in step 3b-2), skip this step.

```
# fsck -F ufs /dev/sfdsk/Class1/rdsk/Volume2
```

6c-2) Backing up to tape

This example shows the procedure when data is backed up to a tape medium of tape device /dev/rmt/0 with the ufsdump(1M) command.

```
# ufsdump 0ucf /dev/rmt/0 /dev/sfdsk/Class1/rdsk/Volume2
```

6d) When backing up the sfxfs file system (GFS Local File System or SynfinityFile file system) with the tar(1) command

6d-1) Reconfiguring partition information

Reconfigure partition information of the sfxfs file system (GFS Local File System) on proxy volume Volume2.

```
# sfxadm /dev/sfdsk/Class1/rdsk/Volume2
```

Information

GFS Local File System in Multi-partition Configuration

For the sfxadm(1M) command, specify device special file names of all the volumes that constitute the file system.

6d-2) Before mounting

Check and repair consistency of the sfxfs file system (GFS Local File System) on proxy volume Volume2 with the fsck(1M) command. When the file system on the master volume was unmounted in step 3b-2), skip this step.

```
# fsck -F sfxfs /dev/sfdsk/Class1/rdsk/Volume2
```

Information

GFS Local File System in Multi-partition Configuration

For the fsck(1M) command, specify a device special file name of the representative partition.

6d-3) Mounting the snapshot

Mount the sfxfs file system (GFS Local File System) on proxy volume Volume2 on temporary mount point /DATA_backup.

```
# mkdir /DATA_backup
# mount -F sfxfs /dev/sfdsk/Class1/dsk/Volume2 /DATA_backup
```


Information

GFS Local File System in Multi-partition Configuration

For the mount(1M) command, specify a device special file name of the representative partition.

6d-4) Backing up to tape

This example shows the procedure when data is backed up to a tape medium of tape device /dev/rmt/0 with the tar(1) command.

```
# cd /DATA_backup
# tar cvf /dev/rmt/0 .
```

6d-5) Unmounting the snapshot

Unmount the file system mounted in step 6b-3).

```
# cd /
# umount /DATA_backup
# rmdir /DATA_backup
```

6e) When backing up the sfxfs file system (GFS Local File System or SynfinityFile file system) with the sfxdump(1) command

6e-1) Reconfiguring partition information

Reconfigure partition information of the sfxfs file system (GFS Local File System) on proxy volume Volume2.

```
# sfxadm /dev/sfdsk/Class1/rdsk/Volume2
```

Information

GFS Local File System in Multi-partition Configuration

For the sfxadm(1M) command, specify device special file names of all the volumes that constitute the file system.

6e-2) Before mounting

Check and repair consistency of the sfxfs file system (GFS Local File System) on proxy volume Volume2 with the fsck(1M) command. When the file system on the master volume was unmounted in step 3b-2), skip this step.

```
# fsck -F sfxfs /dev/sfdsk/Class1/rdsk/Volume2
```

Information

GFS Local File System in Multi-partition Configuration

For the fsck(1M) command, specify a device special file name of the representative partition.

6e-3) Backing up to tape

This example shows the procedure when data is backed up to a tape medium of tape device /dev/rmt/0 with the sfxdump(1M) command.

```
# sfxdump 0ucf /dev/rmt/0 /dev/sfdsk/Class1/rdsk/Volume2
```

Information

GFS Local File System in Multi-partition Configuration

For the `sfxdump(1M)` command, specify a device special file name of the representative partition.

7) Rejoining the proxy volume

To perform online backup again, follow the procedures below on a node that belongs to the scope of class `Class1` and then go back to step 2).

7-1) Stop proxy volume `Volume2`. If `Class1` is a shared class, specify the `"-e allnodes"` option to stop `Volume2` on all nodes.

```
# sdxvolume -F -c Class1 -v Volume2
```

Information

When using GDS Management View, select a proxy volume and execute `[Operation]:[Stop Volume]` in the Main Screen.

7-2) Rejoin proxy volume `Volume2` to master volume `Volume1`.

```
# sdxproxy Rejoin -c Class1 -p Volume2
```

After returning from the command, synchronization copying from `Volume1` to `Volume2` is performed.

See

When using GDS Management View, see ["Rejoin"](#) of ["5.3.2.2 Backup \(by Synchronization\)."](#)

Information

GFS Local File System in Multi-partition Configuration

Rejoin proxy volumes for all the volumes that constitute the file system.

8) Breaking the proxy volume

When no more online backup is to be executed, cancel the relationship between master volume `Volume1` and proxy volume `Volume2`. Execute the following command on a node that belongs to the scope of `Class1`.

```
# sdxproxy Break -c Class1 -p Volume2
```

See

When using GDS Management View, see ["5.5.6 Breaking a Proxy."](#)

Information

GFS Local File System in Multi-partition Configuration

Break proxy volumes for all the volumes that constitute the file system.

6.4.2 Online Backup (Snapshot by OPC)

If GDS Snapshot is installed in a system where a disk array unit with the OPC function is used, backup can be performed through instant snapshot by OPC using proxy volumes without stopping services.

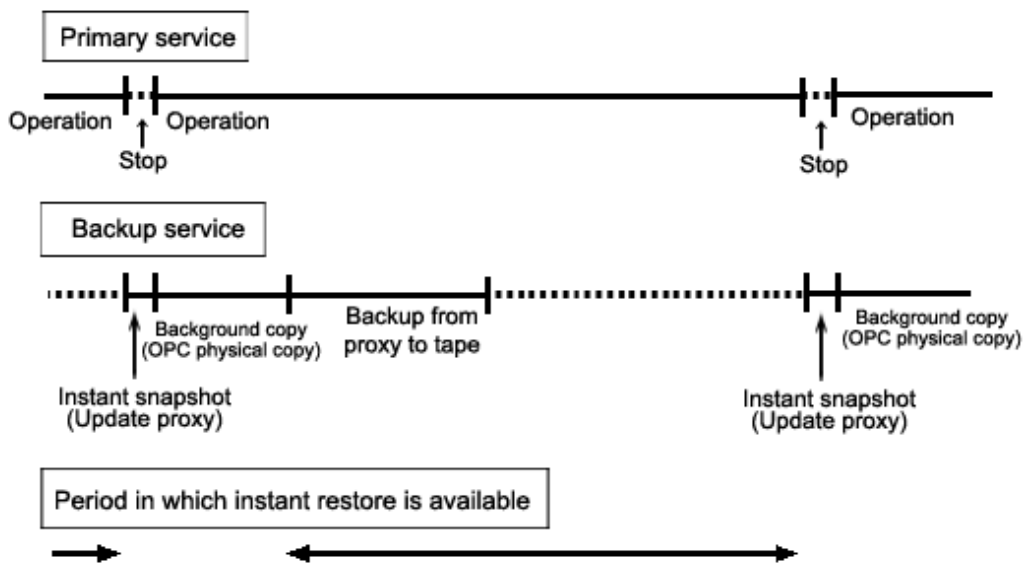
However, to secure consistency of snapshot data, the services must be stopped temporarily when creating a snapshot (when updating proxy volumes).



See

For securing consistency of snapshot data, see "A.2.22 Ensuring Consistency of Snapshot Data."

Figure 6.12 Backup Schedule



Information

Background Copy (OPC Physical Copy) and Backup to Tape

During the copying process, data can be backed up to tape, but the backup will impose the load on the disk array unit and may affect services using master volumes.



Note

Instant Snapshot by OPC

For use conditions of instant snapshot by OPC, see "A.2.19 Instant Snapshot by OPC." If proxy volumes are in mirroring configuration with the multiplicity of two or higher, the OPC function is unavailable for copying data from master volumes to the proxy volumes. For details, see "A.2.18 Using the Advanced Copy Function in a Proxy Configuration."

Procedure

- 1) Relating a proxy volume

Before creating snapshots, relate proxy volume Volume2 as a copy destination to master volume Volume1. In this example, Volume1 and Volume2 belong to class Class1. Execute the following command on a node that belongs to the scope of Class1.

```
# sdxproxy Relate -c Class1 -m Volume1 -p Volume2
```

 See

When using GDS Management View, see "5.2.5.2 Relate."

 Information

GFS Local File System in Multi-partition Configuration

If multiple volumes constitute one GFS Local File System, relate the proxy volume to all of the volumes.

2) Stopping the proxy volume

Stop proxy volume Volume2. If Class1 is a shared class, specify the "-e allnodes" option to stop Volume2 on all nodes.

```
# sdxvolume -F -c Class1 -v Volume2
```

 Information

When using GDS Management View, select a proxy volume and execute [Operation]:[Stop Volume] in the Main Screen.

3) Stopping the services

To secure consistency of snapshot data, stop the services before creating a snapshot and restrain the master volume from being written in.

3a) When the master volume is being used for a cluster application

Inactivate the cluster application.

3b) When the master volume is not being used for a cluster application

3b-1) Stop the services for which the master volume is being used.

3b-2) When the master volume is used as a file system, use either the following methods to prevent write access to the file system. In the following examples, the mount point is /DATA.

- To unmount the file system

```
# cd /  
# umount /DATA
```

- To write-lock the file system

```
# lockfs -w /DATA
```

4) Updating the proxy volume

Copy data from master volume Volume1 to proxy volume Volume2 with the OPC function to update Volume2 with the data of Volume1 at the moment. Execute the following command on a node that belongs to the scope of class Class1.

```
# sdxproxy Update -c Class1 -p Volume2 -e instant
```

When returning from the command, the update of Volume2 will be complete. Subsequently, background OPC physical copying is performed, but you may go on to step 5) without waiting until the copying is complete.



See

When using GDS Management View, see "Update" in "5.3.2.3 Backup (by OPC)." For setting of the update environment, select "Yes" to "Instant Snapshot."



Information

GFS Local File System in Multi-partition Configuration

Update proxy volumes for all the volumes that constitute the file system.

5) Resuming the services

5a) When the master volume is used for a cluster application

Activate the cluster application.

5b) When the master volume is not used for a cluster application

5b-1) When the master volume is used as a file system, use either the following methods to permit write access from applications to the file system on master volume Volume1. In the following examples, the mount point is /DATA.

- When the file system was unmounted in step 3b-2), mount the file system.
- For the ufs file system

```
# mount -F ufs /dev/sfdsk/Class1/dsk/Volume1 /DATA
```

- For the GFS Local File System

```
# mount -F sfxfs /dev/sfdsk/Class1/dsk/Volume1 /DATA
```



Information

GFS Local File System in Multi-partition Configuration

For the mount(1M) command, specify a device special file name of the representative partition.

- When write access to the file system was write-locked in step 3b-2), unlock the file system.

```
# lockfs -u /DATA
```

5b-2) Start the services for which the master volume is used.

6) Starting the proxy volume

Start proxy volume Volume2 on a node where backup to tape is performed.

```
# sdxvolume -N -c Class1 -v Volume2
```

Information

When using GDS Management View, select a proxy volume and execute [Operation]:[Start Volume] in the Main Screen.

7) Confirming the completion of copying

Confirm that the copying is complete.

```
# sdxinfo -S -c Class1 -o Volume2

OBJ      CLASS  GROUP  DISK   VOLUME  STATUS
-----  -----  -----  -----  -----  -----
slice    Class1 Group2  Disk3  Volume2 ACTIVE
slice    Class1 Group2  Disk4  Volume2 ACTIVE
```

If all the displayed slices' STATUS fields are "ACTIVE", copying is complete. If the copying is still in progress, "COPY" will be displayed in the STATUS field.

8) Backing up to tape

Back up the snapshot data on the proxy volume to tape. Execute the following commands on a node that belongs to the scope of class Class1.

See

For details on the backup methods, see the manuals of file systems to be backed up and used commands.

8a) When backing up data held in a raw device with the dd(1M) command

```
# dd if=/dev/sfdsk/Class1/rdisk/Volume2 of=/dev/rmt/0 bs=32768
```

8b) When backing up the ufs file system with the tar(1) command

8b-1) Before mounting

Check and repair consistency of the ufs file system on the proxy volume with the fsck(1M) command. When the file system on the master volume was unmounted in step 3b-2), skip this step.

```
# fsck -F ufs /dev/sfdsk/Class1/rdisk/Volume2
```

8b-2) Mounting the snapshot

Mount the ufs file system on proxy volume Volume2 on temporary mount point /DATA_backup.

```
# mkdir /DATA_backup
# mount -F ufs /dev/sfdsk/Class1/dsk/Volume2 /DATA_backup
```

8b-3) Backing up to tape

This example shows the procedure when data is backed up to a tape medium of tape device /dev/rmt/0 with the tar(1) command.

```
# cd /DATA_backup
# tar cvf /dev/rmt/0 .
```

8b-4) Unmounting the snapshot

Unmount the file system mounted in step 8b-2).

```
# cd /
# umount /DATA_backup
# rmdir /DATA_backup
```

8c) When backing up the ufs file system with the ufsdump(1M) command

8c-1) Before mounting

Check and repair consistency of the ufs file system on the proxy volume with the fsck(1M) command. When the file system on the master volume was unmounted in step 3b-2), skip this step.

```
# fsck -F ufs /dev/sfdsk/Class1/rdsk/Volume2
```

8c-2) Backing up to tape

This example shows the procedure when data is backed up to a tape medium of tape device /dev/rmt/0 with the ufsdump(1M) command

```
# ufsdump 0ucf /dev/rmt/0 /dev/sfdsk/Class1/rdsk/Volume2
```

8d) When backing up the sfxfs file system (GFS Local File System) with the tar(1) command

8d-1) Reconfiguring partition information

Reconfigure partition information of the sfxfs file system (GFS Local File System) on proxy volume Volume2.

```
# sfxadm /dev/sfdsk/Class1/rdsk/Volume2
```



Information

GFS Local File System in Multi-partition Configuration

For the sfxadm(1M) command, specify device special file names of all the volumes that constitute the file system.

8d-2) Before mounting

Check and repair consistency of the sfxfs file system (GFS Local File System) on proxy volume Volume2 with the fsck(1M) command. When the file system on the master volume was unmounted in step 3b-2), skip this step.

```
# fsck -F sfxfs /dev/sfdsk/Class1/rdsk/Volume2
```



Information

GFS Local File System in Multi-partition Configuration

For the fsck(1M) command, specify a device special file name of the representative partition.

8d-3) Mounting the snapshot

Mount the sfxfs file system (GFS Local File System) on proxy volume Volume2 on temporary mount point /DATA_backup.

```
# mkdir /DATA_backup
# mount -F sfxfs /dev/sfdsk/Class1/dsk/Volume2 /DATA_backup
```

Information

GFS Local File System in Multi-partition Configuration

For the mount(1M) command, specify a device special file name of the representative partition.

8d-4) Backing up to tape

This example shows the procedure when data is backed up to a tape medium of tape device /dev/rmt/0 with the tar(1) command.

```
# cd /DATA_backup
# tar cvf /dev/rmt/0 .
```

8d-5) Unmounting the snapshot

Unmount the file system mounted in step 8b-3).

```
# cd /
# umount /DATA_backup
# rmdir /DATA_backup
```

8e) When backing up the sfxfs file system (GFS Local File System) with the sfxdump(1M) command

8e-1) Reconfiguring partition information

Reconfigure partition information of the sfxfs file system (GFS Local File System) on proxy volume Volume2.

```
# sfxadm /dev/sfdsk/Class1/rdsk/Volume2
```

Information

GFS Local File System in Multi-partition Configuration

For the sfxadm(1M) command, specify device special file names of all the volumes that constitute the file system.

8e-2) Before mounting

Check and repair consistency of the sfxfs file system (GFS Local File System) on proxy volume Volume2 with the fsck(1M) command. When the file system on the master volume was unmounted in step 3b-2), skip this step.

```
# fsck -F sfxfs /dev/sfdsk/Class1/rdsk/Volume2
```

Information

GFS Local File System in Multi-partition Configuration

For the fsck(1M) command, specify a device special file name of the representative partition.

8e-3) Backing up to tape

This example shows the procedure when data is backed up to a tape medium of tape device /dev/rmt/0 with the sfxdump(1M) command.

```
# sfxdump 0ucf /dev/rmt/0 /dev/sfdsk/Class1/rdsk/Volume2
```


Information

GFS Local File System in Multi-partition Configuration

For the sfxdump(1M) command, specify a device special file name of the representative partition.

9) Stopping the proxy volume

After backing up to tape, stop Volume2 to protect data in proxy volume Volume2. If Class1 is a shared class, specify the "-e allnodes" option to stop Volume2 on all the nodes.

```
# sdxvolume -F -c Class1 -v Volume2
```

Information

When using GDS Management View, select a proxy volume and execute [Operation]:[Stop Volume] in the Main Screen.

10) Breaking the proxy volume

When no more online backup is to be executed, cancel the relationship between master volume Volume1 and proxy volume Volume2.

```
# sdxproxy Break -c Class1 -p Volume2
```

See

When using GDS Management View, see "5.5.6 Breaking a Proxy."

Information

GFS Local File System in Multi-partition Configuration

Break proxy volumes for all the volumes that constitute the file system

11) Re-executing online backup

To perform online backup again, go back to step 3).

6.4.3 Instant Restore

By keeping proxy volumes parted from master volumes used for services, in the event of a data error in a master volume, data can be restored back from the proxy volume to the master volume (unless background copying is in process after instant snapshot creation).


The master volume must be stopped temporarily for restoration, but you may start the master volume and make it accessible immediately after the restore is initiated without waiting until the copying is complete.

Information

Restore from Tape

If an error occurs in master volume data while the master volume used for the services operation and a proxy volume are joined, the same error occurs in data of the proxy volume and instant restore becomes unavailable. If this happens, restore data back to the master volume following the procedure described in "6.3.3 Restoring."

Figure 6.13 Schedule for Instant Restore



 Note

Instant Restore with Disk Unit's Copy Functions

For restore, only (R)OPC is available, but (R)EC, EMC TimeFinder or EMC SRDF is unavailable. For details, see "A.2.18 Using the Advanced Copy Function in a Proxy Configuration."

Procedure

- 1) Viewing the state of the proxy volume

Confirm that proxy volume Volume2 has been parted from master volume Volume1.

```
# sdxinfo -v -c Class1 -o Volume2 -e long
OBJ      NAME      TYPE      CLASS     GROUP     DISK     MASTER   PROXY ...
-----  -
volume   *         mirror    Class1    Group2    *        *        *     ...
volume   Volume2   mirror    Class1    Group2    *        Volume1 Part   ...
volume   *         mirror    Class1    Group2    *        *        *     ...
```

If Part is displayed in the PROXY field, the proxy volume has been parted.

If Join is displayed in the PROXY field, the proxy volume has been joined to the master volume and instant restore is unavailable. In this situation, data must be restored from tape. For more details, see "Restore from Tape" described above.

- 2) Stopping the services

- 2a) When the master volume is being used for a cluster application

Inactivate the cluster application.

- 2b) When the master volume is not being used for a cluster application

- 2b-1) Stop the services for which the master volume is being used.

- 2b-2) When the master volume is being used as a file system, unmount the file system. This example shows the procedure when the mount point is /DATA.

```
# cd /
# umount /DATA
```

2b-3) Stopping the master volume

Stop master volume Volume1. If Class1 is a shared class, specify the "-e allnodes" option to stop Volume1 on all nodes.

```
# sdxvolume -F -c Class1 -v Volume1
```

Information

When using GDS Management View, select a master volume and execute [Operation]:[Stop Volume] in the Main Screen.

3) Restoring data from the proxy volume

Execute the following commands on a node that belongs to the scope of class Class1.

3a) When the OPC function is unavailable

3a-1) Stop proxy volume Volume2. If Class1 is a shared class, specify the "-e allnodes" option to stop Volume2 on all nodes.

```
# sdxvolume -F -c Class1 -v Volume2
```

Information

When using GDS Management View, select a proxy volume and execute [Operation]:[Stop Volume] in the Main Screen.

3a-2) Restore data from proxy volume Volume2 back to master volume Volume1.

```
# sdxproxy RejoinRestore -c Class1 -p Volume2 -e instant
```

When returning from the command, the restore of Volume1 will be complete. Subsequently, synchronization copying from Volume2 to Volume1 is performed, and you may go on to step 4) without waiting until the copying is complete.

See

When using GDS Management View, see "Restore" in "5.3.3 Restore." For settings of the restore environment, select "Yes" to "Rejoin" and "Yes" to "Instant Restore."

3b) When the OPC function is available

Restore data from proxy volume Volume2 back to master volume Volume1.

```
# sdxproxy Restore -c Class1 -p Volume2 -e instant
```

When returning from the command, the restore of Volume1 will be complete. Subsequently, background OPC physical copying from Volume2 to Volume1 is performed, and you may go on to step 4) without waiting until the copying is complete.

See

When using GDS Management View, see "Restore" in "5.3.3 Restore." For settings of the restore environment, select "No" to "Rejoin" and "yes" to "Instant Restore."

Note

Master Volumes with the Mirroring Multiplicity of Two and Higher

By executing the `sdxproxy Restore` command in step 3b), OPC starts, copying from one of proxy volume's slices to one of master volume's slices. Among slices of the master volume, slices other than the OPC copy destination are excluded from mirroring and thus the data statuses become invalid (INVALID). To recover the master volume mirroring status, perform master volume resynchronization copying by using the `sdxcopy -B` command. If not executing the `sdxcopy -B` command, master volume resynchronization copying automatically starts when starting the master volume in step 4) and data will be copied from the OPC copy destination slice to the other slices with the soft copy function.

Information

GFS Local File System in Multi-partition Configuration

If multiple volumes constitute one GFS Local File System, restore data from the proxy volume to all of the volumes.

4) Resuming the services

Without waiting until the copying is complete, you may resume the services.

Note

Reusing Proxy Volume Data

By executing the `sdxproxy RejoinRestore` command in step 3a), Volume1 and Volume2 are joined and Volume2 will also be updated with data written into Volume1. To reuse data in Volume2 for restore without updating, after the synchronization copying from Volume2 to Volume1 is complete, part Volume2 from Volume1 and then resume the services. When the `sdxproxy Restore` command was executed in step 3b), Volume1 and Volume2 are left parted, and data in Volume2 remains unchanged even if the services are resumed before the copying is complete.

4a) When the master volume is used for a cluster application

Activate the cluster application.

4b) When the master volume is not used for a cluster application

4b-1) Activate master volume Volume1 on the node running the services.

```
# sdxvolume -N -c Class1 -v Volume1
```

Information

When using GDS Management View, select the master volume and execute [Operation]:[Start Volume] in the Main Screen.

Information

When the OPC Function Is Available

If the master volume mirroring multiplicity is two and higher and if restore is performed with the OPC function in step 3b), master volume resynchronization copying automatically starts after command execution. To perform resynchronization copying after OPC physical copying is complete, specify the `-e nosync` option for the `sdxvolume -N` command, and the master volume will start without invoking resynchronization copying. With this method, perform master volume resynchronization copying with the `sdxcopy -B` command after OPC physical copying is complete.

4b-2) When the master volume is used as the sfxfs file system (GFS Local File System), reconfigure the partition information.

```
# sfxadm /dev/sfdsk/Class1/rdsk/Volume1
```

Information

GFS Local File System in Multi-partition Configuration

For the sfxadm(1M) command, specify device file names of all the volumes that constitute the file system.

4b-3) When the master volume is used as a file system, mount the file system. In this example, the mount point is /DATA.

- For the ufs file system

```
# mount -F ufs /dev/sfdsk/Class1/dsk/Volume1 /DATA
```

- For the sfxfs file system (GFS Local File System)

```
# mount -F sfxfs /dev/sfdsk/Class1/dsk/Volume1 /DATA
```

Information

GFS Local File System in Multi-partition Configuration

For the mount(1M) command, specify a device special file name of the representative partition.

4b-4) Resume the servicess using the master volume.

5) Viewing the copy status

The status of the copying from proxy volume Volume2 to master volume Volume1 started in step 3) can be viewed by using the sdxinfo -S command. The copy destination slice is in the COPY status if copying is in progress and it will be in the ACTIVE status after the copy process ends normally.

```
# sdxinfo -S -c Class1 -o Volume1
OBJ   CLASS  GROUP  DISK   VOLUME  STATUS
-----
slice Class1 Group1 Disk1  Volume1 ACTIVE
slice Class1 Group1 Disk2  Volume1 ACTIVE
```

Information

By executing the sdxproxy RejoinRestore command in step 3a), master volume Volume1 and proxy volume Volume2 are joined. In the even of a data error occurs in Volume1, the same data error occurs in Volume2, and restoring data back from tape will be required. Therefore, after resynchronization copying from Volume2 to Volume1 is complete, it is recommended to part Volume2 from Volume1. For the procedures for parting proxy volumes, see the steps 3) through 5) in "[6.4.1 Online Backup \(by Synchronization\)](#)."

Note

When Using (R)EC, EMC TimeFinder, EMC SRDF, for Backup

Restoration of the master by use of the proxy stops sessions of these copying functions. To perform backup using these copying functions, cancel the relation between the master and the proxy with the sdxproxy Break command.

6.5 Backing Up and Restoring through Disk Unit's Copy Functions

Some sophisticated disk devices contain hardware functions to copy disk data within the disk units or to other disk units. For example, Fujitsu Storage Systems ETERNUS provide the Advanced Copy function and EMC's Symmetrix storage systems provide copy functions such as TimeFinder and SRDF.

This section describes the procedures for backing up and restoring object configurations and data of local disks and shared disks through use of these disk unit's copy functions.

Backup and restore can be performed on the following nodes.

- Nodes that operate services
- Nodes that belong to the same cluster domains as those of nodes operating services
- Nodes that do not belong to the same cluster domains as those of nodes operating services

In the following subsections, physical disks named c1t1d1 and c2t1d1 are registered with a shared class named Class1 and mirrored, and a mirror volume named Volume1 is used for services.

6.5.1 Configuring an Environment



Automatic Resource Registration

If the backup server resides in a cluster domain (called a backup domain) that is different from the primary domain, those disks which are registered as resources in the primary domain or are to be registered with classes restored in the backup domain may not be involved in the resource registration in the backup domain. In the backup domain, those relevant disks must be described in the Excluded Device List prior to executing the automatic resource registration. For details on the automatic resource registration, see "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide."

1) Creating an application volume

Create an application mirror volume onto application disks c1t1d1 and c2t1d1. The following settings are necessary on Node1 or Node2 in the primary domain.

1-1) Registering disks

Register disks c1t1d1 and c2t1d1 with shared class Class1 that is shared on Node1 and Node2, and name them Disk1 and Disk2 respectively.

```
# sdxdisk -M -c Class1 -a type=shared,scope=Node1:Node2 -d  
c1t1d1=Disk1,c2t1d1=Disk2
```

1-2) Creating a mirror group

Connect disks Disk1 and Disk2 to mirror group Group1.

```
# sdxdisk -C -c Class1 -g Group1 -d Disk1,Disk2
```

1-3) Creating a mirror volume

Create mirror volume Volume1 into mirror group Group1.

```
# sdxvolume -M -c Class1 -g Group1 -v Volume1 -s 1048576
```

2) Synchronizing backup disks

Synchronize application disks c1t1d1 and c2t1d1, as copy sources, respectively with disks c3t1d1 and c4t1d1, as copy destinations, with copy functions of the disk units.



For the synchronization methods, see the manuals of disk units for information regarding copy functions.



Backup disks must be equivalent in size application disks to be backed up.

6.5.2 Backing Up

3) Backing up the object configuration of the class

On Node1 or Node2 in the primary domain, back up the object configuration of Class1 to be backed up.

3-1) Saving configuration information

Save outputs of the `sdxinfo` command to a file. In this example, the path to a file is `"/var/tmp/Class1.info"`.

```
# sdxinfo -c Class1 -e long > /var/tmp/Class1.info
```

3-2) Creating a configuration file

Output the object configuration within Class1 to a file in configuration table format. In this example, the path to a file is `"/var/tmp/Class1.conf"`.

```
# sdxconfig Backup -c Class1 -o /var/tmp/Class1.conf
```

3-3) Backing up the configuration information and configuration file

Save the files created in steps 3-1) and 3-2) to tape and so on.

4) Detaching the backup disks (suspending synchronization)



In this example, stop services when detaching the backup disks in order to secure consistency of data. If installed software, such as file systems and database systems that manage volume data, provides functions for securing or repairing consistency of data on detached copy destination disks is present, skip steps 4-3) and 4-5). Alternatively, perform operations for securing consistency with software-specific methods. For details, see "[A.2.22 Ensuring Consistency of Snapshot Data](#)."

4-1) Viewing the application volume status

Confirm that data of slices comprising application volume Volume1 is in valid status (ACTIVE or STOP).

```
# sdxinfo -S -c Class1 -o Volume1
OBJ   CLASS  GROUP  DISK   VOLUME  STATUS
-----
slice Class1 Group1 Disk1   Volume1 ACTIVE
slice Class1 Group1 Disk2   Volume1 ACTIVE
```

If it is not in valid status (ACTIVE or STOP), recover the slice status, referring to "[F.1.1 Slice Status Abnormality](#)."

4-2) Checking the statuses of the disk unit's copy functions

Confirm that the application disks and the backup disks are synchronized.



.....
For the confirming methods, see the manuals of disk units for information regarding copy functions.
.....

4-3) Stopping services

To secure consistency of data in the backup disks after they are detached from the application disks, stop applications using application volume Volume1 on Node1 and Node2.

When using Volume1 as a file system, unmount the file system.

4-4) Detaching the backup disks (suspending synchronization)

Detach backup disks c3t1d1 and c4t1d1 from application disks c1t1d1 and c2t1d1.



.....
For the synchronization methods, see the manuals of disk units for information regarding copy functions.
.....

4-5) Resuming the services

When the file system was unmounted in step 4-3), mount it again.

Resume the applications stopped in step 4-3).

5) Creating a backup volume

On backup server Node3, create a backup volume into backup disks c3t1d1 and c4t1d1. The following settings are necessary on backup server Node3.

5-1) Placing the configuration file

Place the configuration file "/var/tmp/Class1.conf" backed up in step 3) onto backup server Node3. In this example, the path to a destination file is "/var/tmp/Class1.conf."

5-2) Changing physical disks in the configuration file

Change the physical disk names of the application disks described in the configuration file "/var/tmp/Class1.conf" from c1t1d1 and c2t1d1 to c3t1d1 and c4t1d1, which are the physical disk names of the backup disks, respectively.

```
# sdxconfig Convert -e replace -c Class1 -p c1t1d1=c3t1d1,c2t1d1=c4t1d1 \  
-i /var/tmp/Class1.conf -o /var/tmp/Class1.conf -e update
```



Physical Disk Sizes

.....

The former physical disks and the latter physical disks must be equivalent in size.
.....

5-3) Editing a class name in the configuration file

Change the class name of the configuration table described in the configuration file "/var/tmp/Class1.conf" from Class1 to Class2 and save the changes to the configuration file "/var/tmp/Class2.conf." If Class1 already exists in a domain to which the backup server belongs, the class must be renamed.

```
# sdxconfig Convert -e rename -c Class1=Class2 -i /var/tmp/Class1.conf -  
o /var/tmp/Class2.conf
```


5-4) Creating a backup volume

According to the configuration table in the configuration file "/var/tmp/Class2.conf" created in step 5-3), create the object configuration of class Class2.

```
# sdxconfig Restore -c Class2 -i /var/tmp/Class2.conf -e chkps,skipsync
```

On backup server Node3, backup disks c3t1d1 and c4t1d1 are registered with local class Class2. Those disks are assigned Disk1 and Disk2 respectively and backup volume Volume1 is created on disks Disk1 and Disk2.

The backup disks to which write access was prevented were detached in step 4-4) and consistency between c3t1d1 and c4t1d1 has been ensured.

Therefore, synchronization copying can be skipped when creating mirror volume Volume1 by specifying the -e skipsync option for the sdxconfig Restore command.

6) Backing up to tape

On backup server Node3, back up data in backup volume Volume1 to a tape medium of tape unit /dev/rmt/0.



See

For details on the backup methods, see the manuals of backed up file systems and used commands.



Information

GFS Shared File System

Back up data with the method shown in step 6a).

6a) When backing up data held in a raw device with the dd(1M) command

```
# dd if=/dev/sfdsk/Class2/rdsk/Volume1 of=/dev/rmt/0 bs=32768
```

6b) When backing up the ufs file system with the tar(1) command

6b-1) Check and repair consistency of the ufs file system on backup volume Volume1. If the file system was unmounted when the backup disks were detached in step 3), skip this step.

```
# fsck -F ufs -y /dev/sfdsk/Class2/rdsk/Volume1
```

6b-2) Mount the ufs file system on backup volume Volume1 on /mnt1, a temporary mount point, in the read only mode.

```
# mkdir /mnt1
# mount -F ufs -o ro /dev/sfdsk/Class2/dsk/Volume1 /mnt1
```

6b-3) Back up data held in the file system to tape.

```
# cd /mnt1
# tar cvf /dev/rmt/0 .
```

6b-4) Unmount the file system mounted in step 6b-2).

```
# cd /
# umount /mnt1
# rmdir /mnt1
```

6c) When backing up the ufs file system with the ufsdump(1M) command

6c-1) Check and repair consistency of the ufs file system on backup volume Volume1. If the file system was unmounted when the backup disks were detached in step 3), skip this step.

```
# fsck -F ufs -y /dev/sfdsk/Class2/rdsk/Volume1
```

6c-2) Back up data held in the file system to tape.

```
# ufsdump 0ucf /dev/rmt/0 /dev/sfdsk/Class2/rdsk/Volume1
```

6d) When backing up the sfxfs file system (GFS Local File System) with the tar(1) command

6d-1) Reconfigure the partition information of the sfxfs file system on backup volume Volume1.

```
# sfxadm /dev/sfdsk/Class2/rdsk/Volume1
```

6d-2) Check and repair consistency of the sfxfs file system on backup volume Volume1. If the file system was unmounted when the backup disks were detached in step 3), skip this step.

```
# fsck -F sfxfs -y /dev/sfdsk/Class2/rdsk/Volume1
```

6d-3) Mount the sfxfs file system on backup volume Volume1 on /mnt1, a temporary mount point, in the read only mode.

```
# mkdir /mnt1
# mount -F sfxfs -o ro /dev/sfdsk/Class2/dsk/Volume1 /mnt1
```

6d-4) Back up data held in the file system to tape.

```
# cd /mnt1
# tar cvf /dev/rmt/0 .
```

6d-5) Unmount the file system mounted in step 6d-3).

```
# cd /
# umount /mnt1
# rmdir /mnt1
```

6e) When backing up the sfxfs file system (GFS Local File System) with the sfxdump(1M) command

6e-1) Reconfigure the partition information of the sfxfs file system on backup volume Volume1.

```
# sfxadm /dev/sfdsk/Class2/rdsk/Volume1
```

6e-2) Check and repair consistency of the sfxfs file system on backup volume Volume1. If the file system was unmounted when the backup disks were detached in step 3), skip this step.

```
# fsck -F sfxfs -y /dev/sfdsk/Class2/rdsk/Volume1
```

6e-3) Back up data held in the file system to tape.

```
# sfxdump 0ucf /dev/rmt/0 /dev/sfdsk/Class2/rdsk/Volume1
```

7) Removing the backup volumes

After the backup process is complete, delete the object configuration of Class2 created for backup. On backup server Node3, perform the following procedures.

7-1) Stopping the backup volume

Stop all the volumes in Class2.

```
# sdxvolume -F -c Class2
```

7-2) Deleting the object configuration of Class2

Delete the object configuration of Class2.

```
# sdxconfig Remove -c Class2
```

8) Resynchronizing the backup disks

Preparatory to the next backup, resynchronize application disks c1t1d1 and c2t1d1, as copy sources, respectively with backup disks c3t1d1 and c4t1d1, as copy destinations, with copy functions of the disk units.



See

.....
For the resynchronization methods, see the manuals of disk units for information regarding copy functions.
.....

6.5.3 Restoring from Backup Disks

9) Stopping services

Stop applications using application volume Volume1 on nodes Node1 and Node2 in the primary domain.

When using Volume1 as a file system, unmount the file system.

10) Stopping the application volume

On Node1 and Node2 in the primary domain, stop application volume Volume1. Execute the following command on Node1 or Node2

```
# sdxvolume -F -c Class1 -v Volume1 -e allnodes
```

11) Deleting the class

In the primary domain, delete class Class1 to which application volume Volume1 belongs. Execute the following command on Node1 or Node2 in the primary domain.

```
# sdxvolume -R -c Class1 -v Volume1  
# sdxgroup -R -c Class1 -g Group1  
# sdxclass -R -c Class1
```

12) Restoring data back from the backup disks

Restore data from backup disks c3t1d1 and c4t1d1, copy sources, as respectively back to application disks c1t1d1 and c2t1d1, as copy destinations, with disk unit's copy functions.



See

For the restore methods, see the manuals of disk units for information regarding copy functions.

13) Parting the backup disks (suspending synchronization)

After the restore process is completed, part backup disks c3t1d1 and c4t1d1 from application disks c1t1d1 and c2t1d1.



See

For the methods for suspending synchronization, see the manuals of disk units for information regarding copy functions.

14) Restoring the object configuration of the class

On Node1 or Node2 in the primary domain, according to the configuration table in the configuration file "/var/tmp/Class1.conf" created in step 3) of "6.5.2 Backing Up," restore the object configuration of Class1.

After restoring the object configuration, reboot the restored node.

```
# sdxconfig Restore -c Class1 -i /var/tmp/Class1.conf
# shutdown -y -i6 -g0
```

15) Changing the class type and expanding the class scope

If the backed up class, Class1, is a shared class, change the type and scope attributes of Class1. For the scope of backed up class Class1, check the SCOPE field of the class information output by the sdxinfo command and saved in step 3) of "6.5.2 Backing Up." In this example, the scope of backed up class Class1 is Node1:Node2.

15-1) Stop the volume in the class.

```
# sdxvolume -F -c Class1
```

15-2) Change the class type and expand the class scope.

```
# sdxattr -C -c Class1 -a type=shared,scope=Node1:Node2
```

16) Starting the application volume

On Node1 or Node2 in the primary domain, start application volume Volume1. Execute the following command on Node1 or Node2.

16-1) Starting the application volume

Start application volume Volume1.

```
# sdxvolume -N -c Class1 -v Volume1
```

16-2) Reconfiguring the partition information of the GFS Local File System

When using application volume Volume1 as the GFS Local File System, reconfigure the partition information.

```
# sfxadm /dev/sfdsk/Class1/rdsk/Volume1
```

17) Resuming the services

When the file system on application volume Volume1 was unmounted in step 8) of "6.5.2 Backing Up," unmount it again.

Start the applications using Volume1.

6.6 Backing Up and Restoring through an External Server

This section discusses the method of backing up data from and restoring data back to logical volumes (called application volumes in this manual) in a local or shared class in the primary domain through a server in a domain different from the primary domain (called an external backup server in this manual).

The backup and restore operations through an external backup server can be categorized into 4 patterns.

1. Backing up and restoring a logical volume with no replication



See

For details, see "6.6.1 Backing Up and Restoring a Logical Volume with No Replication."

2. Backing up and restoring through snapshot by slice detachment



See

For details, see "6.6.2 Backing Up and Restoring through Snapshot by Slice Detachment."

3. Backing up and restoring using snapshots from a proxy volume



See

For details, see "6.6.3 Backing Up and Restoring Using Snapshots from a Proxy Volume."

4. Backing up and restoring by the disk unit's copy function



See

For details, see "6.6.4 Backing Up and Restoring by the Disk Unit's Copy Function."

The following table summarizes characteristics of the respective operations.

Patterns	Online backup	Backed up target	Disk unit's copy function	Application volume type (*1)	Required component (primary domain) (*2)
1	Not available	Application volumes	-	Any	GDS (*8)
2	Available	Slices detached temporarily	-	One of (*3): mirror (concat + mirror) (stripe + mirror)	GDS (*8)
3	Available	Proxy volumes	-	One of (*4): single mirror concat + mirror stripe + mirror	GDS GDS Snapshot (*9)

Patterns	Online backup	Backed up target	Disk unit's copy function	Application volume type (*1)	Required component (primary domain) (*2)
			Advanced Copy function	One of (*5): single mirror	
			EMC TimeFinder or EMC SRDF (*6)	One of (*6): mirror	
4	Available	Non-SDX disks as the disk unit's copy function destinations	EMC TimeFinder or EMC SRDF (*7)	Any	GDS (*8)

(*1):

The table above describes the volume types according to the following classification.

Type	Description
single	Single volume created in a single disk.
mirror	Mirror volume connected to a mirror group to which one or more disks are connected. This is excluded when a lower level group is connected to the mirror group. Making disk data redundant can improve the services continuity.
concat	Volume created in a concatenation group. When no large volume exists, it can be created by concatenating multiple disks.
stripe	Stripe volume created in a stripe group. I/O loads of the services can be balanced up on multiple disks.
concat + mirror	Mirror volume created in a mirror group to which one or more concatenation groups are connected. This is excluded when a disk is connected to the mirror group. The effects of both concat and mirror can be gained.
stripe + mirror	Mirror volume created in a mirror group to which one or more stripe group are connected. This is excluded when a disk is connected to the mirror group. The effects of both stripe and mirror can be gained.
(concat + mirror)	Mirror volume created in a mirror group to which one disk and one or more concatenation groups are connected. Since the size of this volume is limited by the size of the disk connected to the mirror group, the effects of concat cannot be gained.
(stripe + mirror)	Mirror volume created in a mirror group to which one disk and one or more stripe group are connected. Since I/O loads to the disk connected to the mirror group are not balanced up, the effects of stripe cannot be gained. However, while the slice is temporarily detached from the disk, the effects of stripe can be gained since I/O loads to the volume are balanced up on multiple disks connected to the stripe group. Note, however, that a slice cannot be temporarily detached from a master volume related to a proxy volume.

(*2):

For an external backup server, GDS and GDS Snapshot must be installed for creating shadow volumes.

(*3):

See ["A.2.12 Creating a Snapshot by Slice Detachment."](#)

(*4):

See ["A.1.8 Proxy Configuration Preconditions."](#)

(*5):

See ["A.1.8 Proxy Configuration Preconditions"](#) and ["A.2.18 Using the Advanced Copy Function in a Proxy Configuration."](#)

(*6):

See ["A.1.8 Proxy Configuration Preconditions"](#) and ["A.2.21 Using EMC TimeFinder or EMC SRDF in a Proxy Configuration."](#)

(*7):

There are two operation patterns that use EMC TimeFinder or EMC SRDF. The features of respective types are as follows.

- Backup and restore using snapshots from proxy volumes

Snapshot operations are available through only GDS and GDS Snapshot commands without using a SYMCLI command. When striping or concatenation has been applied to the application volume, copying to proxy volumes is conducted by the soft copy function instead of TimeFinder and SRDF.

- Backup and restore using the hard copy function

Disk areas that are copy destinations of TimeFinder or SRDF can be backed up, regardless of the type of the application volume.

(*8):

SynfinityDisk can substitute.

(*9):

SynfinityDisk and SafeDISK Snapshot can substitute.

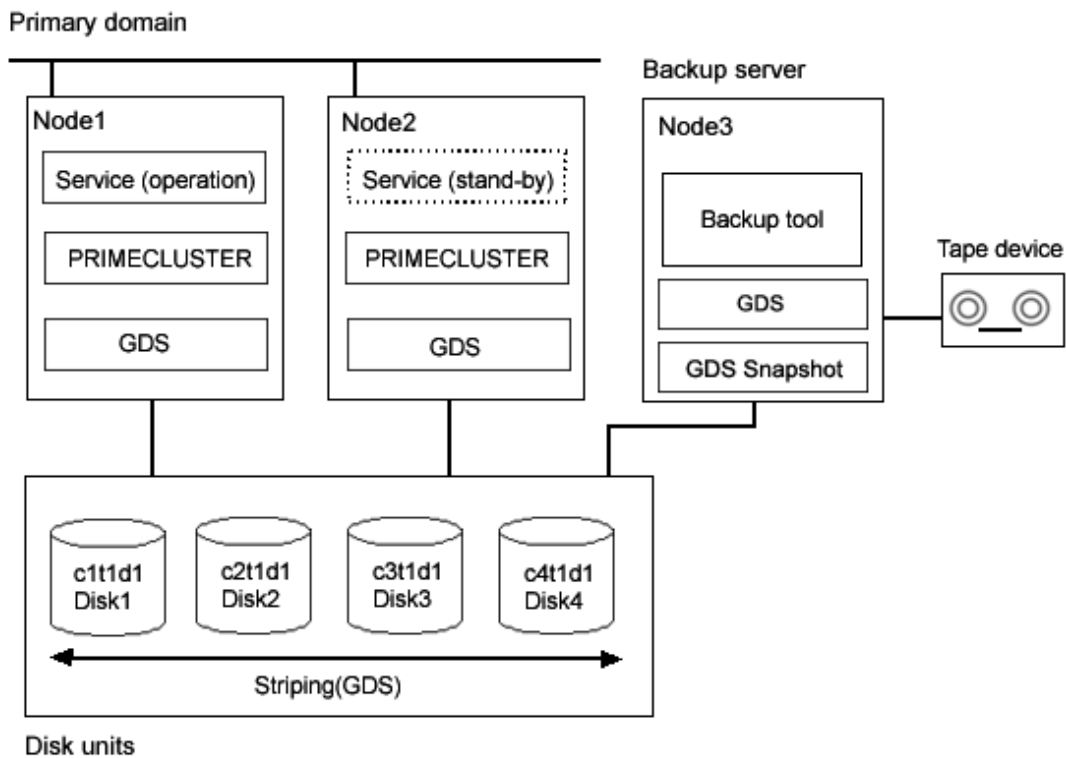
6.6.1 Backing Up and Restoring a Logical Volume with No Replication

This sub-section describes the method of backing up data from and restoring data back to logical volumes in the primary domain through a backup server in another domain.

The following is an example of backing up and restoring a stripe volume. Mirror volumes, single volumes, and volumes in concatenation groups can also be backed up and restored in a similar manner. However, for backing up mirror volumes, the method described in ["6.6.2 Backing Up and Restoring through Snapshot by Slice Detachment"](#) is recommended.

6.6.1.1 System Configuration

Figure 6.14 System Configuration



Information

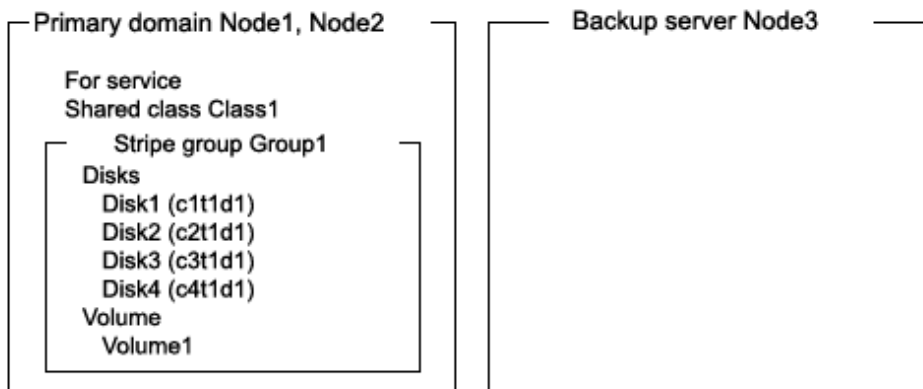
Not only GDS and PRIMECLUSTER but also SynfinityDisk and SynfinityCluster are available in the primary domain.

Note

Physical Device Name

Different physical device names (such as c1t1d1) may be assigned to the identical physical disk in the primary domain and the backup server.

Figure 6.15 Object Configuration in Normal Operation



6.6.1.2 Summary of Backup

Backups can be created while the service is stopped and the application volume is not in use.

Figure 6.16 Backup

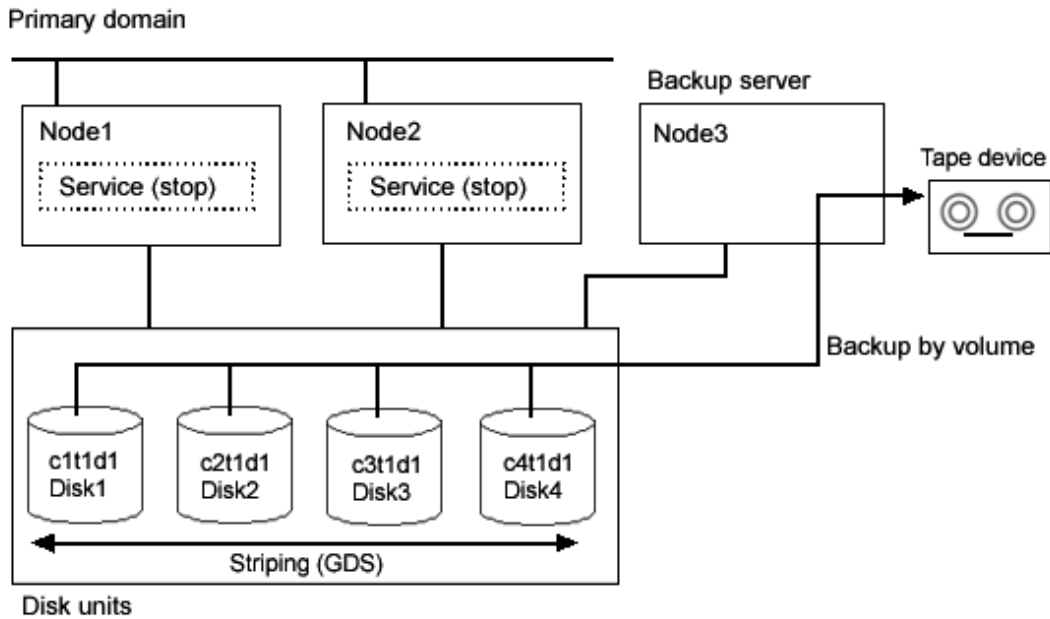


Figure 6.17 Object Configuration When Backing Up

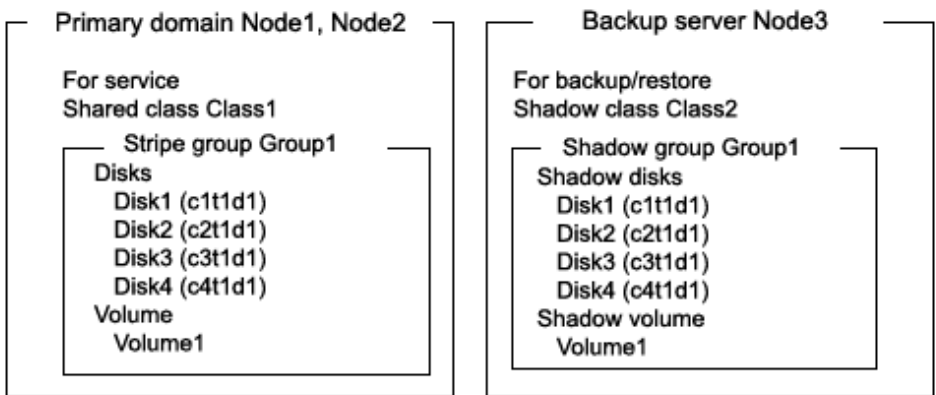
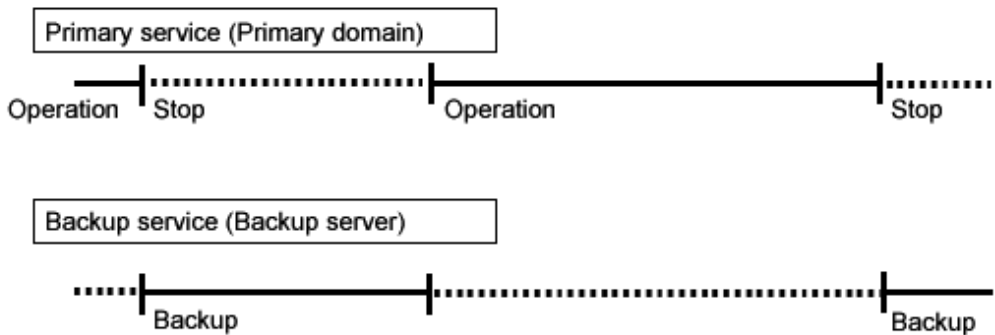


Figure 6.18 Backup Schedule



6.6.1.3 Summary of Restore

If volume data is damaged, it can be restored from tape. Data can be restored while the services is stopped and the application volume is not in use.

Figure 6.19 Restore

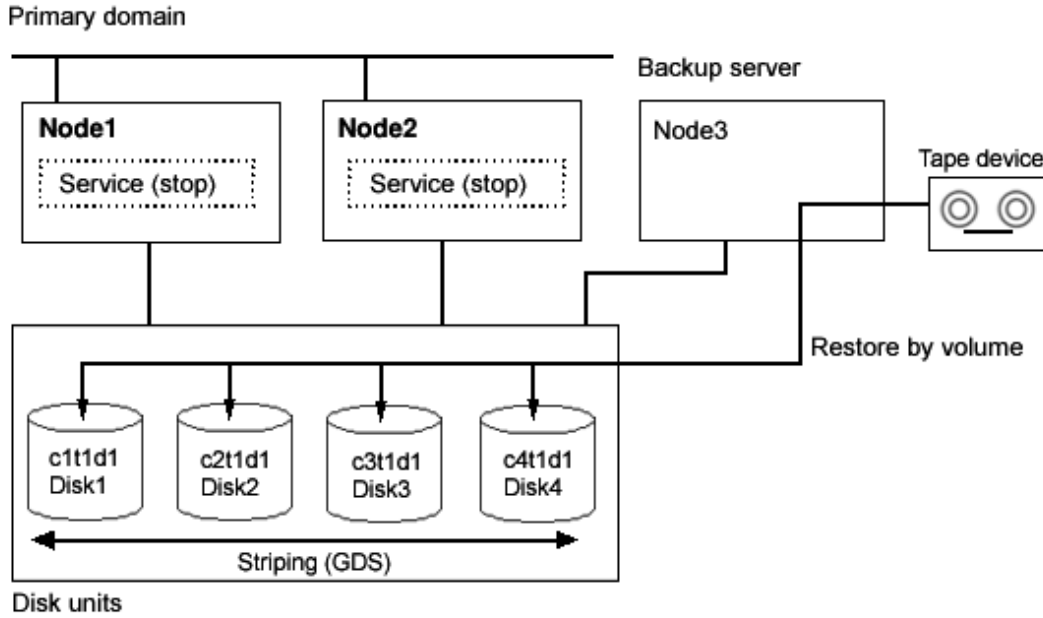


Figure 6.20 Object Configuration When Restoring

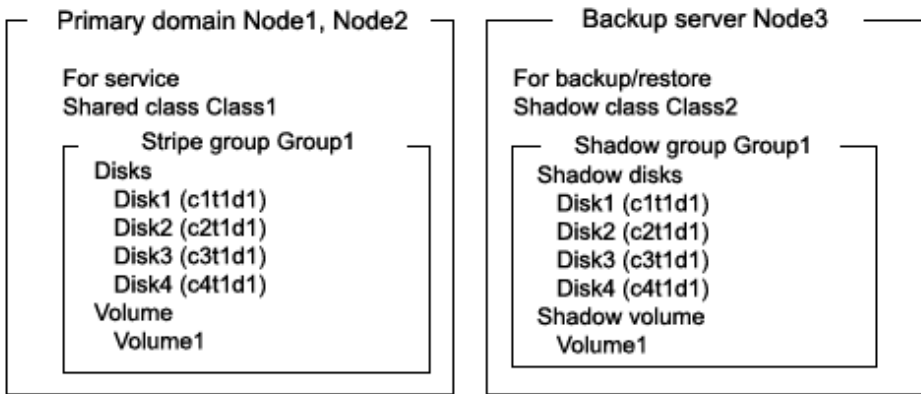
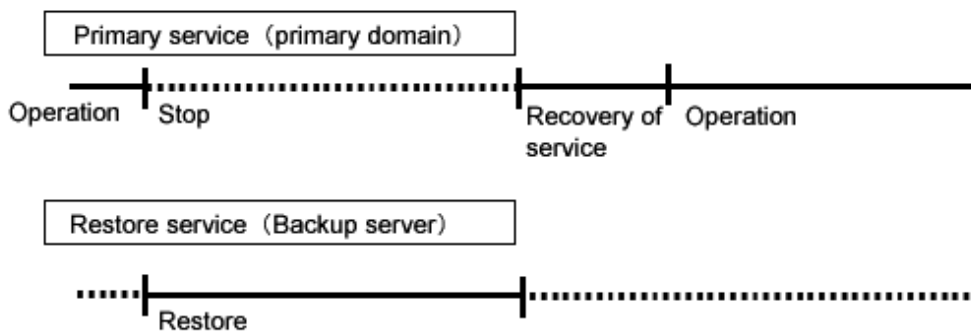


Figure 6.21 Restore Schedule



6.6.1.4 Summary of Procedure

Figure 6.22 Outline of the Configuration Procedure



Figure 6.23 Outline of the Backup Procedure

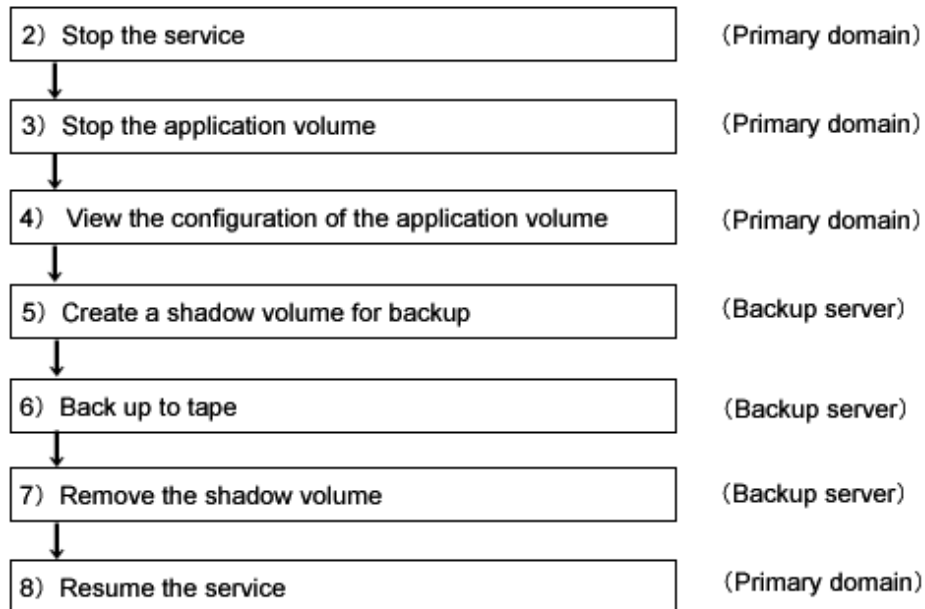
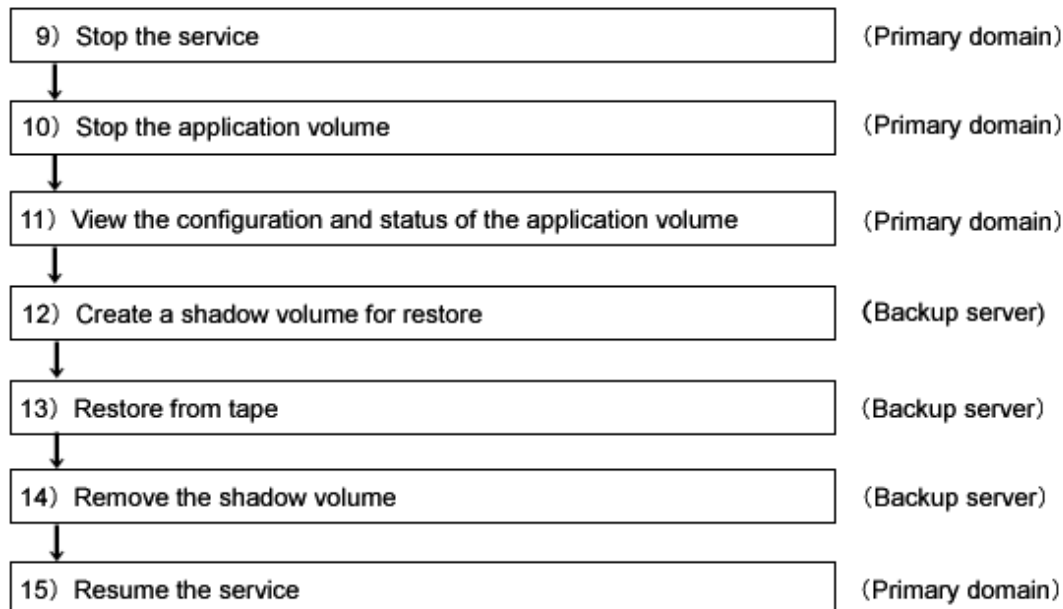


Figure 6.24 Outline of the Restore Procedure



6.6.1.5 Configuring an Environment



Automatic Resource Registration

If the backup server resides in a cluster domain (called a backup domain), those disks that are registered as resources in the primary domain or are to be registered with a shadow class in the backup domain may not be involved in the resource registration in the backup domain. In the backup domain, those relevant disks must be described in the Excluded Device List prior to executing the automatic resource registration. For details on the automatic resource registration, see "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide."

1) Creating an application volume

Create a stripe volume used for the services on disks c1t1d1, c2t1d1, c3t1d1, and c4t1d1. The following settings are necessary on Node1 or Node2 in the primary domain.

1-1) Registering disks

Register disks c1t1d1, c2t1d1, c3t1d1, and c4t1d1 with shared class Class1 that is shared on Node1 and Node2, and name them Disk1, Disk2, Disk3, and Disk4 respectively.

```
# sdxdisk -M -c Class1 -a type=shared,scope=Node1:Node2 -d
c1t1d1=Disk1,c2t1d1=Disk2,c3t1d1=Disk2,c4t1d1=Disk4
```

1-2) Creating a stripe group

Connect disks Disk1, Disk2, Disk3, and Disk4 to stripe group Group1.

```
# sdxdisk -C -c Class1 -g Group1 -d Disk1,Disk2,Disk3,Disk4 -a
type=stripe,width=256
```

1-3) Creating a stripe volume

Create stripe volume Volume1 to stripe group Group1.

```
# sdxvolume -M -c Class1 -g Group1 -v Volume1 -s 1048576 -a pslice=off
```

6.6.1.6 Backing Up

2) Stopping the services

Exit all applications accessing the application volume Volume1 in the primary domain on Node1 and Node2.

When Volume1 is used as a file system, it should be unmounted.

3) Stopping the application volume

To write-lock volume Volume1, inactivate Volume1 on Node1 and Node2 in the primary domain. Execute the following command on Node1 or Node2.

```
# sdxvolume -F -c Class1 -v Volume1 -e allnodes
```

4) Viewing the configuration of the application volume

On Node1 or Node2 in the primary domain, view the configuration of application volume Volume1 that is the backup target.

```
# sdxinfo -c Class1
OBJ      NAME      TYPE      SCOPE      SPARE
-----
class   Class1   shared   Node1:Node2  0
```

OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	DEVCONNECT	STATUS	
disk	Disk1	stripe	Class1	Group1	c1t1d1	8380800	Node1:Node2	ENABLE	
disk	Disk2	stripe	Class1	Group1	c2t1d1	8380800	Node1:Node2	ENABLE	
disk	Disk3	stripe	Class1	Group1	c3t1d1	8380800	Node1:Node2	ENABLE	
disk	Disk4	stripe	Class1	Group1	c4t1d1	8380800	Node1:Node2	ENABLE	

OBJ	NAME	CLASS	DISKS	BLKS	FREEBLKS	SPARE
group	Group1	Class1	Disk1:Disk2:Disk3:Disk4	32964608	31850496	*

OBJ	NAME	CLASS	GROUP	SKIP	JRM	1STBLK	LASTBLK	BLOCKS	STATUS
volume	*	Class1	Group1	*	*	0	65535	65536	PRIVATE
volume	Volume1	Class1	Group1	*	*	65536	1114111	1048576	STOP
volume	*	Class1	Group1	*	*	1114112	32964607	31850496	FREE

OBJ	CLASS	GROUP	DISK	VOLUME	STATUS
slice	Class1	Group1	*	Volume1	STOP

If application volume Volume1 belongs to a stripe group, also pay attention to the stripe width.

```
# sdxinfo -G -c Class1 -o Group1 -e long
```

OBJ	NAME	CLASS	DISKS	BLKS	FREEBLKS	SPARE	MASTER	TYPE	WIDTH
group	Group1	Class1	Disk1:Disk2:Disk3:Disk4	32964608	31850496	*	*	stripe	

256

5) Creating a shadow volume for backup

On backup server Node3, create a backup volume (shadow volume) in the same configuration as the application volume found in step 4). The following settings are necessary on the backup server.

Note

Application volume data may be damaged if data is written into a shadow volume in incorrect configuration. Be sure to confirm that the shadow volume configuration is correct in step 5-4).

5-1) Registering shadow disks

Register disks c1t1d1, c2t1d1, c3t1d1, and c4t1d1 with shadow class Class2, and name them Disk1, Disk2, Disk3, and Disk4 respectively.

```
# sdxshadowdisk -M -c Class2 -d
c1t1d1=Disk1,c2t1d1=Disk2,c3t1d1=Disk3,c4t1d1=Disk4
```

Point

- The disk names must correspond to the disk names assigned in step 1-1). The disk names assigned in 1-1) can be viewed in the NAME field for disk information displayed with the sdxinfo command in step 4).
- The class can be assigned any name.

5-2) Creating a shadow group

Connect shadow disks Disk1, Disk2, Disk3, and Disk4 to stripe type shadow group Group1.

```
# sdxshadowdisk -C -c Class2 -g Group1 -d Disk1,Disk2,Disk3,Disk4 -a
type=stripe,width=256
```

Point

- If the application volume belongs to a stripe group or a concatenation group, the order of connecting shadow disks to a shadow group must correspond to the order of connecting disks to a group in step 1-2). The order of connecting disks in step 1-2) can be viewed in the DISKS field for group information displayed with the sdxinfo command in step 4).
- When the application volume belongs to a stripe group, the stripe width of a shadow group must correspond to the stripe width specified in step 1-2). The stripe width specified in step 1-2) can be viewed in the WIDTH field for group information displayed with the sdxinfo -e long command in step 4).
- The group can be assigned any name.

5-3) Create a shadow volume.

Create shadow volume Volume1 to shadow group Group1.

```
# sdxshadowvolume -M -c Class2 -g Group1 -v Volume1 -s 1048576
```

Point

- The volume must be created in the size corresponding to the volume size in step 1-3). The size of a volume created in step 1-3) can be viewed in the BLOCKS field for volume information displayed with the sdxinfo command in step 4).
- If there are multiple volumes, the corresponding shadow volumes must be created in the order of ascending values (first block numbers) in the 1STBLK field for volume information displayed with the sdxinfo command in step 4).
- The volume can be assigned any name.

5-4) Viewing the shadow volume configuration

Using the sdxinfo command, confirm that the group configuration and the volume configuration are correct based on group information in the DISKS field, volume information in the 1STBLK field and in the BLOCKS field and so on.

```
# sdxinfo -c Class2
OBJ      NAME      TYPE      SCOPE      SPARE
-----
class    Class2    shared    Node3      0

OBJ      NAME      TYPE      CLASS      GROUP      DEVNAM      DEVBLKS      DEVCONNECT      STATUS
-----
disk     Disk1    stripe    Class2     Group1     c1t1d1      8380800      Node3           ENABLE
disk     Disk2    stripe    Class2     Group1     c2t1d1      8380800      Node3           ENABLE
disk     Disk3    stripe    Class2     Group1     c3t1d1      8380800      Node3           ENABLE
disk     Disk4    stripe    Class2     Group1     c4t1d1      8380800      Node3           ENABLE

OBJ      NAME      CLASS      DISKS              BLKS      FREEBLKS      SPARE
-----
group    Group1    Class2     Disk1:Disk2:Disk3:Disk4  32964608  31850496      *
```

OBJ	NAME	CLASS	GROUP	SKIP	JRM	1STBLK	LASTBLK	BLOCKS	STATUS
volume	*	Class2	Group1	*	*	0	65535	65536	PRIVATE
volume	Volume1	Class2	Group1	*	*	65536	1114111	1048576	ACTIVE
volume	*	Class2	Group1	*	*	1114112	32964607	31850496	FREE

OBJ	CLASS	GROUP	DISK	VOLUME	STATUS
slice	Class2	Group1	*	Volume1	ACTIVE

For a stripe volume, see also the stripe width.

```
# sdxinfo -G -c Class2 -o Group1 -e long
```

OBJ	NAME	CLASS	DISKS	BLKS	FREEBLKS	SPARE	MASTER	TYPE	WIDTH
group	Group1	Class2	Disk1:Disk2:Disk3:Disk4	32964608	31850496	*	*	stripe	256

6) Backing up to tape

On backup server Node3, back up data from the shadow volume to tape. In the following examples, data in shadow volume Volume1 is backed up to a tape medium of tape device /dev/rmt/0.



See

For details on the backup methods, see the manuals of file systems to be backed up and used commands.



Information

In a GFS Shared File System

Back up through the method as described in step 6a).

6a) When backing up data held in a raw device with the dd(1M) command

```
# dd if=/dev/sfdsk/Class2/rdsk/Volume1 of=/dev/rmt/0 bs=32768
```

6b) When backing up a ufs file system with the tar(1) command

6b-1) Mount the ufs file system on shadow volume Volume1 on /mnt1, a temporary mount point, in the read only mode.

```
# mkdir /mnt1
# mount -F ufs -o ro /dev/sfdsk/Class2/dsk/Volume1 /mnt1
```

6b-2) Back up data held in the file system to tape.

```
# cd /mnt1
# tar cvf /dev/rmt/0 .
```

6b-3) Unmount the file system mounted in step 6b-1).

```
# cd /  
# umount /mnt1  
# rmdir /mnt1
```

6c) When backing up a ufs file system with the ufsdump(1M) command

```
# ufsdump 0ucf /dev/rmt/0 /dev/sfdsk/Class2/rdisk/Volume1
```

6d) When backing up the sfxfs file system (GFS Local File System or SynfinityFile file system) with tar(1) command

6d-1) Activate shadow volume Volume1 in the read and write access mode (rw).

```
# sdxshadowvolume -F -c Class2 -v Volume1  
# sdxshadowvolume -N -c Class2 -v Volume1 -e mode=rw
```

6d-2) Reconfigure the partition information of the sfxfs file system on shadow volume Volume1.

```
# sfxadm /dev/sfdsk/Class2/rdisk/Volume1
```

6d-3) Mount the sfxfs file system on shadow volume Volume1 on /mnt1, a temporary mount point, in read only mode.

```
# mkdir /mnt1  
# mount -F sfxfs -o ro /dev/sfdsk/Class2/dsk/Volume1 /mnt1
```

6d-4) Back up data held in the file system to tape.

```
# cd /mnt1  
# tar cvf /dev/rmt/0 .
```

6d-5) Unmount the file system mounted in step 6d-3).

```
# cd /  
# umount /mnt1  
# rmdir /mnt1
```

6e) When backing up the sfxfs file system (GFS Local File System or SynfinityFile file system) with the sfxdump(1M) command

6e-1) Activate shadow volume Volume1 in the access mode for write and read (rw).

```
# sdxshadowvolume -F -c Class2 -v Volume1  
# sdxshadowvolume -N -c Class2 -v Volume1 -e mode=rw
```

6e-2) Reconfigure the partition information of the sfxfs file system on shadow volume Volume1.

```
# sfxadm /dev/sfdsk/Class2/rdisk/Volume1
```

6e-3) Back up data held in the file system to tape.

```
# sfxdump 0ucf /dev/rmt/0 /dev/sfdsk/Class2/rdisk/Volume1
```

7) Removing the shadow volume

After the backup process is complete, remove the shadow volume to prevent improper access to it. The following procedure must be performed on backup server Node3.

7-1) Stopping the shadow volume

Stop shadow volume Volume1.

```
# sdxshadowvolume -F -c Class2 -v Volume1
```

7-2) Removing the shadow volume

Remove shadow volume Volume1.

```
# sdxshadowvolume -R -c Class2 -v Volume1
```

7-3) Removing the shadow group

Remove shadow group Group1.

```
# sdxshadowgroup -R -c Class2 -g Group1
```

7-4) Removing the shadow disks

Remove shadow disks Disk1, Disk2, Disk3, and Disk4.

```
# sdxshadowdisk -R -c Class2 -d Disk1
# sdxshadowdisk -R -c Class2 -d Disk2
# sdxshadowdisk -R -c Class2 -d Disk3
# sdxshadowdisk -R -c Class2 -d Disk4
```

8) Resuming the services

Resume the services in the primary domain. The following procedure must be performed on the node that runs the services.

8-1) Activating the application volume

Activate application volume Volume1.

```
# sdxvolume -N -c Class1 -v Volume1
```

8-2) Reconfiguring the partition information of the sfxfs file system

When application volume Volume1 is used as the sfxfs file system, reconfigure the partition information.

```
# sfxadm /dev/sfdsk/Class1/rdsk/Volume1
```

8-3) Resuming the services

When the file system on application volume Volume1 was unmounted in step 2), mount it again.

Start the applications using Volume1.

6.6.1.7 Restoring

9) Stopping the services

Exit all applications accessing application volume Volume1 on Node1 and Node2 in the primary domain.

When Volume1 is used as a file system, it should be unmounted.

10) Stopping the application volume

To write-lock application volume Volume1, inactivate Volume1 on Node1 and Node2 in the primary domain. Execute the following command on Node1 or Node2.

```
# sdxvolume -F -c Class1 -v Volume1 -e allnodes
```

11) Viewing the configuration and status of the application volume

On Node1 or Node2 in the primary domain, view the configuration and status of application volume Volume1 that is the restore target. Confirm that Volume1 is in the STOP status. If the volume status is invalid, repair it referencing to "F.1.3 Volume Status Abnormality."

```
# sdxinfo -c Class1
OBJ      NAME      TYPE      SCOPE      SPARE
-----
class    Class1    shared    Node1:Node2 0

OBJ      NAME      TYPE      CLASS      GROUP      DEVDNAM      DEVBLKS      DEVCONNECT      STATUS
-----
disk     Disk1    stripe    Class1     Group1     c1t1d1       8380800      Node1:Node2     ENABLE
disk     Disk2    stripe    Class1     Group1     c2t1d1       8380800      Node1:Node2     ENABLE
disk     Disk3    stripe    Class1     Group1     c3t1d1       8380800      Node1:Node2     ENABLE
disk     Disk4    stripe    Class1     Group1     c4t1d1       8380800      Node1:Node2     ENABLE

OBJ      NAME      CLASS      DISKS              BLKS      FREEBLKS      SPARE
-----
group    Group1    Class1     Disk1:Disk2:Disk3:Disk4 32964608 31850496 *

OBJ      NAME      CLASS      GROUP      SKIP  JRM  1STBLK  LASTBLK      BLOCKS      STATUS
-----
volume  *         Class1    Group1     *    *    0       65535       65536      PRIVATE
volume  Volume1   Class1    Group1     *    *    65536  1114111     1048576    STOP
volume  *         Class1    Group1     *    *    1114112 32964607    31850496    FREE

OBJ      CLASS      GROUP      DISK      VOLUME      STATUS
-----
slice   Class1     Group1     *         Volume1     STOP
```

If application volume Volume1 belongs to a stripe group, also pay attention to the stripe width.

```
# sdxinfo -G -c Class1 -o Group1 -e long
OBJ      NAME      CLASS      DISKS              BLKS      FREEBLKS      SPARE      MASTER      TYPE      WIDTH
-----
group    Group1    Class1     Disk1:Disk2:Disk3:Disk4 32964608 31850496 *          *           stripe
256
```

12) Creating a shadow volume for restore

On backup server Node3, create a volume for restore (shadow volume) in the same configuration as the application volume found in step 11). The following settings are necessary on backup server Node3. A shadow volume for restore and a shadow volume for backup are common. When it has already been created, simply change the access mode as described in step 12-4).



Application volume data may be damaged if data is written into a shadow volume in incorrect configuration. Be sure to confirm that the shadow volume configuration is correct in step 12-5).

12-1) Registering shadow disks

Register disks c1t1d1, c2t1d1, c3t1d1, and c4t1d1 with shadow class Class2, and name them Disk1, Disk2, Disk3, and Disk4 respectively.

```
# sdxshadowdisk -M -c Class2 -d  
c1t1d1=Disk1,c2t1d1=Disk2,c3t1d1=Disk3,c4t1d1=Disk4
```

Point

- The disk names must correspond to the disk names assigned in step 1-1). The disk names assigned in 1-1) can be viewed in the NAME field for disk information displayed with the sdxinfo command in step 11).
- The class can be assigned any name.

12-2) Creating a shadow group

Connect shadow disks Disk1, Disk2, Disk3, and Disk4 to stripe type shadow group Group1.

```
# sdxshadowdisk -C -c Class2 -g Group1 -d Disk1,Disk2,Disk3,Disk4 -a  
type=stripe,width=256
```

Point

- If the application volume belongs to a stripe group or a concatenation group, the order of connecting shadow disks to a shadow group must correspond to the order of connecting disks to a group in step 1-2). The order of connecting disks in step 1-2) can be viewed in the DISKS field for group information displayed with the sdxinfo command in step 11).
- If the application volume belongs to a stripe group, the stripe width of a shadow group must correspond to the stripe width specified in step 1-2). The stripe width specified in step 1-2) can be viewed in the WIDTH field for group information displayed with the sdxinfo -e long command in step 11).
- The group can be assigned any name.

12-3) Creating a shadow volume

Create shadow volume Volume1 to Group1.

```
# sdxshadowvolume -M -c Class2 -g Group1 -v Volume1 -s 1048576
```

Point

- The volume must be created in the size corresponding to the volume size in step 1-3). The size of a volume created in step 1-3) can be viewed in the BLOCKS field for volume information displayed with the sdxinfo command in step 11).
- If there are multiple volumes, the corresponding shadow volumes must be created in the order of ascending values (first block numbers) in the 1STBLK field for volume information displayed with the sdxinfo command in step 11).
- The volume can be assigned any name.

12-4) Setting the access mode of the shadow volume

Activate shadow volume Volume1 in read and write access mode (rw).

```
# sdxshadowvolume -F -c Class2 -v Volume1  
# sdxshadowvolume -N -c Class2 -v Volume1 -e mode=rw
```

12-5) Viewing the shadow volume configuration

Using the `sdxinfo` command, confirm that the group configuration and the volume configuration are correct based on group information in the DISKS field, volume information in the 1STBLK field and in the BLOCKS field and so on.

```
# sdxinfo -c Class2
OBJ      NAME      TYPE      SCOPE      SPARE
-----
class    Class2    shared    Node3      0

OBJ      NAME      TYPE      CLASS      GROUP      DEVNAM      DEVBLKS      DEVCONNECT      STATUS
-----
disk     Disk1    stripe    Class2     Group1     c1t1d1      8380800      Node3            ENABLE
disk     Disk2    stripe    Class2     Group1     c2t1d1      8380800      Node3            ENABLE
disk     Disk3    stripe    Class2     Group1     c3t1d1      8380800      Node3            ENABLE
disk     Disk4    stripe    Class2     Group1     c4t1d1      8380800      Node3            ENABLE

OBJ      NAME      CLASS      DISKS              BLKS      FREEBLKS      SPARE
-----
group    Group1    Class2     Disk1:Disk2:Disk3:Disk4  32964608  31850496      *

OBJ      NAME      CLASS      GROUP      SKIP  JRM  1STBLK  LASTBLK      BLOCKS      STATUS
-----
volume  *          Class2    Group1     *    *      0       65535        65536      PRIVATE
volume  Volume1    Class2    Group1     *    *      65536   1114111      1048576    ACTIVE
volume  *          Class2    Group1     *    *      1114112 32964607    31850496    FREE

OBJ      CLASS      GROUP      DISK      VOLUME      STATUS
-----
slice    Class2     Group1     *          Volume1     ACTIVE
```

For a stripe volume, see also the stripe width.

```
# sdxinfo -G -c Class2 -o Group1 -e long
OBJ      NAME      CLASS      DISKS              BLKS      FREEBLKS      SPARE      MASTER      TYPE      WIDTH
-----
group    Group1    Class2     Disk1:Disk2:Disk3:Disk4  32964608  31850496      *          *          stripe
256
```

13) Restoring from tape

On backup server Node3, restore shadow volume data from the tape to which it was backed up in step 6). In the following examples, data held in shadow volume Volume1 is restored from a tape medium of tape device /dev/rmt/0.



For details on the restore methods, see the manuals of file systems to be restored and used commands.



In a GFS Shared File System

Restore through the method as described in step 13a).

13a) When restoring data held in a raw device with the dd(1M) command

```
# dd if=/dev/rmt/0 of=/dev/sfdsk/Class2/rdisk/Volume1 bs=32768
```

13b) When restoring a ufs file system with the tar(1) command

13b-1) Create a ufs file system on Volume1.

```
# newfs /dev/sfdsk/Class2/rdisk/Volume1
```

13b-2) Mount the ufs file system on shadow volume Volume1 on /mnt1, a temporary mount point.

```
# mkdir /mnt1
# mount -F ufs /dev/sfdsk/Class2/dsk/Volume1 /mnt1
```

13b-3) Restore data held in the file system from tape.

```
# cd /mnt1
# tar xvf /dev/rmt/0
```

13b-4) Unmount the file system mounted in step 13b-3).

```
# cd /
# umount /mnt1
# rmdir /mnt1
```

13c) When restoring a ufs file system with the ufsrestore(1M) command

13c-1) Create a ufs file system on shadow volume Volume1.

```
# newfs /dev/sfdsk/Class2/rdisk/Volume1
```

13c-2) Mount the ufs file system on shadow volume Volume1 on /mnt1, a temporary mount point.

```
# mkdir /mnt1
# mount -F ufs /dev/sfdsk/Class2/dsk/Volume1 /mnt1
```

13c-3) Restore data held in the file system from tape.

```
# cd /mnt1
# ufsrestore rvf /dev/rmt/0
```

13c-4) Delete the temporary file created by the ufsrestore(1M) command.

```
# rm /mnt1/restoresymtable
```

13c-5) Unmount the file system mounted in step 13c-2).

```
# cd /
# umount /mnt1
# rmdir /mnt1
```

13d) When restoring the sfxfs file system (GFS Local File System or SynfinityFile file system) with the tar(1) command

13d-1) Create the sfxfs file system on shadow volume Volume1.

```
# sfxnewfs /dev/sfdsk/Class2/rdisk/Volume1
```

13d-2) Mount the sfxfs file system on shadow volume Volume1 on /mnt1, a temporary mount point.

```
# mkdir /mnt1
# mount -F sfxfs /dev/sfdsk/Class2/dsk/Volume1 /mnt1
```

13d-3) Restore data held in the file system from tape.

```
# cd /mnt1
# tar xvf /dev/rmt/0
```

13d-4) Unmount the file system mounted in step 13d-2).

```
# cd /
# umount /mnt1
# rmdir /mnt1
```

13e) When restoring the sfxfs file system (GFS Local File System or SynfinityFile file system) with the sfxrestore(1M) command

13e-1) Create the sfxfs file system on shadow volume Volume1.

```
# sfxnewfs /dev/sfdsk/Class2/rdisk/Volume1
```

13e-2) Mount the sfxfs file system on shadow volume Volume1 on /mnt1, a temporary mount point.

```
# mkdir /mnt1
# mount -F sfxfs /dev/sfdsk/Class2/dsk/Volume1 /mnt1
```

13e-3) Restore data held in the file system from tape.

```
# cd /mnt1
# sfxrestore rvf /dev/rmt/0
```

13e-4) Delete the temporary file created by the sfxrestore(1M) command.

```
# rm /mnt1/restoresymtable
```

13e-5) Unmount the file system mounted in step 13e-2).

```
# cd /
# umount /mnt1
# rmdir /mnt1
```

14) Removing the shadow volume

After the restore process is complete, remove the shadow volume to prevent improper access to it. The following procedure must be performed on backup server Node3.

14-1) Stopping the shadow volume

Stop shadow volume Volume1.

```
# sdxshadowvolume -F -c Class2 -v Volume1
```

14-2) Removing the shadow volume

Remove shadow volume Volume1.

```
# sdxshadowvolume -R -c Class2 -v Volume1
```

14-3) Removing the shadow group

Remove shadow group Group1.

```
# sdxshadowgroup -R -c Class2 -g Group1
```

14-4) Removing the shadow disks

Remove shadow disks Disk1, Disk2, Disk3, and Disk4.

```
# sdxshadowdisk -R -c Class2 -d Disk1
# sdxshadowdisk -R -c Class2 -d Disk2
# sdxshadowdisk -R -c Class2 -d Disk3
# sdxshadowdisk -R -c Class2 -d Disk4
```

15) Resuming the services

Resume the services in the primary domain. The following procedure must be performed on the node that runs the services. Step 15-2) is necessary only when the application volume is used as the sfxfs file system.

15-1) Activating the application volume

Activate application volume Volume1.

```
# sdxvolume -N -c Class1 -v Volume1
```

15-2) Reconfiguring the partition information of the sfxfs file system

When application volume Volume1 is used as the sfxfs file system, reconfigure the partition information.

```
# sfxadm /dev/sfdisk/Class1/rdisk/Volume1
```

15-3) Resuming the services

When the file system on application volume Volume1 was unmounted in step 9), mount it again.

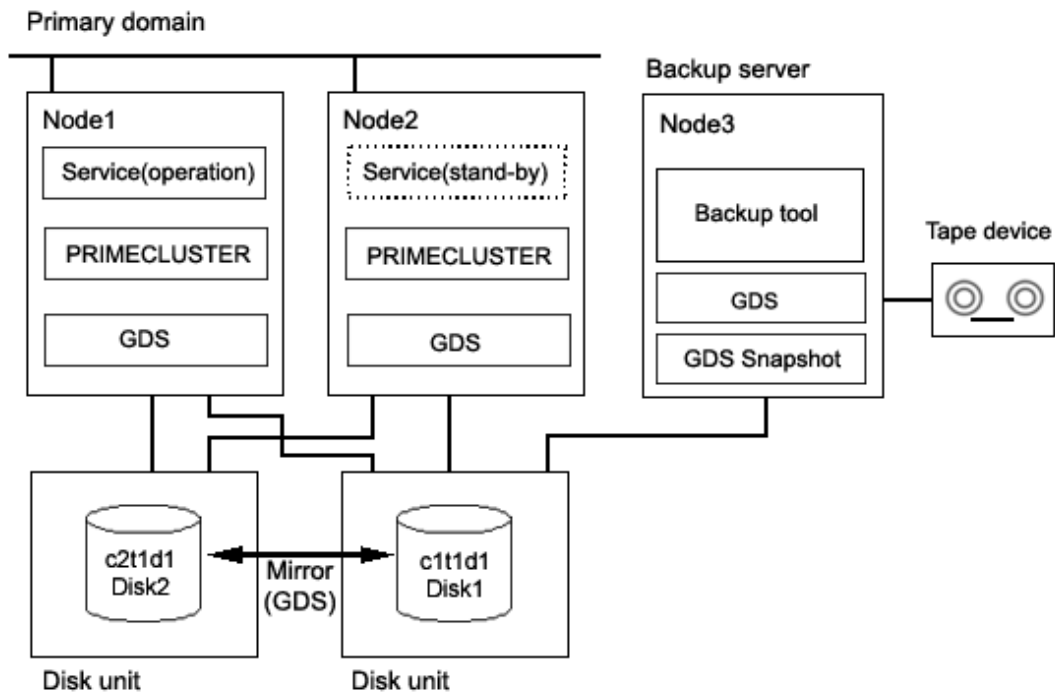
Start the applications using Volume1.

6.6.2 Backing Up and Restoring through Snapshot by Slice Detachment

This sub-section describes the method of backing up data from and restoring data back to mirror volumes in the primary domain through a backup server in another domain.

6.6.2.1 System Configuration

Figure 6.25 System Configuration



Information

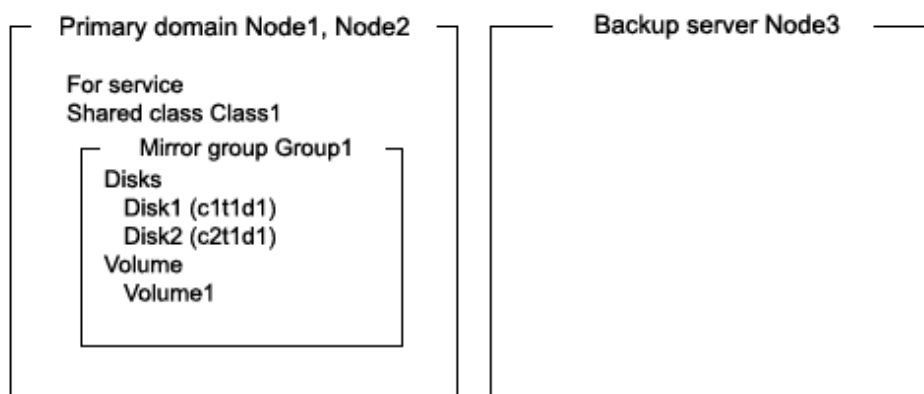
In the primary domain, SynfinityDisk and SynfinityCluster are also available instead of GDS and PRIMECLUSTER.

Note

Physical Device Name

Different physical device names (such as c1t1d1) may be assigned to the identical physical disk in the primary domain and the backup server.

Figure 6.26 Object Configuration in Normal Operation



6.6.2.2 Summary of Backup

Data in a slice temporarily detached from a volume can be backed up to tape during the service operation.

To secure consistency of data in a detached slice, the services must be stopped temporarily when detaching the slice.

Information

Consistency of Snapshot Data

When detaching a slice while the services are operating, data consistency must be secured through the method specific to that software, such as a file system and a database system, which manages volume data. For details, see "[A.2.22 Ensuring Consistency of Snapshot Data](#)."

Figure 6.27 Backup

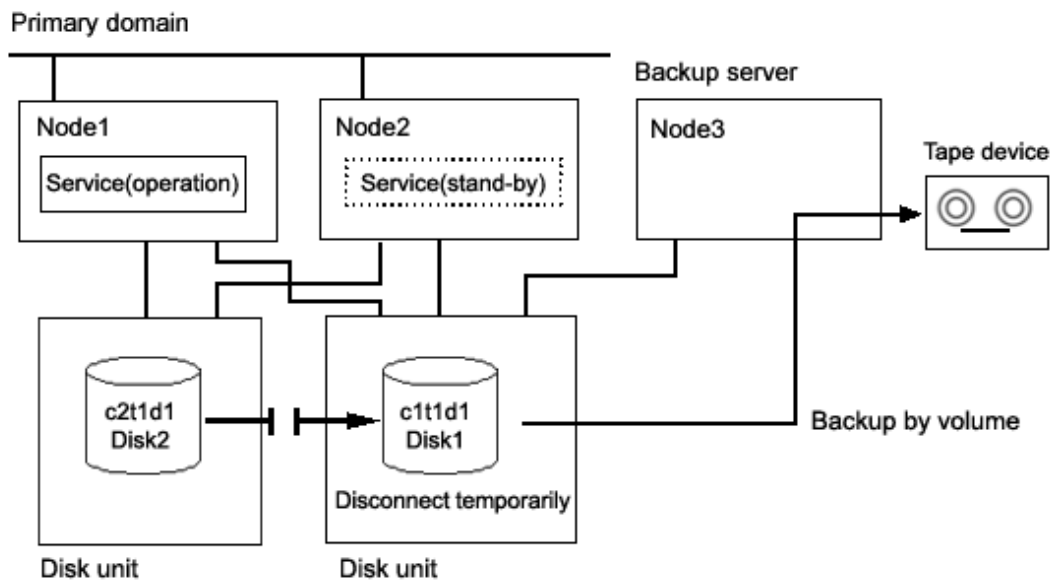


Figure 6.28 Object Configuration When Backing Up





Figure 6.29 Backup Schedule

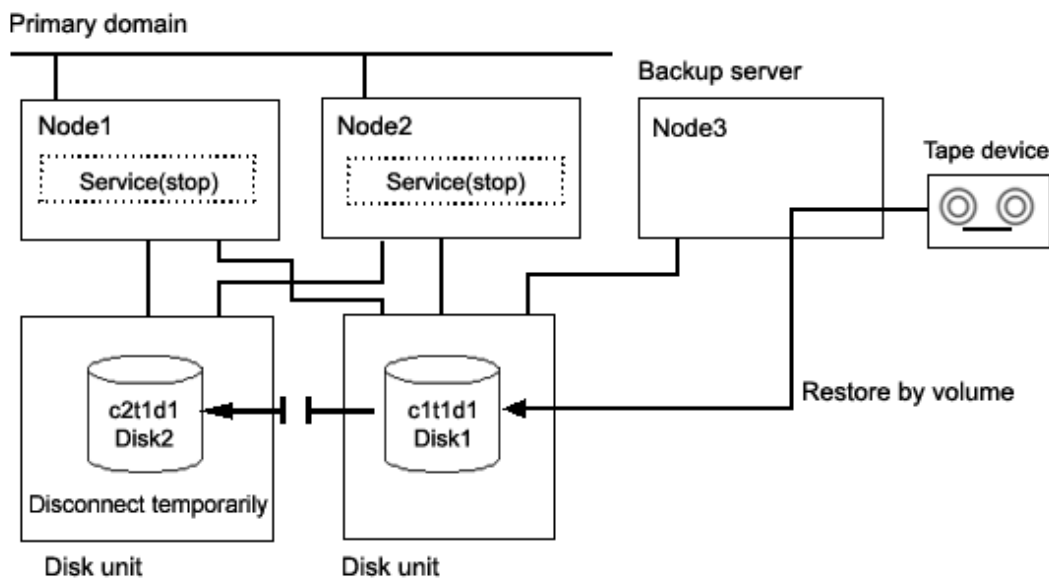


6.6.2.3 Summary of Restore

If volume data is damaged, it can be restored from tape.

Data can be restored while service is stopped and the application volume is not in use.

Figure 6.30 Restore



Information

In this configuration access cannot be gained from backup server Node3 to disk c2t1d1. Therefore, after data is restored from tape back to c1t1d1 while c2t1d1 is detached temporarily, resynchronization copying from c1t1d1 to c2t1d1 must be performed by reattaching c2t1d1. When access can be gained from Node3 to both c1t1d1 and c2t1d1, it is not required that c2t1d1 be detached temporarily since data can be restored from tape back to both c1t1d1 and c2t1d1. For details on this restore method, see "[6.6.1 Backing Up and Restoring a Logical Volume with No Replication](#)"

Figure 6.31 Object Configuration When Restoring

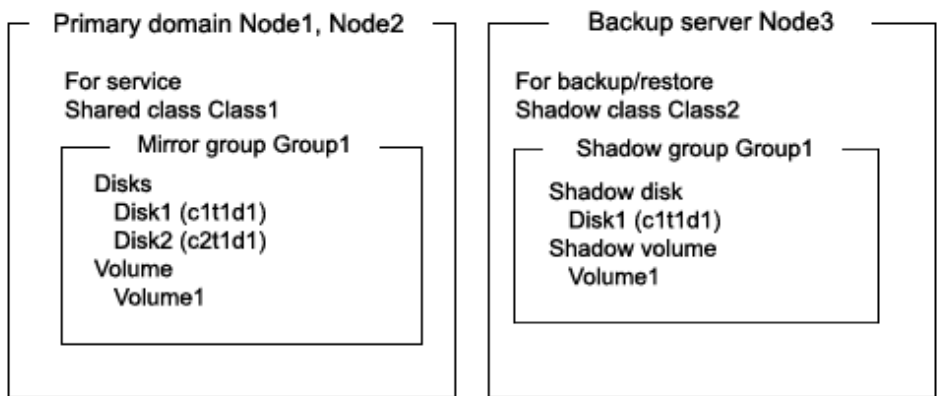
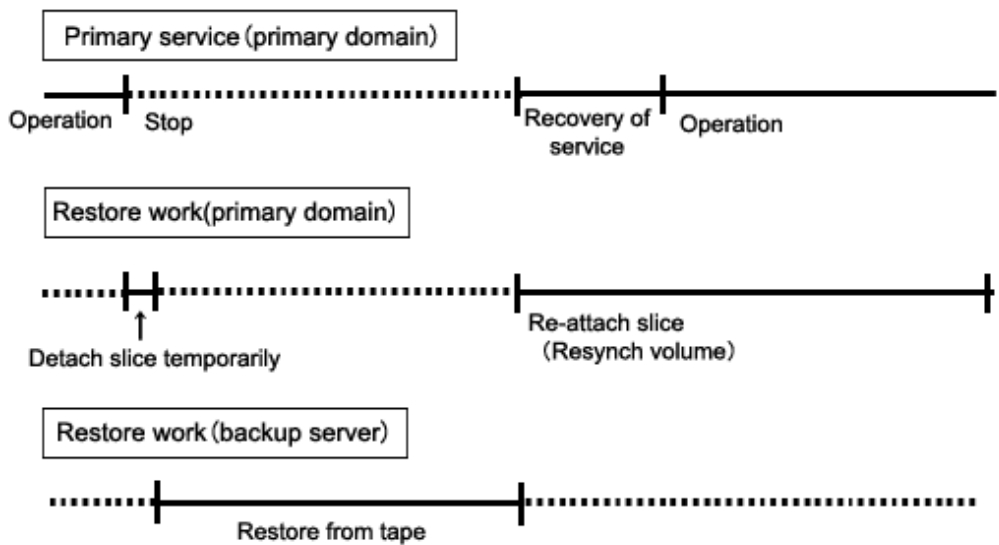


Figure 6.32 Restore Schedule



6.6.2.4 Summary of Procedure

Figure 6.33 Outline of the Configuration Procedure



Figure 6.34 Outline of the Backup Procedure





Figure 6.35 Outline of the Restore Procedure



6.6.2.5 Configuring an Environment



Note

Automatic Resource Registration

If the backup server resides in a cluster domain (called a backup domain), those disks that are registered as resources in the primary domain or are to be registered with a shadow class in the backup domain may not be involved in the resource registration in the backup domain. In the backup domain, those relevant disks must be described in the Excluded Device List prior to executing the automatic resource registration. For details on the automatic resource registration, see "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide."

1) Creating an application volume

Create a mirror volume used for the services on disks c1t1d1 and c2t1d1. The following settings are necessary on Node1 or Node2 in the primary domain.

1-1) Registering disks

Register disks c1t1d1 and c2t1d1 with shared class Class1 that is shared on Node1 and Node2, and name them Disk1 and Disk2 respectively.

```
# sdxdisk -M -c Class1 -a type=shared,scope=Node1:Node2 -d  
c1t1d1=Disk1,c2t1d1=Disk2
```

1-2) Creating a mirror group

Connect Disk1 and Disk2 to mirror group Group1.

```
# sdxdisk -C -c Class1 -g Group1 -d Disk1,Disk2
```

1-3) Creating a mirror volume

Create mirror volume Volume1 to mirror group Group1.

```
# sdxvolume -M -c Class1 -g Group1 -v Volume1 -s 1048576
```

6.6.2.6 Backing Up

2) Detaching the backup target slice

Temporarily detach the slice on Disk1 that is the backup target, among slices in application volume Volume1. The following procedure must be performed on Node1 or Node2 in the primary domain.



Information

The following example secures data consistency by stopping the services when a slice is detached. Steps 2-1) and 2-3) are not required if your software, such as a file system and a database system, that manages volume data provides functionality ensuring data consistency or repairing consistency for a detached slice. Alternatively, data consistency must be secured with the method specific to that software. For details, see "[A.2.22 Ensuring Consistency of Snapshot Data.](#)"

2-1) Stopping the services

To secure consistency of data in a detached slice, exit all applications accessing application volume Volume1 on Node1 and Node2.

When Volume1 is used as a file system, it should be unmounted.

2-2) Detaching the slice

Temporarily detach the slice on disk Disk1 from Volume1. To write-lock the detached slice, set the access mode of the slice to ro (read only).

```
# sdxslice -M -c Class1 -d Disk1 -v Volume1 -a jrm=off,mode=ro
```



Note

Just Resynchronization Mode for Silice

On backup server Node3, data may be written from Node3 into Disk1 when data in Disk1 is backed up to tape. GDS in the primary domain cannot recognize the write occurrence from Node3. Consequently, if the JRM mode of the detached slice is "on", the portions updated

from Node3 may not be involved in resynchronization copying performed when the slice is reattached. If this happens, synchronization of Volume1 is no longer ensured. For this reason, the JRM mode of a detached slice must be set to off in advance.

2-3) Resuming the services

When the file system was unmounted in step 2-1), mount it again.

Resume the application stopped in step 2-1).

3) Viewing the configuration of the application volume

On Node1 or Node2 in the primary domain, view the configuration of application volume Volume1 that is the backup target.

```
# sdxinfo -c Class1
OBJ      NAME      TYPE      SCOPE      SPARE
-----
class    Class1    shared    Node1:Node2 0

OBJ      NAME      TYPE      CLASS      GROUP      DEVNAM      DEVBLKS      DEVCONNECT      STATUS
-----
disk     Disk1    mirror    Class1     Group1     c1t1d1      8380800      Node1:Node2     ENABLE
disk     Disk2     mirror    Class1     Group1     c2t1d1      8380800      Node1:Node2     ENABLE

OBJ      NAME      CLASS      DISKS      BLKS      FREEBLKS      SPARE
-----
group    Group1    Class1     Disk1:Disk2      8290304      7176192      0

OBJ      NAME      CLASS      GROUP      SKIP      JRM      1STBLK      LASTBLK      BLOCKS      STATUS
-----
volume   *         Class1     Group1     *         *         0           65535         65536        PRIVATE
volume   Volume1   Class1     Group1     off        on        65536      1114111      1048576     ACTIVE
volume   *         Class1     Group1     *         *         1114112     8290303      7176192      FREE

OBJ      CLASS      GROUP      DISK      VOLUME      STATUS
-----
slice    Class1     Group1     Disk1    Volume1     TEMP
slice    Class1     Group1     Disk2     Volume1     ACTIVE
```

4) Creating a shadow volume for backup

Create a volume for backup (shadow volume) to disk c1t1d1 on backup server Node3. The following settings are necessary on backup server Node3.



Application volume data may be damaged if data is written into a shadow volume in incorrect configuration. Be sure to confirm that the shadow volume configuration is correct in step 4-4).

4-1) Registering a shadow disk

Register disk c1t1d1 with shadow class Class2, and name it Disk1.

```
# sdxshadowdisk -M -c Class2 -d c1t1d1=Disk1
```

Point

- The disk name must correspond to the disk name assigned to disk c1t1d1 in step 1-1). The disk names assigned in 1-1) can be viewed in the NAME field for disk information displayed with the `sdxinfo` command in step 3).
- The class can be assigned any name.

4-2) Creating a shadow group

Connect shadow disk Disk1 to mirror type shadow group Group1.

```
# sdxshadowdisk -C -c Class2 -g Group1 -d Disk1
```

4-3) Creating a shadow volume

Create shadow volume Volume1 to shadow group Group1.

```
# sdxshadowvolume -M -c Class2 -g Group1 -v Volume1 -s 1048576
```

Point

- The volume must be created in the size corresponding to the volume size in step 1-3). The size of a volume created in step 1-3) can be viewed in the BLOCKS field for volume information displayed with the `sdxinfo` command in step 3).
- If there are multiple volumes, the corresponding shadow volumes must be created in the order of ascending values (first block numbers) in the 1STBLK field for volume information displayed with the `sdxinfo` command in step 3).
- The volume can be assigned any name.

4-4) Viewing the configuration of the shadow volume

Using the `sdxinfo` command, confirm that the group configuration and the volume configuration are correct based on group information in the DISKS field, volume information in the 1STBLK field and in the BLOCKS field and so on.

```
# sdxinfo -c Class2
OBJ  NAME  TYPE  SCOPE  SPARE
-----
class Class2 local  Node3  0

OBJ  NAME  TYPE  CLASS  GROUP  DEVNAM  DEVBLKS  DEVCONNECT  STATUS
-----
disk  Disk1 mirror Class2  Group1  c1t1d1  8380800  Node3      ENABLE

OBJ  NAME  CLASS  DISKS  BLKS  FREEBLKS  SPARE
-----
group Group1  Class2  Disk1  8290304  7176192  0

OBJ  NAME  CLASS  GROUP  SKIP  JRM  1STBLK  LASTBLK  BLOCKS  STATUS
-----
volume *      Class2  Group1  *      *      0      65535    65536    PRIVATE
volume Volume1 Class2  Group1  off    off    65536  1114111  1048576  ACTIVE
volume *      Class2  Group1  *      *      1114112 8290303  7176192  FREE
```

OBJ	CLASS	GROUP	DISK	VOLUME	STATUS
slice	Class2	Group1	<u>Disk1</u>	Volume1	ACTIVE

5) Backing up to tape

On backup server Node3, back up data in the shadow volume to tape. The following shows examples of backing up data in shadow volume Volume1 to a tape medium of tape device /dev/rmt/0.



For details on the backup method, see the manuals of file systems to be backed up and used command.

Information

In a GFS Shared File System

Back up through the method as described in step 5a).

5a) When backing up data held in a raw device with the dd(1M) command

```
# dd if=/dev/sfdsk/Class2/rdsk/Volume1 of=/dev/rmt/0 bs=32768
```

5b) When backing up a ufs file system with the tar(1) command

5b-1) Activate shadow volume Volume1 in the read and write access mode (rw).

```
# sdxshadowvolume -F -c Class2 -v Volume1
# sdxshadowvolume -N -c Class2 -v Volume1 -e mode=rw
```

5b-2) Check and repair consistency of the ufs file system on shadow volume Volume1.

If the file system was unmounted when the slice was detached in step 2), this step can be skipped.

```
# fsck -F ufs -y /dev/sfdsk/Class2/rdsk/Volume1
```

5b-3) Mount the ufs file system on shadow volume Volume1 on /mnt1, a temporary mount point, in the read only mode.

```
# mkdir /mnt1
# mount -F ufs -o ro /dev/sfdsk/Class2/dsk/Volume1 /mnt1
```

5b-4) Back up data held in the file system to tape.

```
# cd /mnt1
# tar cvf /dev/rmt/0 .
```

5b-5) Unmount the file system mounted in step 5b-3).

```
# cd /
# umount /mnt1
# rmdir /mnt1
```

5c) When backing up a ufs file system with the ufsdump(1M) command

5c-1) Activate shadow volume Volume1 in the read and write access mode (rw).

```
# sdxshadowvolume -F -c Class2 -v Volume1
# sdxshadowvolume -N -c Class2 -v Volume1 -e mode=rw
```

5c-2) Check and repair consistency of the ufs file system on shadow volume Volume1.

If the file system was unmounted when the slice was detached in step 2), this step can be skipped.

```
# fsck -F ufs -y /dev/sfdsk/Class2/rdsk/Volume1
```

5c-3) Back up data held in the file system to tape.

```
# ufsdump 0ucf /dev/rmt/0 /dev/sfdsk/Class2/rdsk/Volume1
```

5d) When backing up the sfxfs file system (GFS Local File System or SynfinityFile file system) with the tar(1) command

5d-1) Activate shadow volume Volume1 in the read and write access mode (rw).

```
# sdxshadowvolume -F -c Class2 -v Volume1
# sdxshadowvolume -N -c Class2 -v Volume1 -e mode=rw
```

5d-2) Reconfigure the partition information of the sfxfs file system on shadow volume Volume1.

```
# sfxadm /dev/sfdsk/Class2/rdsk/Volume1
```

5d-3) Check and repair consistency of the sfxfs file system on shadow volume Volume1.

If the file system was unmounted when the slice was detached in step 2), this step can be skipped.

```
# fsck -F sfxfs -y /dev/sfdsk/Class2/rdsk/Volume1
```

5d-4) Mount the sfxfs file system on shadow volume Volume1 on /mnt1, a temporary mount point, in the read only mode.

```
# mkdir /mnt1
# mount -F sfxfs -o ro /dev/sfdsk/Class2/dsk/Volume1 /mnt1
```

5d-5) Back up data held in the file system to tape.

```
# cd /mnt1
# tar cvf /dev/rmt/0 .
```

5d-6) Unmount the file system mounted in step 5d-4).

```
# cd /
# umount /mnt1
# rmdir /mnt1
```

5e) When backing up the sfxfs file system (GFS Local File System or SynfinityFile file system) with the sfxdump(1M) command

5e-1) Activate shadow volume Volume1 in the read and write access mode (rw).

```
# sdxshadowvolume -F -c Class2 -v Volume1
# sdxshadowvolume -N -c Class2 -v Volume1 -e mode=rw
```

5e-2) Reconfigure the partition information of the sfxfs file system on shadow volume Volume1.

```
# sfxadm /dev/sfdsk/Class2/rdsk/Volume1
```

5e-3) Check and repair consistency of the sfxfs file system on shadow volume Volume1.

If the file system was unmounted when the slice was detached in step 2), this step can be skipped.

```
# fsck -F sfxfs -y /dev/sfdsk/Class2/rdsk/Volume1
```

5e-4) Back up data held in the file system to tape.

```
# sfxdump 0ucf /dev/rmt/0 /dev/sfdsk/Class2/rdsk/Volume1
```

6) Removing the shadow volume

After the backup process is complete, remove the shadow volume to prevent improper access to it. The following procedure must be performed on backup server Node3.

6-1) Stopping the shadow volume

Stop shadow volume Volume1.

```
# sdxshadowvolume -F -c Class2 -v Volume1
```

6-2) Removing the shadow volume

Remove shadow volume Volume1.

```
# sdxshadowvolume -R -c Class2 -v Volume1
```

6-3) Removing the shadow group

Remove shadow group Group1.

```
# sdxshadowgroup -R -c Class2 -g Group1
```

6-4) Removing the shadow disk

Remove shadow disk Disk1.

```
# sdxshadowdisk -R -c Class2 -d Disk1
```

7) Reattaching the backup target slice

Reattach the slice temporarily detached from the application volume back to it. The following procedure must be performed on Node1 or Node2 in the primary domain.

7-1) Reattaching the backup target slice

Reattach slice Volume1.Disk1 temporarily detached from application volume Volume1 in step 2-2).

```
# sdxslice -R -c Class1 -d Disk1 -v Volume1
```

After returning from the command, synchronization copying from the slice on Disk1 of volume Volume1 to the slice on Disk2 is executed.

7-2) Viewing the copy status

The status of synchronization copying can be viewed using the `sdxinfo -S` command. The copy destination slice is in the `COPY` status if copying is in progress and it will be in the `ACTIVE` status after the copy process ends normally (note, however, that it will be in the `STOP` status when Volume1 is in the `STOP` status).

```
# sdxinfo -S -c Class1 -o Volume1
OBJ      CLASS   GROUP   DISK     VOLUME   STATUS
-----
slice    Class1  Group1  Disk1    Volume1  ACTIVE
slice    Class1  Group1  Disk2    Volume1  COPY
```

6.6.2.7 Restoring

Information

If access can be gained from backup server Node3 to all of the disks constituting Volume1 (c1t1d1 and c2t1d1) data can be restored from tape back to both c1t1d1 and c2t1d1 on Node3. Under these circumstances, detaching a slice should not be performed as described in step 10).

8) Stopping the services

Exit all applications using application volume Volume1 on Node1 and Node2 in the primary domain.

When Volume1 is used as a file system, it should be unmounted.

9) Stopping the application volume

To write-lock volume Volume1, inactivate Volume1 on Node1 and Node2 in the primary domain. Execute the following command on Node1 or Node2.

```
# sdxvolume -F -c Class1 -v Volume1 -e allnodes
```

10) Detaching any nonrelevant slice from the application volume

Temporarily detach the slice on any disk (Disk2) other than Disk1 that is the restore target from Volume1, among slices in application volume Volume1. Execute the following command on Node1 or Node2 in the primary domain.

```
# sdxslice -M -c Class1 -d Disk2 -v Volume1 -a jrm=off
```

Note

Just Resynchronization Mode for Slice

On backup server Node3, after data is restored from tape back to Disk1, the slice on Disk2 is supposed to be reattached to application volume Volume1 in the primary domain. At this point the entire volume data must be copied to the attached slice. For this reason, the JRM mode of a detached slice must be set to off in advance.

Information

If access can be gained from backup server Node3 to all of the disks constituting Volume1 (c1t1d1 and c2t1d1) this procedure (detaching a slice) should not be performed.

11) Viewing the configuration and status of the application volume

On Node1 or Node2 in the primary domain, view the configuration and status of application volume Volume1 that is the restore target. Confirm that Volume1 is in STOP status and that only restore target slice Volume1.Disk1 is in STOP status among the slices constituting the volume and the other slices are in TEMP or TEMP-STOP status. If the volume or slice status is invalid, repair it referencing to "F.1.3 Volume Status Abnormality" and "F.1.1 Slice Status Abnormality."

```
# sdxinfo -c Class1
OBJ      NAME      TYPE      SCOPE      SPARE
-----
class    Class1    shared    Node1:Node2 0

OBJ      NAME      TYPE      CLASS      GROUP      DEVNAM      DEVBLKS      DEVCONNECT      STATUS
-----
disk     Disk1    mirror    Class1     Group1     c1t1d1      8380800      Node1:Node2     ENABLE
disk     Disk2     mirror    Class1     Group1     c2t1d1      8380800      Node1:Node2     ENABLE

OBJ      NAME      CLASS      DISKS              BLKS      FREEBLKS      SPARE
-----
group    Group1    Class1     Disk1:Disk2        8290304    7176192      0

OBJ      NAME      CLASS      GROUP      SKIP  JRM      1STBLK      LASTBLK      BLOCKS      STATUS
-----
volume  *          Class1     Group1     *     *          0           65535         65536      PRIVATE
volume  Volume1    Class1     Group1     off   on        65536       1114111       1048576    STOP
volume  *          Class1     Group1     *     *        1114112     8290303       7176192    FREE

OBJ      CLASS      GROUP      DISK      VOLUME      STATUS
-----
slice    Class1     Group1     Disk1    Volume1     STOP
slice    Class1     Group1     Disk2     Volume1     TEMP
```

Information

If access can be gained from backup server Node3 to all of the disks constituting Volume1 (c1t1d1 and c2t1d1) you must confirm that all of the slices of Volume1 are in STOP status.

12) Creating a shadow volume for restore

On backup server Node3, create a volume for restore (shadow volume) on disk c1t1d1. The following settings are necessary on backup server Node3. A shadow volume for restore and a shadow volume for backup are common. If one already exists, this procedure is not required.

Note

Application volume data may be damaged if data is written into a shadow volume in incorrect configuration. Be sure to confirm that the shadow volume configuration is correct in step 12-5).

Information

If access can be gained from backup server Node3 to all of the disks constituting Volume1 (c1t1d1 and c2t1d1) a shadow volume for restore must be created in the same configuration as Volume1. Under these circumstances, those shadow volumes for restore and backup are not common.

12-1) Registering a shadow disk

Register disk c1t1d1 with shadow class Class2, and name it Disk1.

```
# sdxshadowdisk -M -c Class2 -d c1t1d1=Disk1
```

Point

- The disk name must correspond to the disk name assigned to c1t1d1 in step 1-1). The disk names assigned in 1-1) can be viewed in the NAME field for disk information displayed with the sdxinfo command in step 11).
- The class can be assigned any name.

Information

If access can be gained from backup server Node3 to all of the disks constituting Volume1 (c1t1d1 and c2t1d1) all of those disks (c1t1d1 and c2t1d1) must be registered with a shadow class.

12-2) Creating a shadow group

Connect shadow disk Disk1 to mirror type shadow group Group1.

```
# sdxshadowdisk -C -c Class2 -g Group1 -d Disk1
```

Information

If access can be gained from backup server Node3 to all of the disks constituting Volume1 (c1t1d1 and c2t1d1) all of those disks (c1t1d1 and c2t1d1) must be connected to a shadow group.

12-3) Creating a shadow volume

Create shadow volume Volume1 to shadow group Group1.

```
# sdxshadowvolume -M -c Class2 -g Group1 -v Volume1 -s 1048576
```

Point

- The volume must be created in the size corresponding to the volume size in step 1-3). The size of a volume created in step 1-3) can be viewed in the BLOCKS field for volume information displayed with the sdxinfo command in step 11).
- If there are multiple volumes, the corresponding shadow volumes must be created in the order of ascending values (first block numbers) in the 1STBLK field for volume information displayed with the sdxinfo command in step 11).
- The volume can be assigned any name.

12-4) Setting the access mode of the shadow volume

Activate shadow volume Volume1 in the read and write access mode (rw).

```
# sdxshadowvolume -F -c Class2 -v Volume1
# sdxshadowvolume -N -c Class2 -v Volume1 -e mode=rw
```

12-5) Viewing the configuration of the shadow volume

Using the `sdxinfo` command, confirm that the group configuration and the volume configuration are correct based on group information in the DISKS field, volume information in the 1STBLK field and in the BLOCKS field and so on.

```
# sdxinfo -c Class2
OBJ      NAME      TYPE      SCOPE      SPARE
-----
class    Class2    local     Node3      0

OBJ      NAME      TYPE      CLASS      GROUP      DEVNAM      DEVBLKS      DEVCONNECT      STATUS
-----
disk     Disk1    mirror    Class2     Group1     c1t1d1      8380800      Node3           ENABLE

OBJ      NAME      CLASS      DISKS      BLKS      FREEBLKS      SPARE
-----
group    Group1    Class2     Disk1      8290304    7176192      0

OBJ      NAME      CLASS      GROUP      SKIP      JRM      1STBLK      LASTBLK      BLOCKS      STATUS
-----
volume   *         Class2     Group1     *         *         0           65535        65536      PRIVATE
volume   Volume1   Class2     Group1     off      off      65536      1114111      1048576    ACTIVE
volume   *         Class2     Group1     *         *         1114112    8290303      7176192    FREE

OBJ      CLASS      GROUP      DISK      VOLUME      STATUS
-----
slice    Class2     Group1     Disk1    Volume1     ACTIVE
```

13) Restoring from tape

On backup server Node3, restore shadow volume data from tape to which it was backed up in step 5). In the following examples, restore data held in shadow volume Volume1 from a tape medium of tape device /dev/rmt/0.



See

For details on the restore method, see the manuals of file systems to be restored and used commands.



Information

In a GFS Shared File System

Restore through the method as described in step 13a).

13a) When restoring data held in a raw device with the dd(1M) command

```
# dd if=/dev/rmt/0 of=/dev/sfdsk/Class2/rdisk/Volume1 bs=32768
```

13b) When restoring a ufs file system with the tar(1) command

13b-1) Create a ufs file system to shadow volume Volume1.

```
# newfs /dev/sfdsk/Class2/rdsk/Volume1
```

13b-2) Mount the ufs file system on shadow volume Volume1 on /mnt1, a temporary mount point.

```
# mkdir /mnt1
# mount -F ufs /dev/sfdsk/Class2/dsk/Volume1 /mnt1
```

13b-3) Restore data held in the file system from tape.

```
# cd /mnt1
# tar xvf /dev/rmt/0
```

13b-4) Unmount the file system mounted in step 13b-3).

```
# cd /
# umount /mnt1
# rmdir /mnt1
```

13c) When restoring a ufs file system with the ufsrestore(1M) command

13c-1) Create a ufs file system to shadow volume Volume1.

```
# newfs /dev/sfdsk/Class2/rdsk/Volume1
```

13c-2) Mount the ufs file system on shadow volume Volume1 on /mnt1, a temporary mount point.

```
# mkdir /mnt1
# mount -F ufs /dev/sfdsk/Class2/dsk/Volume1 /mnt1
```

13c-3) Restore data held in the file system from tape.

```
# cd /mnt1
# ufsrestore rvf /dev/rmt/0
```

13c-4) Delete the temporary file created by the ufsrestore(1M) command.

```
# rm /mnt1/restoresytable
```

13c-5) Unmount the file system mounted in step 13c-2).

```
# cd /
# umount /mnt1
# rmdir /mnt1
```

13d) When restoring the sfxfs file system (GFS Local File System or SynfinityFile file system) with the tar(1) command

13d-1) Create the sfxfs file system System to shadow volume Volume1.

```
# sfxnewfs /dev/sfdsk/Class2/rdsk/Volume1
```

13d-2) Mount the sfxfs file system on shadow volume Volume1 on /mnt1, a temporary mount point.

```
# mkdir /mnt1
# mount -F sfxfs /dev/sfdsk/Class2/dsk/Volume1 /mnt2
```

13d-3) Restore data held in the file system from tape.

```
# cd /mnt1
# tar xvf /dev/rmt/0
```

13d-4) Unmount the file system mounted in step 13d-2).

```
# cd /
# umount /mnt1
# rmdir /mnt1
```

13e) When restoring the sfxfs file system (GFS Local File System or SynfinityFile file system) with the sfxrestore(1M) command

13e-1) Create the sfxfs file system to shadow volume Volume1.

```
# sfxnewfs /dev/sfdsk/Class2/rdisk/Volume1
```

13e-2) Mount the sfxfs file system on shadow volume Volume1 on /mnt1, a temporary mount point.

```
# mkdir /mnt1
# mount -F sfxfs /dev/sfdsk/Class2/dsk/Volume1 /mnt1
```

13e-3) Restore data held in the file system from tape.

```
# cd /mnt1
# sfxrestore rvf /dev/rmt/0
```

13e-4) Delete the temporary file created by the sfxrestore(1M) command.

```
# rm /mnt1/restoresymtable
```

13e-5) Unmount the file system mounted in step 13e-2).

```
# cd /
# umount /mnt1
# rmdir /mnt1
```

14) Removing the shadow volume

After the restore process is complete, remove the shadow volume to prevent improper access to it. The following procedure must be performed on backup server Node3.

14-1) Stopping the shadow volume

Stop shadow volume Volume1.

```
# sdxshadowvolume -F -c Class2 -v Volume1
```

14-2) Removing the shadow volume

Remove shadow volume Volume1.

```
# sdxshadowvolume -R -c Class2 -v Volume1
```

14-3) Removing the shadow group

Remove shadow group Group1.

```
# sdxshadowgroup -R -c Class2 -g Group1
```

14-4) Removing the shadow disk

Remove shadow disk Disk1.

```
# sdxshadowdisk -R -c Class2 -d Disk1
```



Information

.....
If access can be gained from backup server Node3 to all of the disks constituting Volume1 (c1t1d1 and c2t1d1) all of the disks registered with shadow class Class2 in step 12) (c1t1d1 and c2t1d1) must be removed.
.....

15) Resuming the services and reattaching the slice to the application volume

Resume service in the primary domain. The following procedure should be performed on the node that runs the services.



Information

.....
In the following example resuming the services is put above resynchronizing the application volume. Through this procedure the service is resumed first and then resynchronization of the volume is secured during the services operation. If resynchronizing the volume should be put above resuming the services, the procedure should be followed in the order of steps 15-1), 15-2), 15-4), 15-5) (confirming that the synchronization copying is complete), and 15-3).
.....

15-1) Activating the application volume

Activate application volume Volume1.

```
# sdxvolume -N -c Class1 -v Volume1
```

15-2) Reconfiguring the partition information of the sfxfs file system

When application volume Volume1 is used as the sfxfs file system, reconfigure the partition information.

```
# sfxadm /dev/sfdsk/Class1/rdsk/Volume1
```

15-3) Resuming the services

When the file system on application volume Volume1 was unmounted in step 8), mount it again.

Start the applications using Volume1.

15-4) Reattaching the slice to the application volume

Reattach slice Volume1.Disk2 that was temporarily detached from application volume Volume1 in step 10) back to Volume1.

```
# sdxslice -R -c Class1 -d Disk2 -v Volume1
```

After returning from the command, synchronization copying from the slice on Disk1 of volume Volume1 to the slice on Disk2 is executed.

15-5) Viewing the copy status

The status of synchronization copying can be viewed using the `sdxinfo -S` command. The copy destination slice is in COPY status if copying is in progress and it will be in ACTIVE status after the copy process ends normally (note, however, that it will be in STOP status when Volume1 is in STOP status).

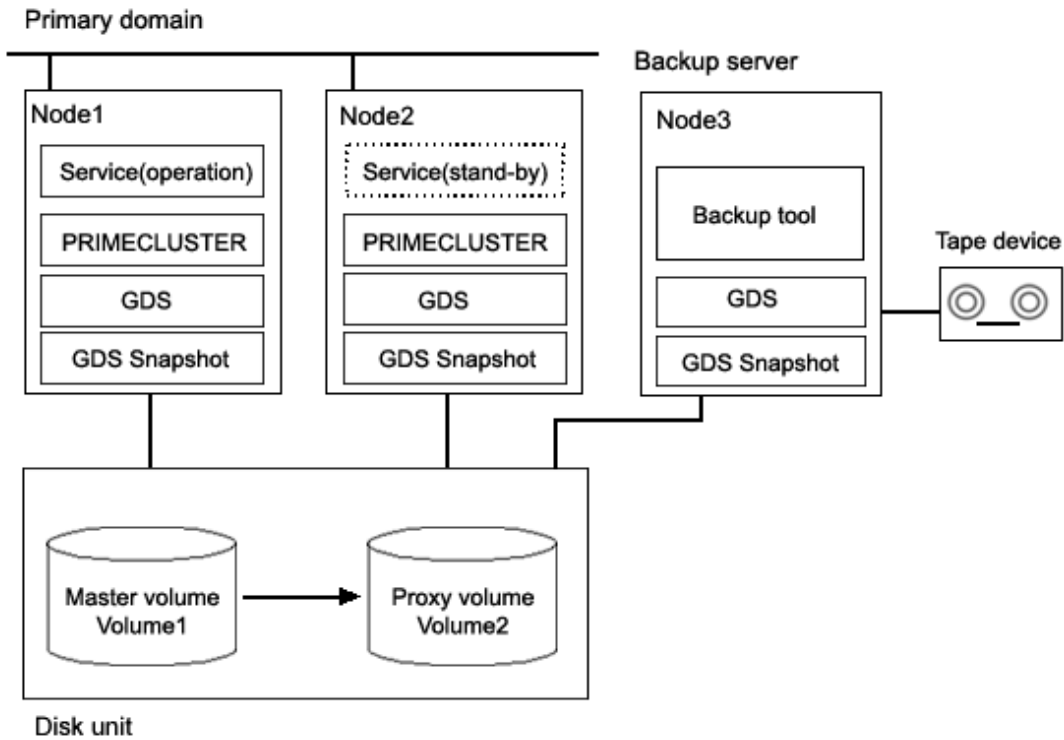
```
# sdxinfo -S -c Class1 -o Volume1
OBJ    CLASS  GROUP  DISK    VOLUME  STATUS
-----
slice  Class1  Group1 Disk1    Volum1  ACTIVE
slice  Class1  Group1 Disk2    Volum1  COPY
```

6.6.3 Backing Up and Restoring Using Snapshots from a Proxy Volume

This sub-section describes the method of backing up data from and restoring data back to logical volumes in the primary domain through a backup server in another domain by use of snapshots from the proxy volume.

6.6.3.1 System Configuration


Figure 6.36 System Configuration



Information

In the primary domain, SynfinityDisk and SynfinityCluster are also available instead of GDS , GDS Snapshot and PRIMECLUSTER.

Figure 6.37 Object Configuration in Normal Operation



6.6.3.2 Summary of Backup

Data in a proxy volume parted from the master volume during the services operation can be backed up to tape.

To secure consistency of data in the proxy volume to be parted, the services should be stopped temporarily when it is parted.

Information

Consistency of Snapshot Data

When detaching a proxy volume while the service is operating, data consistency must be secured through the method specific to that software, such as a file system and a database system, which manages volume data. For details, see "[A.2.22 Ensuring Consistency of Snapshot Data.](#)"

Figure 6.38 Backup

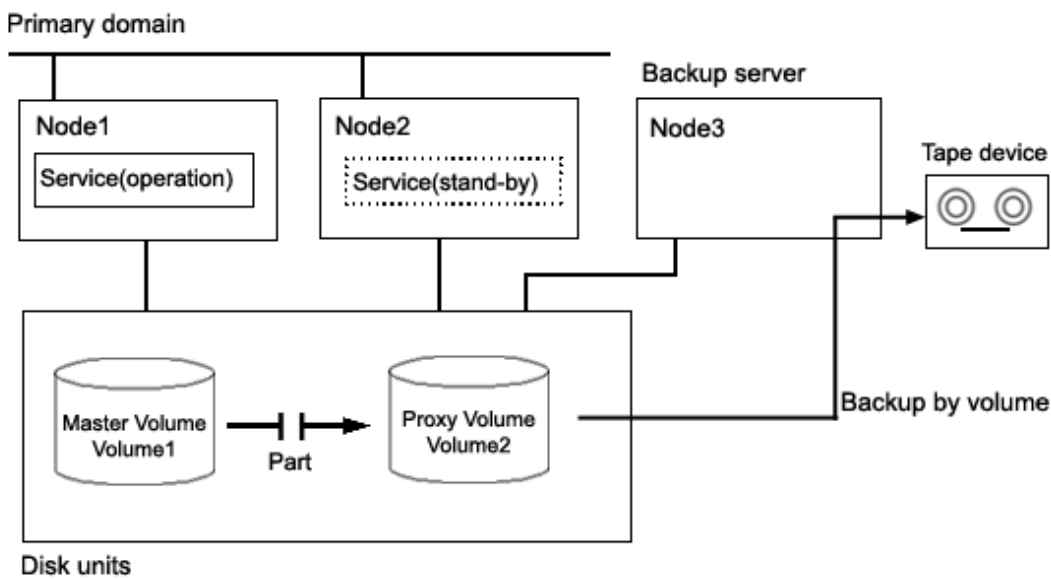


Figure 6.39 Object Configuration When Backing Up





Figure 6.40 Backup Schedule



6.6.3.3 Summary of Restore from a Proxy Volume

If master volume data is damaged while a proxy volume is parted from the master volume used for the services operation, data can be restored from the proxy volume back to the master volume.

When restoring, access to the volume must be suspended temporarily.

Figure 6.41 Restore from a Proxy Volume

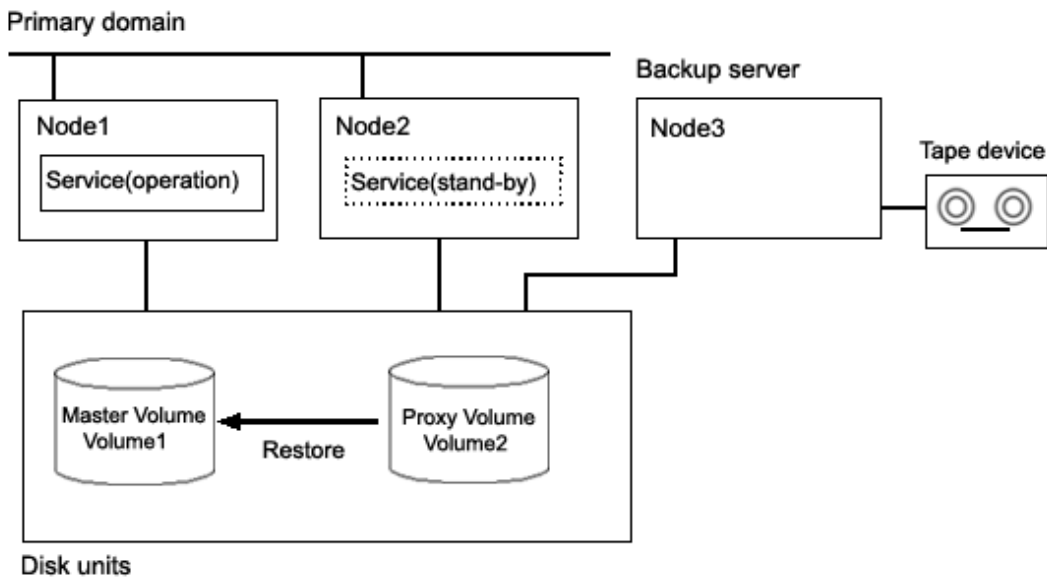


Figure 6.42 Object Configuration When Restoring from a Proxy Volume


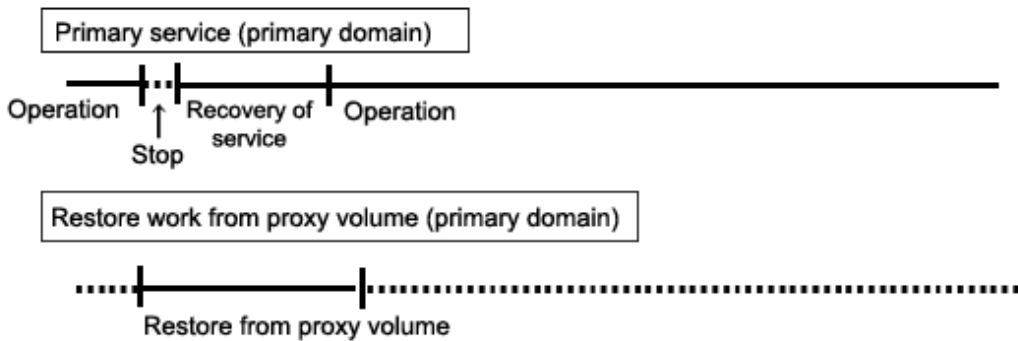


Figure 6.43 Schedule for Restore a Proxy Volume

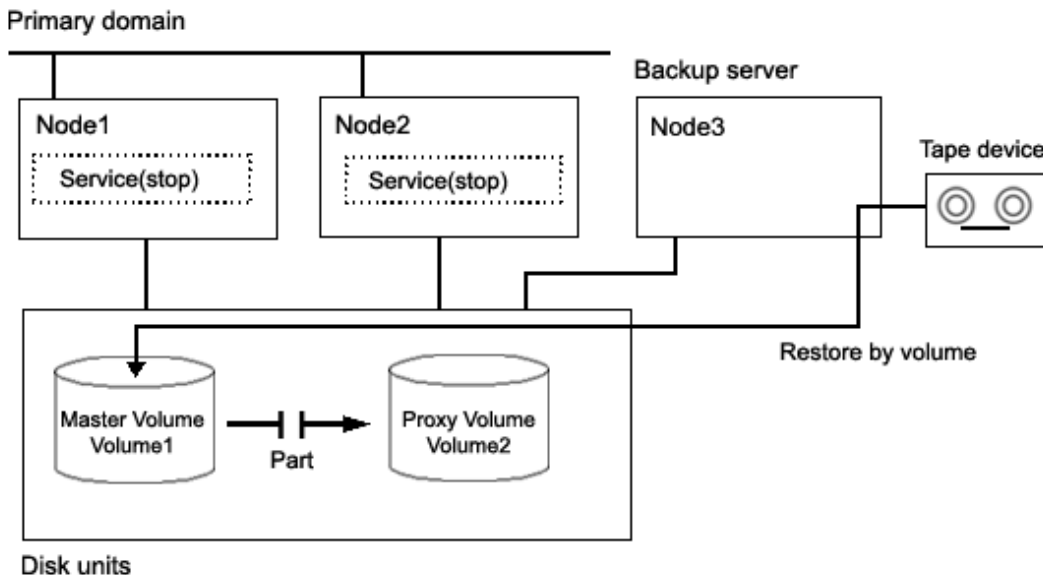


6.6.3.4 Summary of Restore from Tape

If master volume data is damaged while the master volume for the services operation and the proxy volume are in the joined state, the proxy data is also damaged. In this case data can be restored from tape back to the master volume.

Data can be restored while the service is stopped and the master volume is not in use.

Figure 6.44 Restore from Tape



Information

This sub-section shows an example when access can be gained from backup server Node3 to all of the disks constituting master volume Volume1.

Information

When access can be gained from backup server Node3 to all of the disks constituting master volume Volume1 and the disk unit's copy function is used to synchronize a master and a proxy, parting proxy volume Volume2 is not required to restore data from tape.

Information

When access cannot be gained from the backup server to the disks constituting the master volume, while proxy volume Volume2 is parted, copy data from tape to the proxy volume, and then restore master volume data using the proxy volume.

Figure 6.45 Object Configuration When Restoring from Tape


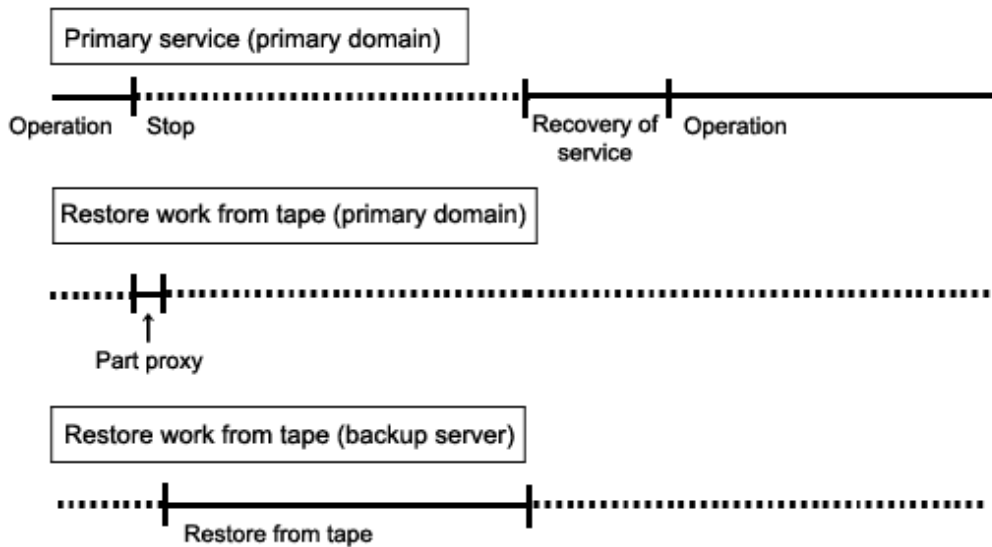


Figure 6.46 Schedule for Restore from Tape



6.6.3.5 Summary of Procedure

Figure 6.47 Outline of the Configuration Procedure




Figure 6.48 Outline of the Backup Procedure

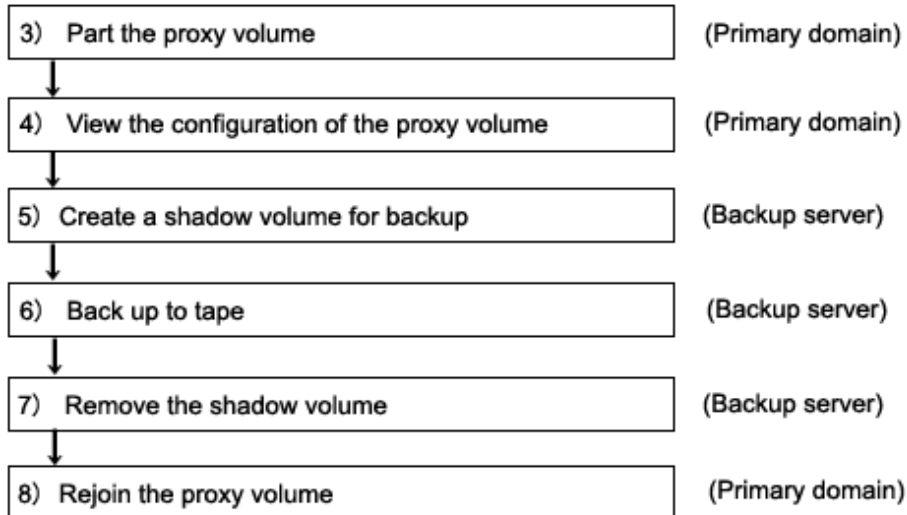


Figure 6.49 Outline of the Procedure for Restoring from a Proxy Volume





Figure 6.50 Outline of the Procedure for Restoring from Tape



6.6.3.6 Configuring an Environment



Note

Automatic Resource Registration

If the backup server resides in a cluster domain (called a backup domain), those disks that are registered as resources in the primary domain or are to be registered with a shadow class in the backup domain may not be involved in the resource registration in the backup domain. In the backup domain, those relevant disks must be described in the Excluded Device List prior to executing the automatic resource registration. For details on the automatic resource registration, see "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide."

1) Creating a master volume

In the primary domain, create the master volume that is used for the services operation.

The following example creates mirror group Group1 that consists of disks c1t1d1 and c1t1d2 to shared class Class1 that is shared on nodes Node1 and Node2 and creates mirror volume Volume1.

```

# sdxdisk -M -c Class1 -a type=shared,scope=Node1:Node2 -d
c1t1d1=Disk1,c1t1d2=Disk2
# sdxdisk -C -c Class1 -g Group1 -d Disk1,Disk2
# sdxvolume -M -c Class1 -g Group1 -v Volume1 -s 1048576
  
```

2) Creating and joining a proxy group

Create a proxy volume as the copy destination of the master volume and join it to the master volume in the primary domain. The following settings are necessary on Node1 or Node2 in the primary domain.

2-1) Creating a proxy volume

Create a proxy volume in the same size as master volume Volume1 to shared class Class1 to which Volume1 belongs.

The following example creates mirror group Group2 that consists of only disk c1t1d3 and creates mirror volume Volume2.

```
# sdxdisk -M -c Class1 -d c1t1d3=Disk3
# sdxdisk -C -c Class1 -g Group2 -d Disk3
# sdxvolume -M -c Class1 -g Group2 -v Volume2 -s 1048576
```

2-2) Stopping the proxy volume

Stop proxy volume Volume2 on all nodes.

```
# sdxvolume -F -c Class1 -v Volume2 -e allnodes
```

2-3) Joining the proxy volume

Relate and join proxy volume Volume2 to master volume Volume1.

```
# sdxproxy Join -c Class1 -m Volume1 -p Volume2
```

After returning from the command, synchronization copying from Volume1 to Volume2 is executed, and as a result, they become synchronized.

6.6.3.7 Backing Up

3) Parting the proxy volume

Part the proxy volume from the master volume. The following procedure must be performed on Node1 or Node2 in the primary domain.



Information

The following example secures data consistency by stopping the services when the proxy volume is parted. Steps 3-2) and 3-4) are not required if software, such as a file system and a database system, that manages volume data provides functionality ensuring data consistency or repairing consistency for a parted volume is present. Alternatively, data consistency must be secured with the method specific to that software. For details, see "[A.2.22 Ensuring Consistency of Snapshot Data.](#)"

3-1) Viewing the status of the proxy volume

Confirm that master volume Volume1 and proxy volume Volume2 are in sync with each other.

Confirm that proxy volume Volume2 is in the joined state. If Join is displayed in the PROXY field, the proxy volume is in the joined state.

```
# sdxinfo -v -c Class1 -o Volume2 -e long
OBJ      NAME      TYPE      CLASS     GROUP     DISK     MASTER  PROXY  ...
-----
volume *      mirror  Class1   Group2    *         *        *       *     ...
volume Volume2 mirror  Class1   Group2    *         Volume1 Join  ...
volume *      mirror  Class1   Group2    *         *        *       *     ...
```

Confirm that data in all the slices of proxy volume Volume2 is valid (STOP).

```
# sdxinfo -s -c Class1 -o Volume2
OBJ      CLASS     GROUP     DISK      VOLUME     STATUS
-----
slice   Class1   Group2    Disk3     Volume2    STOP
```

If data is not valid (STOP), repair the slice status referring to "[F.1.1 Slice Status Abnormality.](#)"

3-2) Stopping the services

To secure consistency of data in proxy volume Volume2 parted, exit all applications using master volume Volume1 on Node1 and Node2.

When Volume1 is used as a file system, it should be unmounted.

3-3) Parting the proxy volume

Part proxy volume Volume2 from master volume Volume1.

```
# sdxproxy Part -c Class1 -p Volume2 -a pjrm=off
```



Just Resynchronization Mode for Proxy

On backup server Node3, data may be written from Node3 into Volume2 when data in Volume2 is backed up to tape. GDS in the primary domain cannot recognize the write occurrence from Node3. Consequently, if the JRM mode for proxies of the parted volume is on, the portions updated from Node3 may not be involved in resynchronization copying performed when the proxy volume is rejoined or restored. If this happens, synchronization of master volume Volume1 and proxy volume Volume2 is no longer ensured. For this reason, the JRM mode of a parted proxy volume must be set to off in advance.

When synchronization copying between a master and a proxy is conducted by the disk unit's copy function, the disk unit's copy function recognizes such a written occurrence from Node3. In this case, only the difference between the master and the proxy is copied through synchronization copying with the disk unit's copy function when the proxy is rejoined regardless of the value specified to the JRM mode for proxies. However, synchronization copying on restore is conducted by the soft copy function. Therefore, the JRM mode of a parted proxy volume should be set to off in advance.

3-4) Resuming the services

When the file system was unmounted in step 3-2), mount it again.

Resume the application stopped in step 3-2).

3-5) Stopping the proxy volume

To prevent improper access to proxy volume Volume2, stop Volume2.

```
# sdxvolume -F -c Class1 -v Volume2 -e allnodes
```

4) Viewing the configuration of the proxy volume

On Node1 or Node2 in the primary domain, view the configuration of proxy volume Volume2 that is the backup target.

```
# sdxinfo -c Class1 -o Volume2
# sdxinfo -c Class1 -o Volume2 -e long
```

5) Creating a shadow volume for backup

On backup server Node3, create a volume for backup (shadow volume).

```
# sdxshadowdisk -M -c Class2 -d c1t1d3=Disk3
# sdxshadowdisk -C -c Class2 -g Group2 -d Disk3
# sdxshadowvolume -M -c Class2 -g Group2 -v Volume2 -s 1048576
```



Master volume data may be damaged if data is written into a shadow volume in incorrect configuration. Be sure to confirm that the shadow volume configuration is correct using the sdxinfo command.

Point

- The shadow volume must be created in the same configuration as the proxy volume created in step 2).
- The shadow disk name must correspond to the disk name assigned in the primary domain. The disk names assigned in the primary domain can be viewed in the NAME field for disk information displayed with the `sdxinfo` command in step 4).
- The class, the group, and the volume can be assigned any name.
- The order of connecting shadow disks to a shadow group must correspond to the order of connecting disks to a group in the primary domain. The order of connecting disks in the primary domain can be viewed in the DISKS field for group information displayed with the `sdxinfo` command in step 4).
- The stripe width of a shadow group must correspond to the stripe width in the primary domain. The stripe width specified in the primary domain can be viewed in the WIDTH field for group information displayed with the `sdxinfo -e long` command in step 4).
- The shadow volume must be created in the size corresponding to the proxy volume size. The proxy volume size can be viewed in the BLOCKS field for volume information displayed with the `sdxinfo` command in step 4).
- If there are multiple volumes, the corresponding shadow volumes must be created in the order of ascending values (first block numbers) in the 1STBLK field for volume information displayed with the `sdxinfo` command in step 4).

6) Backing up to tape

On backup server Node3, back up data in the shadow volume to tape. In the following examples, back up data in shadow volume Volume2 to a tape medium of tape device `/dev/rmt/0`.

See

For details on the backup method, see the manuals of file systems to be backed up and used commands.

Information

In a GFS Shared File System

Back up through the method as described in step 6a).

6a) When backing up data held in a raw device with the `dd(1M)` command

```
# dd if=/dev/sfdsk/Class2/rdsk/Volume2 of=/dev/rmt/0 bs=32768
```

6b) When backing up a ufs file system with the `tar(1)` command

6b-1) Activate shadow volume Volume2 in the read and write access mode (rw).

```
# sdxshadowvolume -F -c Class2 -v Volume2
# sdxshadowvolume -N -c Class2 -v Volume2 -e mode=rw
```

6b-2) Check and repair consistency of the ufs file system on shadow volume Volume2.

If the file system was unmounted when the proxy volume was parted in step 3), this step can be skipped.

```
# fsck -F ufs /dev/sfdsk/Class2/rdsk/Volume2
```

6b-3) Mount the ufs file system on shadow volume Volume2 on `/mnt1`, a temporary mount point, in the read only mode.

```
# mkdir /mnt1
# mount -F ufs -o ro /dev/sfdsk/Class2/dsk/Volume2 /mnt1
```

6b-4) Back up data held in the file system to tape.

```
# cd /mnt1
# tar cvf /dev/rmt/0 .
```

6b-5) Unmount the file system mounted in step 6b-3).

```
# cd /
# umount /mnt1
# rmdir /mnt1
```

6c) When backing up a ufs file system with the ufsdump(1M) command

6c-1) Activate shadow volume Volume2 in the read and write access mode (rw).

```
# sdxshadowvolume -F -c Class2 -v Volume2
# sdxshadowvolume -N -c Class2 -v Volume2 -e mode=rw
```

6c-2) Check and repair consistency of the ufs file system on shadow volume Volume2.

If the file system was unmounted when the proxy volume was parted in step 3), this step can be skipped.

```
# fsck -F ufs -y /dev/sfdsk/Class2/rdsk/Volume2
```

6c-3) Back up data held in the file system to tape.

```
# ufsdump 0ucf /dev/rmt/0 /dev/sfdsk/Class2/rdsk/Volume2
```

6d) When backing up the sfxfs file system (GFS Local File System or SynfinityFile file system) with the tar(1) command

6d-1) Activate shadow volume Volume2 in the read and write access mode (rw).

```
# sdxshadowvolume -F -c Class2 -v Volume2
# sdxshadowvolume -N -c Class2 -v Volume2 -e mode=rw
```

6d-2) Reconfigure the partition information of the sfxfs file system on shadow volume Volume2.

```
# sfxadm /dev/sfdsk/Class2/rdsk/Volume2
```

6d-3) Check and repair consistency of the sfxfs file system on shadow volume Volume2.

If the file system was unmounted when the proxy volume was parted in step 3), this step can be skipped.

```
# fsck -F sfxfs -y /dev/sfdsk/Class2/rdsk/Volume2
```

6d-4) Mount the sfxfs file system on shadow volume Volume2 on /mnt1, a temporary mount point, in the read only mode.

```
# mkdir /mnt1
# mount -F sfxfs -o ro /dev/sfdsk/Class2/dsk/Volume2 /mnt1
```

6d-5) Back up data held in the file system to tape.

```
# cd /mnt1
# tar cvf /dev/rmt/0 .
```

6d-6) Unmount the file system mounted in step 6d-4).

```
# cd /
# umount /mnt1
# rmdir /mnt1
```

6e) When backing up the sfxfs file system (GFS Local File System or SynfinityFile file system) with the sfxdump(1M) command

6e-1) Activate shadow volume Volume2 in the read and write access mode (rw).

```
# sdxshadowvolume -F -c Class2 -v Volume2
# sdxshadowvolume -N -c Class2 -v Volume2 -e mode=rw
```

6e-2) Reconfigure the partition information of the sfxfs file system on shadow volume Volume2.

```
# sfxadm /dev/sfdsk/Class2/rdsk/Volume2
```

6e-3) Check and repair consistency of the sfxfs file system on shadow volume Volume2.

If the file system was unmounted when the proxy volume was parted in step 3), this step can be skipped.

```
# fsck -F sfxfs -y /dev/sfdsk/Class2/rdsk/Volume2
```

6e-4) Back up data held in the file system to tape.

```
# sfxdump 0ucf /dev/rmt/0 /dev/sfdsk/Class2/rdsk/Volume2
```

7) Removing the shadow volume

After the backup process is complete, remove the shadow volume to prevent improper access to it. The following settings are necessary on backup server Node3.

```
# sdxshadowvolume -F -c Class2 -v Volume2
# sdxshadowvolume -R -c Class2 -v Volume2
# sdxshadowgroup -R -c Class2 -g Group2
# sdxshadowdisk -R -c Class2 -d Disk3
```

8) Rejoining the proxy volume

Rejoin the proxy volume to the master volume. The following procedure must be performed on Node1 or Node2 in the primary domain.

8-1) Rejoining the proxy volume

Rejoin proxy volume Volume2 to master volume Volume1.

```
# sdxproxy Rejoin -c Class1 -p Volume2
```

After returning from the command, synchronization copying from Volume1 to Volume2 is executed.

8-2) Viewing the copy status

The status of synchronization copying can be viewed using the `sdxinfo -S` command. The slice of proxy volume Volume2 as the copy destination is in COPY status if copying is in progress and it will be in STOP status after the copy process ends normally.

```
# sdxinfo -S -c Class1 -o Volume2
OBJ    CLASS   GROUP   DISK    VOLUME  STATUS
-----
slice  Class1   Group2  Disk3   Volume2 STOP
```

6.6.3.8 Restoring from a Proxy Volume

9) Stopping the services

Exit all applications using master volume Volume1 on Node1 and Node2 in the primary domain.

When Volume1 is used as a file system, it should be unmounted.

10) Stopping the master volume

Stop master volume Volume1 on Node1 and Node2 in the primary domain. Execute the following command on Node1 or Node2.

```
# sdxvolume -F -c Class1 -v Volume1 -e allnodes
```

11) Restoring from the proxy volume

In the primary domain, restore data from proxy volume Volume2 back to master volume Volume1. Execute the following command on Node1 or Node2.

```
# sdxproxy RejoinRestore -c Class1 -p Volume2
```

After returning from the command, synchronization copying from Volume2 to Volume1 is executed.

12) Resuming the services

After synchronization copying is started from proxy volume Volume2 to master volume Volume1 in step 11), the services can be resumed before the copy process is completed.

The following procedure must be performed on the node that runs the services.

12-1) Activating the master volume

Activate master volume Volume1.

```
# sdxvolume -N -c Class1 -v Volume1
```

12-2) Reconfiguring the partition information of the sfxfs file system

When master volume Volume1 is used as the sfxfs file system, reconfigure the partition information.

```
# sfxadm /dev/sfdsk/Class1/rdsk/Volume1
```

12-3) Check and repair consistency of the sfxfs file system

When master volume Volume1 is used as the sfxfs file system, check and repair consistency of the file system.

```
# fsck -F sfxfs -y /dev/sfdsk/Class1/rdsk/Volume1
```

12-4) Resuming the services

When the file system on master volume Volume1 was unmounted in step 9), mount it again.

Start the applications using Volume1.

12-5) Viewing the copy status

The status of synchronization copying from proxy volume Volume2 to master volume Volume1 executed in step 11), can be viewed using the `sdxinfo -S` command. The slice of master volume Volume1 as the copy destination is in COPY status if copying is in progress and it will be in ACTIVE status after the copy process ends normally.

```
# sdxinfo -S -c Class1 -o Volume1
OBJ    CLASS  GROUP  DISK   VOLUME  STATUS
-----
slice  Class1  Group1  Disk1  Volume1 ACTIVE
slice  Class1  Group1  Disk2  Volume1 ACTIVE
```

Information

Master volume Volume1 and proxy volume Volume2 are made the joined state through step 11). If data in Volume1 is damaged while they are in the joined state, the proxy data is also damaged. Thus data cannot be restored from Volume2 back to Volume1. Therefore, once the synchronization copying from Volume2 to Volume1 is complete, it is recommended to part Volume2 from Volume1. For details on the procedure for parting a proxy volume, see step 3) described in "6.6.3.7 Backing Up."

6.6.3.9 Restoring from Tape

This sub-section shows an example that access can be gained from backup server Node3 to all of the disks constituting master volume Volume1.

Information

When Using a Disk Unit's Copy Function

When access can be gained from backup server Node3 to all of the disks constituting master volume Volume1 and the disk unit's copy function is used to synchronize a master and a proxy, parting the proxy volume in step 15) is not required.

Information

When access cannot be gained from the backup server to the disks constituting the master volume

While proxy volume Volume2 is parted, copy data from tape to the proxy volume, and then restore master volume data using the proxy volume.

13) Stopping the services

Exit all applications using master volume Volume1 on Node1 and Node2 in the primary domain.

When Volume1 is used as a file system, it should be unmounted.

14) Stopping the master volume

On Node1 and Node2 in the primary domain, stop master volume Volume1 to prevent improper access to it. Execute the following command on Node1 or Node2.

```
# sdxvolume -F -c Class1 -v Volume1 -e allnodes
```

15) Parting the proxy volume

In the primary domain, part proxy volume Volume2 from master volume Volume1. Execute the following command on Node1 or Node2.

```
# sdxproxy Part -c Class1 -p Volume2 -a pjrm=off
```

Note

Just Resynchronization Mode for Proxy

After data held in master volume Volume1 is restored from tape on backup server Node3, proxy volume Volume2 is supposed to be rejoined to master volume Volume1. At this point the entire Volume1 data must be copied to Volume2. For this reason, the JRM mode of a parted proxy volume must be set to off in advance.

Information

When Using a Disk Unit's Copy Function

When the disk unit's copy function is used to synchronize a master and a proxy, this procedure (parting a proxy) is not required.

Information

When access cannot be gained from the backup server to the disks constituting the master volume

After proxy volume Volume2 is parted, inactivate Volume2 on Node1 and Node2 to prevent Volume2 from being written in improperly.

16) Viewing the status and configuration of the master volume

On Node1 and Node2 in the primary domain, view the configuration and status of master volume Volume1 that is the restore target. Confirm that all of the slices constituting Volume1 are in STOP status. If the status of a slice is invalid, repair it referring to "[F.1.1 Slice Status Abnormality](#)."

```
# sdxinfo -S -c Class1 -o Volume1
OBJ    CLASS  GROUP  DISK    VOLUME  STATUS
-----
slice  Class1 Group1  Disk1   Volume1 STOP
slice  Class1 Group1  Disk2   Volume1 STOP
```

Information

When Using a Disk Unit's Copy Function

When the proxy volume was not parted in step 15) because the disk unit's copy function is used for synchronizing a master and a proxy, you must also confirm that all of the slices constituting proxy volume Volume2 are in STOP status.

Information

When access cannot be gained from the backup server to the disks constituting the master volume

View the configuration and the status of proxy volume Volume2 that is the restore target.

17) Creating a shadow volume for restore

On backup server Node3, create a volume for restore (shadow volume).

17-1) Creating a shadow volume

```
# sdxshadowdisk -M -c Class2 -d c1t1d1=Disk1,c1t1d2=Disk2
# sdxshadowdisk -C -c Class2 -g Group1 -d Disk1,Disk2
# sdxshadowvolume -M -c Class2 -g Group1 -v Volume1 -s 1048576
```

Note

Master volume data may be damaged if data is written into a shadow volume in incorrect configuration. Be sure to confirm that the shadow volume configuration is correct using the `sdxinfo` command.

Point

- The shadow volume must be created in the same configuration as the master volume created in step 1).
- The shadow disk name must correspond to the disk name assigned in the primary domain. The disk names assigned in the primary domain can be viewed in the NAME field for disk information displayed with the `sdxinfo` command in step 16).
- The class, the group, and the volume can be assigned any name.
- The order of connecting shadow disks to a shadow group must correspond to the order of connecting disks to a group in the primary domain. The order of connecting disks in the primary domain can be viewed in the DISKS field for group information displayed with the `sdxinfo` command in step 16).
- The stripe width of a stripe type shadow group must correspond to the stripe width of a stripe group in the primary domain. The stripe width in the primary domain can be viewed in the WIDTH field for group information displayed with the `sdxinfo -e long` command in step 16).
- A shadow volume must be created in the size corresponding to the master volume size. The master volume size can be viewed in the BLOCKS field for volume information displayed with the `sdxinfo` command in step 16).
- If there are multiple volumes, the corresponding shadow volumes must be created in the order of ascending values (first block numbers) in the 1STBLK field for volume information displayed with the `sdxinfo` command in step 16).

Information

When access cannot be gained from the backup server to the disks constituting the master volume

Create a shadow volume for restore in a similar procedure that created a proxy volume in step 2).

17-2) Setting the access mode of the shadow volume

Activate shadow volume Volume1 in the read and write access mode (rw).

```
# sdxshadowvolume -F -c Class2 -v Volume1
# sdxshadowvolume -N -c Class2 -v Volume1 -e mode=rw
```

17-3) Viewing the configuration of the shadow volume

Using the `sdxinfo` command, confirm that the group configuration and the volume configuration are correct based on group information in the DISKS field, volume information in the 1STBLK field and in the BLOCKS field and so on.

```
# sdxinfo -c Class2
```

18) Restoring from tape

On backup server Node3, restore shadow volume data from tape to which it was backed up in step 6). In the following examples, restore data held in shadow volume Volume1 from a tape medium of tape device /dev/rmt/0.



For details on the restore method, see the manuals of file systems to be restored and used commands.



In a GFS Shared File System

Restore through the method as described in step 18a).

18a) When restoring data held in a raw device with the dd(1M) command

```
# dd if=/dev/rmt/0 of=/dev/sfdsk/Class2/rdisk/Volume1 bs=32768
```

18b) When restoring a ufs file system with the tar(1) command

18b-1) Create a ufs file system to shadow volume Volume1.

```
# newfs /dev/sfdsk/Class2/rdisk/Volume1
```

18b-2) Mount the ufs file system on shadow volume Volume1 on /mnt1, a temporary mount point.

```
# mkdir /mnt1
# mount -F ufs /dev/sfdsk/Class2/dsk/Volume1 /mnt1
```

18b-3) Restore data held in the file system from tape.

```
# cd /mnt1
# tar xvf /dev/rmt/0
```

18b-4) Unmount the file system mounted in step 18b-3).

```
# cd /
# umount /mnt1
# rmdir /mnt1
```

18c) When restoring a ufs file system with the ufsrestore(1M) command

18c-1) Create a ufs file system to shadow volume Volume1.

```
# newfs /dev/sfdsk/Class2/rdisk/Volume1
```

18c-2) Mount the ufs file system on shadow volume Volume1 on /mnt1, a temporary mount point.

```
# mkdir /mnt1
# mount -F ufs /dev/sfdsk/Class2/dsk/Volume1 /mnt1
```

18c-3) Restore data held in the file system from tape.

```
# cd /mnt1
# ufsrestore rvf /dev/rmt/0
```

18c-4) Delete the temporary file created by the ufsrestore(1M) command.

```
# rm /mnt1/restoresymtable
```

18c-5) Unmount the file system mounted in step 18c-2).

```
# cd /
# umount /mnt1
# rmdir /mnt1
```

18d) When restoring the sfxfs file system (GFS Local File System or SynfinityFile file system) with the tar(1) command

18d-1) Create the sfxfs file system to shadow volume Volume1.

```
# sfxnewfs /dev/sfdsk/Class2/rdsk/Volume1
```

18d-2) Mount the sfxfs file system on shadow volume Volume1 on /mnt1, a temporary mount point.

```
# mkdir /mnt1
# mount -F sfxfs /dev/sfdsk/Class2/dsk/Volume1 /mnt1
```

18d-3) Restore data held in the file system from tape.

```
# cd /mnt1
# tar xvf /dev/rmt/0
```

18d-4) Unmount the file system mounted in step 18d-2).

```
# cd /
# umount /mnt1
# rmdir /mnt1
```

18e) When restoring the sfxfs file system (GFS Local File System or SynfinityFile file system) the sfxrestore(1M) command

18e-1) Create the sfxfs file system to shadow volume Volume1.

```
# sfxnewfs /dev/sfdsk/Class2/rdsk/Volume1
```

18e-2) Mount the sfxfs file system on shadow volume Volume1 on /mnt1, a temporary mount point.

```
# mkdir /mnt1
# mount -F sfxfs /dev/sfdsk/Class2/dsk/Volume1 /mnt1
```

18e-3) Restore data held in the file system from tape.

```
# cd /mnt1
# sfxrestore rvf /dev/rmt/0
```

18e-4) Delete the temporary file created by the sfxrestore(1M) command.

```
# rm /mnt1/restoresymtable
```

18e-5) Unmount the file system mounted in step 18e-2).

```
# cd /
# umount /mnt1
# rmdir /mnt1
```

19) Removing the shadow volume

After the restore process is complete, remove the shadow volume to prevent improper access to it. The following settings are necessary on backup server Node3.

```
# sdxshadowvolume -F -c Class2 -v Volume1
# sdxshadowvolume -R -c Class2 -v Volume1
# sdxshadowgroup -R -c Class2 -g Group1
# sdxshadowdisk -R -c Class2 -d Disk1
# sdxshadowdisk -R -c Class2 -d Disk2
```

20) Resuming the services

Resume services in the primary domain. The following procedure must be performed on the node that runs the services.



Information

.....

When access cannot be gained from the backup server to the disks constituting the master volume

Before the services are resumed, restore data from proxy volume Volume2 to master volume Volume1. For the procedure see "[6.6.3.8 Restoring from a Proxy Volume.](#)"

.....

20-1) Activating the master volume

Activate master volume Volume1.

```
# sdxvolume -N -c Class1 -v Volume1
```

20-2) Reconfiguring the partition information of the sfxfs file system

When master volume Volume1 is used as the sfxfs file system, reconfigure the partition information.

```
# sfxadm /dev/sfdsk/Class1/rdsk/Volume1
```

20-3) Resuming the services

When the file system on Volume1 was unmounted in step 13), mount it again.

Start the applications using Volume1.

6.6.4 Backing Up and Restoring by the Disk Unit's Copy Function

This sub-section describes the method of backing up data from and restoring data back to mirror volumes in the primary domain through a backup server in another domain using the disk unit's copy function provided by a disk array unit.

If volume data is updated with copy functions of disk units, GDS does not recognize the update. If mirrored disk data is updated with copy functions of disk units, synchronization of the mirrors are no longer ensured. Therefore, when restoring mirror volume data using the disk unit's copy function, it is necessary to disconnect the other disk from mirroring once and reconnect it to mirroring after restoration.

The following example illustrates using an EMC's Symmetrix storage unit as a disk array unit and EMC TimeFinder as a copy function.

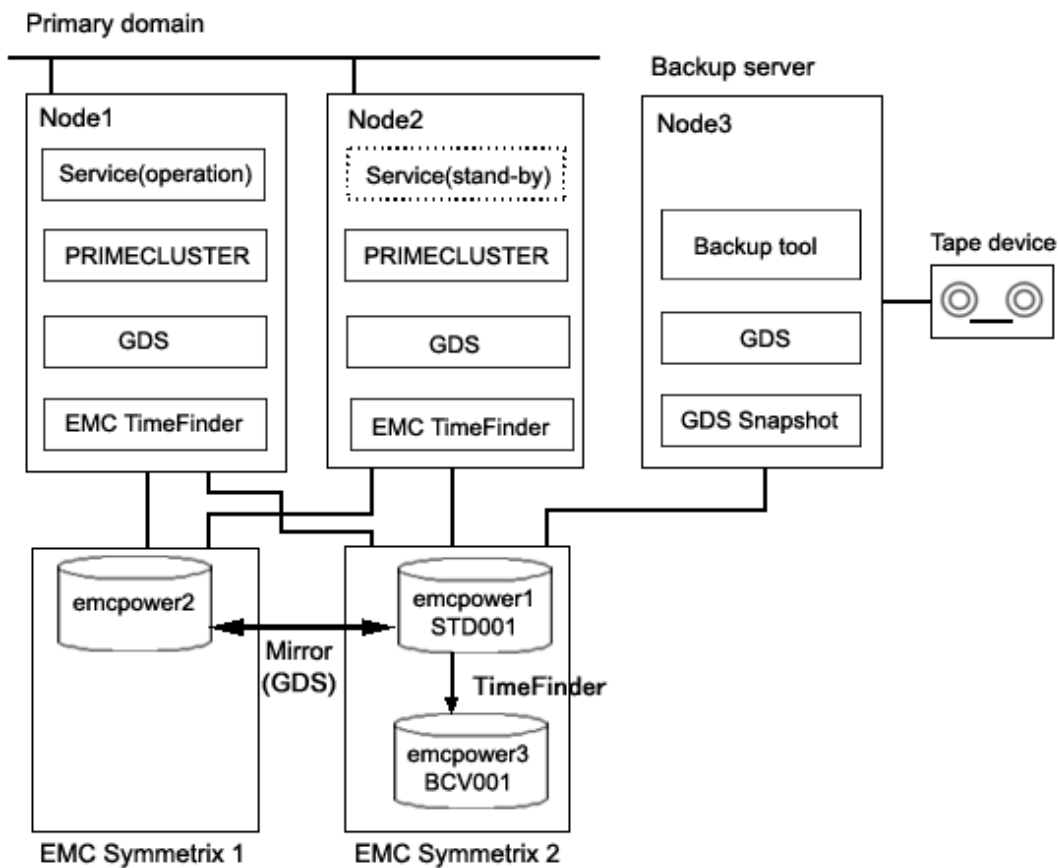
When restoration is performed with TimeFinder, configuration information within the private slice is also restored. For this reason, simultaneous backup and restore for all disks within the class is required. Additionally, the object configuration and status must match at backup and at restore, and so as in restore, it is necessary to disconnect the disks from mirroring before backup and reconnect after.

 See

- For using EMC's Symmetrix storage units, refer to notes described in "A.2.20 To Use EMC Symmetrix."
- For details on the EMC TimeFinder command, refer to the EMC's SYMCLI manual.

6.6.4.1 System Configuration

Figure 6.51 System Configuration



 Information

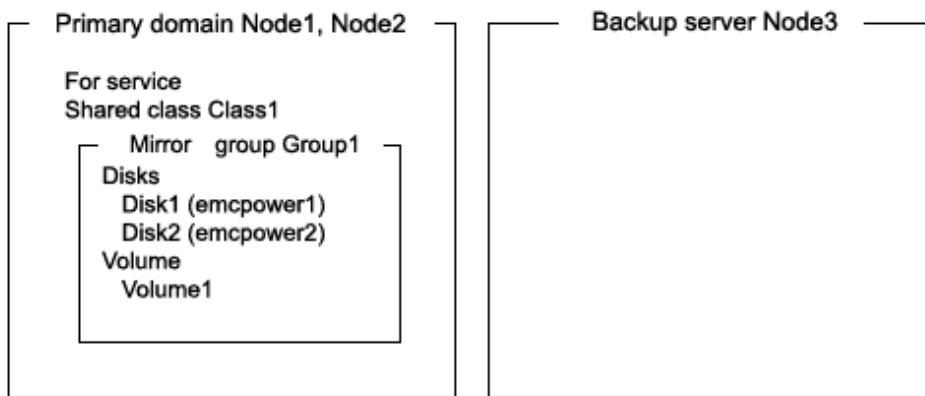
- In the primary domain, SynfinityDisk and SynfinityCluster are also available instead of GDS and PRIMECLUSTER.
- A configuration that makes a node in the primary domain (e.g. Node2) work as a backup server is also available.

 Note

Physical Device Name

Different physical device names (such as emcpower1) may be assigned to the identical physical disk in the primary domain and the backup server.

Figure 6.52 Object Configuration in Normal Operation



6.6.4.2 Summary of Backup

Data in a split BCV can be backed up to tape during the services.

To secure consistency of BCV data, the services should be stopped temporarily when the BCV is split.

Information

Consistency of Snapshot Data

When detaching BCV while the services are operating, data consistency must be secured through the method specific to that software, such as a file system and a database system, which manages volume data. For details, see "[A.2.22 Ensuring Consistency of Snapshot Data.](#)"

Figure 6.53 Backup

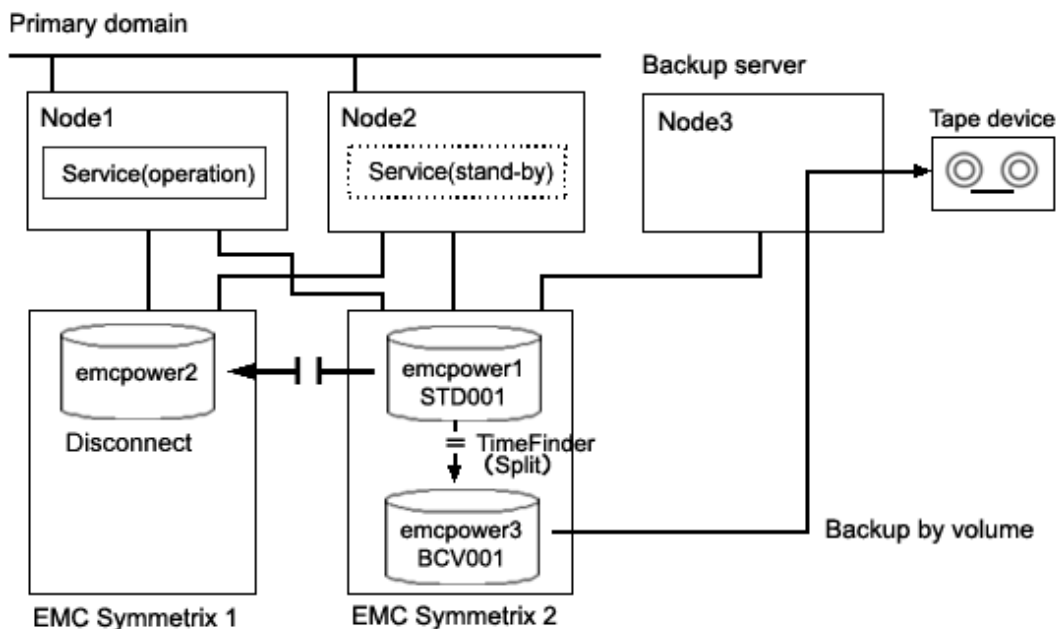


Figure 6.54 Object Configuration When Backing Up


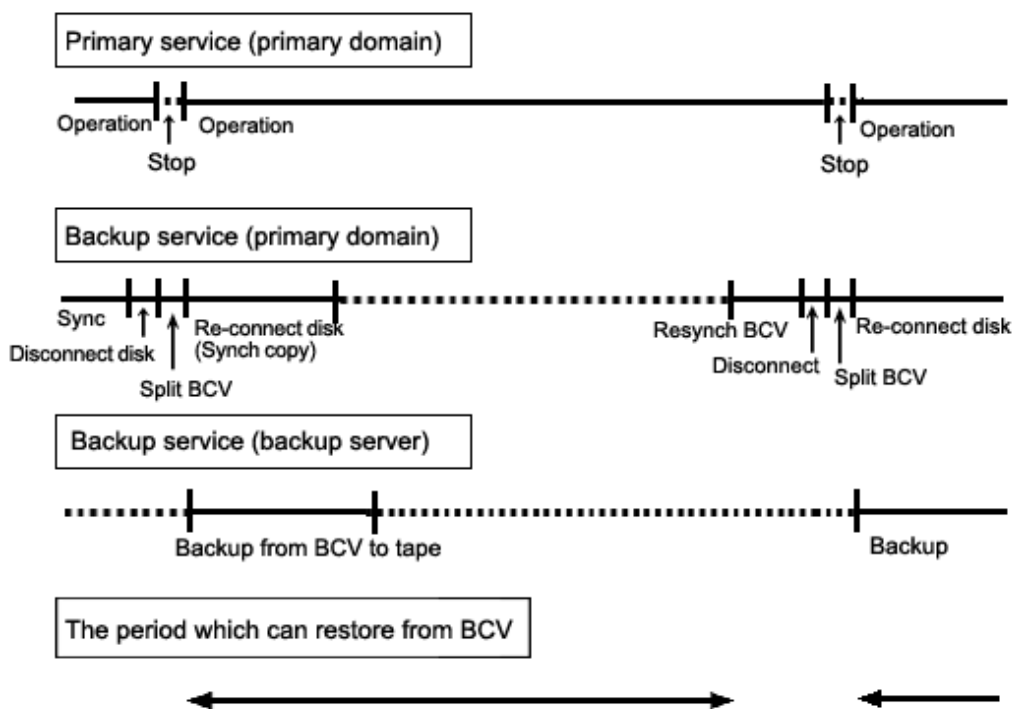


Figure 6.55 Backup Schedule



Note

When Restoring from the BCV

The following conditions must be met when performing a backup to the BCV (disconnecting the BCV).

- Back up all disks registered with the backup target class (except for disks that are disconnected and in SWAP status) to the BCV. For disks not backed up to the BCV, it is necessary to disconnect from the class before the BCV is detached.
- Before backup to the BCV is completed for all disks within the class (except for disks that are disconnected and in SWAP status), do not change the object configuration or status in the class.

Information

When Not Restoring from the BCV

When restore is always from tape instead of BCV, it is not necessary to perform disk disconnection and reconnection in the backup process.

6.6.4.3 Summary of Restore from a BCV

If data in a standard device is damaged while a BCV is split from the disk (standard disk) used for the services operation, the standard device data can be restored from the BCV.

Data can be restored while the service is stopped and the application volume is not in use.

Note

Conditions for Restore from the BCV

The following conditions must be met when performing a restore from the BCV.

- To all disks within the restore target class (except for disks that are disconnected and in SWAP status), restore data from the BCV. For disks not backed up to the BCV, it is necessary to disconnect from the class before restore from the BCV is performed.
- The configuration of objects within the restore target class must be the same as it was when the backup to the BCV (disconnection of the BCV) was performed.
- Before restore to the BCV is completed for all disks within the class (except for disks that are disconnected and in SWAP status), do not reboot any node in the primary domain.

Figure 6.56 Restore from a BCV

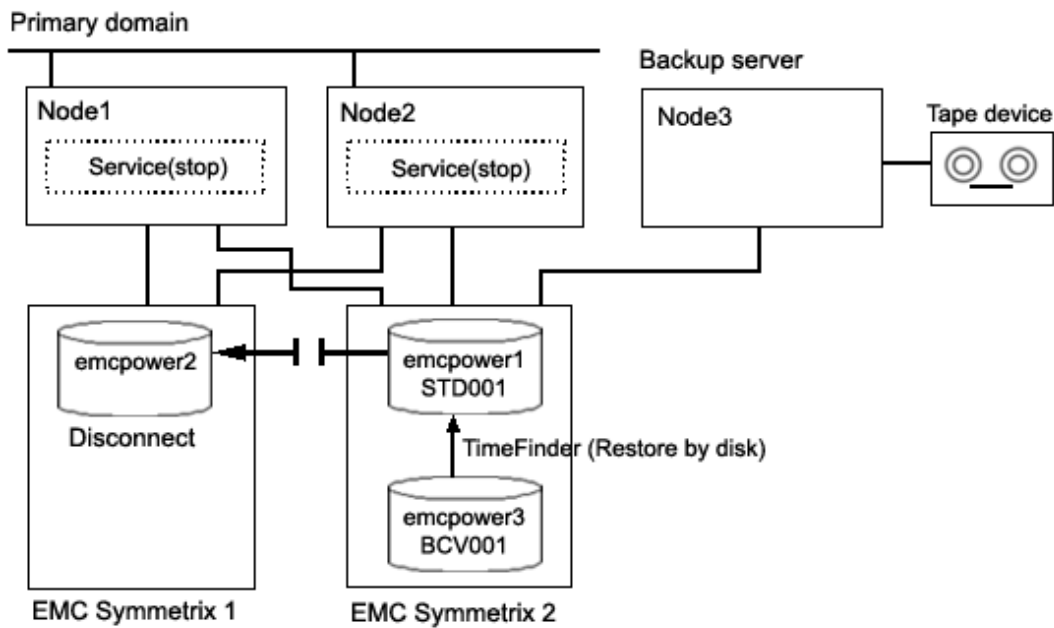


Figure 6.57 Object Configuration When Backing Up from a BCV


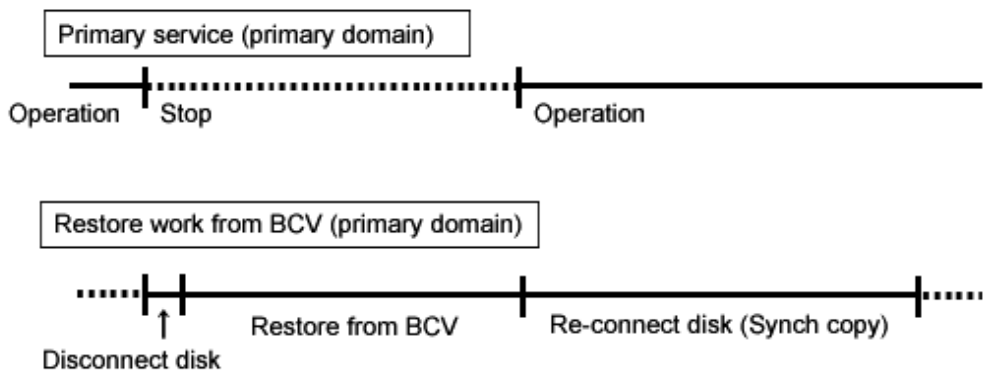


Figure 6.58 Schedule for Restore from a BCV




6.6.4.4 Summary of Restore from Tape

If data in a standard device used for the service is damaged while it is in sync with a BCV, data in the BCV is also damaged. In this case data can be restored from tape back to the standard device.

Data can be restored while the services are stopped and the application volume is not in use.

Figure 6.59 Restore from Tape



 **Information**

In this configuration, access cannot be gained from backup server Node3 to disk emcpower2. Therefore, after data held in emcpower1 is restored from tape while emcpower2 is detached temporarily, synchronization copying from emcpower1 to emcpower2 must be performed by reattaching emcpower2. When access can be gained from Node3 to both emcpower1 and emcpower2, it is not required that emcpower2 be detached temporarily since data can be restored from tape back to both emcpower1 and emcpower2.

Figure 6.60 Object Configuration When Restoring from Tape

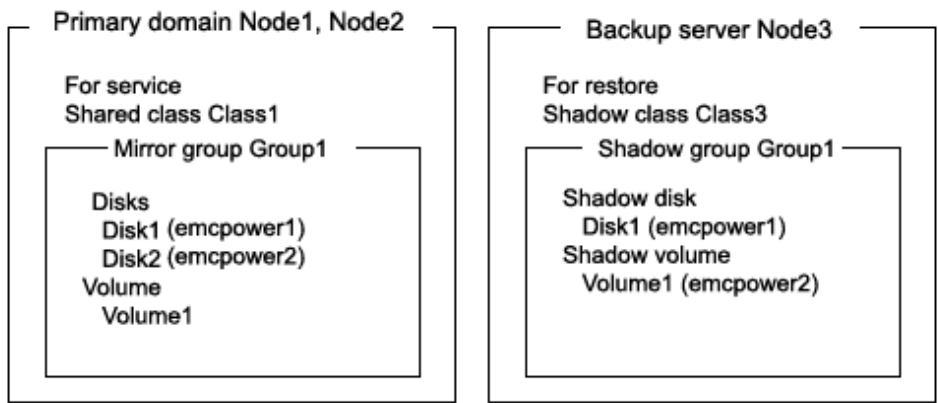



Figure 6.61 Schedule for Restore from Tape



6.6.4.5 Summary of Procedure

Figure 6.62 Outline of the Configuration Procedure




Figure 6.63 Outline of the Backup Procedure




Figure 6.64 Outline of the Procedure for Restoring from a BCV


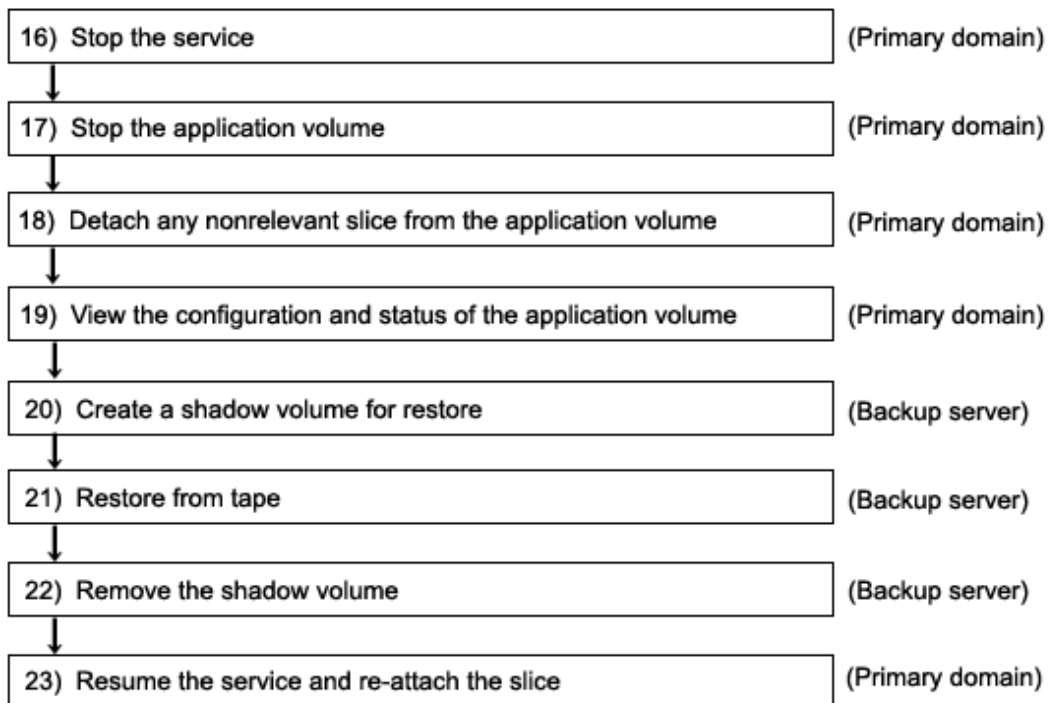


Figure 6.65 Outline of the Procedure for Restoring from Tape



6.6.4.6 Configuring an Environment



Note

Automatic Resource Registration

If the backup server resides in a cluster domain (called a backup domain), those disks that are registered as resources in the primary domain or are to be registered with a shadow class in the backup domain may not be involved in the resource registration in the backup domain. In the backup domain, those relevant disks must be described in the Excluded Device List prior to executing the automatic resource registration. For details on the automatic resource registration, see "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide."

1) Creating an application volume

Create a mirror volume used for the services operation on disks (standard devices) emcpower1 and emcpower2. The following settings are necessary on Node1 or Node2 in the primary domain.

1-1) Registering disks

Register disks (standard devices) emcpower1 and emcpower2 with shared class Class1 that is shared on Node1 and Node2, and name them Disk1 and Disk2 respectively.

```
# sdxdisk -M -c Class1 -a type=shared,scope=Node1:Node2 -d  
emcpower1=Disk1,emcpower2=Disk2
```

1-2) Creating a mirror group

Connect disks Disk1 and Disk2 to mirror group Group1.

```
# sdxdisk -C -c Class1 -g Group1 -d Disk1,Disk2
```

1-3) Creating a mirror volume

Create mirror volume Volume1 to mirror group Group1.

```
# sdxvolume -M -c Class1 -g Group1 -v Volume1 -s 1048576
```

2) Synchronizing a BCV

Relate standard device emcpower1 to BCV device emcpower3 that will be the copy destination. The following settings are necessary on both nodes Node1 and Node2 in the primary domain.

2-1) Creating a device group

Create device group DevGroup.

```
# symdg create DevGroup
```

2-2) Registering a standard device

Register standard device emcpower1 with device group DevGroup, and name it STD001 as a logical device.

```
# symld -g DevGroup add pd /dev/rdisk/emcpower1 STD001
```

2-3) Relating a BCV device

Relate BCV device emcpower3 to device group DevGroup, and name it BCV001 as a logical device.

```
# symbcv -g DevGroup associate pd /dev/rdisk/emcpower3 BCV001
```

2-4) Establishing a BCV pair (synchronized)

Synchronize standard device STD001 with BCV device BCV001.

```
# symmir -g DevGroup -full establish STD001 bcv ld BCV001
```

6.6.4.7 Backing Up

3) Disconnecting a disk of the application volume

In the primary domain, among disks registered with class Class1 to which application volume Volume1 belongs, disconnect a disk (Disk2) other than the backup target disk Disk1 from Class1. Execute the following command on node Node1 or Node2 in the primary domain.

```
# sdxswap -O -c Class1 -d Disk2
```

4) Splitting the BCV

Split BCV device emcpower3 from standard device emcpower1. The following procedure must be performed on Node1 or Node2 in the primary domain.

Information

The following example secures data consistency by stopping the services when a BCV is split. Steps 4-3) and 4-5) are not required if your software, such as a file system and a database system, that manages volume data provides functionality ensuring data consistency or repairing consistency for a split BCV. Alternatively, data consistency must be secured with the method specific to that software. For details, see "[A.2.22 Ensuring Consistency of Snapshot Data.](#)"

4-1) Viewing the status of the application volume

Check the slice on standard device emcpower1 (Disk1) that is the copy source of BCV device emcpower3 among the slices of application volume Volume1 for the data validity (ACTIVE or STOP). Additionally, check that the slice of Disk2 disconnected in step 3) is in SWAP status.

```
# sdxinfo -S -c Class1 -o Volume1
OBJ    CLASS  GROUP  DISK    VOLUME  STATUS
-----
slice  Class1  Group1 Disk1    Volume1 ACTIVE
slice  Class1  Group1 Disk2    Volume1 NOUSE
```

If the data is invalid (not ACTIVE or STOP), repair the slice status referring to "[F.1.1 Slice Status Abnormality.](#)"

4-2) Viewing the condition of the BCV pair

Confirm that standard device STD001 (emcpower1) and BCV device BCV001 (emcpower3) are in sync with each other (synchronized).

```
# symmir -g DevGroup query
Device Group (DG) Name: DevGroup
DG's Type           : REGULAR
DG's Symmetrix ID   : 000285502123

Standard Device          BCV Device          status
-----
Inv.                      Inv.
Logical   Sym Tracks   Logical           Sym Tracks   STD <=> BCV
-----
STD001    005          0 BCV001          073 * 61754   Synchronized
```

4-3) Stopping the services

To secure consistency of data in the split BCV device, exit all applications using application volume Volume1 on Node1 and Node2.

When Volume1 is used as a file system, it should be unmounted.

4-4) Splitting the BCV pair (disconnect)

Split the BCV pair (standard device STD001 and BCV device BCV001).

```
# symmir -g DevGroup split
```

4-5) Resuming the services

When the file system was unmounted in step 4-3), mount it again.

Resume the application stopped in step 4-3).

5) Reconnecting the disk of the application volume

Reconnect disk Disk2 disconnected in step 3) from class Class1 to which application volume Volume1 belongs to Class1.

```
# sdxswap -I -c Class1 -d Disk2 -e nowaitsync
```

After returning from the command, synchronization copying from the slice on Disk1 of volume Volume1 to the slice on Disk2 is executed.

6) Viewing the configuration of the application volume

On Node1 or Node2 in the primary domain, view the configuration of services volume Volume1 that is the backup target.

```
# sdxinfo -c Class1
OBJ      NAME      TYPE      SCOPE      SPARE
-----
class    Class1    shared    Node1:Node2 0

OBJ      NAME      TYPE      CLASS      GROUP      DEVDNAM      DEVBLKS      DEVCONNECT      STATUS
-----
disk     Disk1    mirror    Class1     Group1     emcpower1    8380800      Node1:Node2     ENABLE
disk     Disk2     mirror    Class1     Group1     emcpower2    8380800      Node1:Node2     ENABLE

OBJ      NAME      CLASS      DISKS              BLKS      FREEBLKS      SPARE
-----
group    Group1    Class1     Disk1:Disk2        8290304    7176192      0

OBJ      NAME      CLASS      GROUP      SKIP  JRM      1STBLK      LASTBLK      BLOCKS      STATUS
-----
volume  *         Class1     Group1     *     *         0           65535         65536      PRIVATE
volume  Volume1   Class1     Group1     off   on        65536       1114111       1048576    ACTIVE
volume  *         Class1     Group1     *     *         1114112     8290303       7176192    FREE

OBJ      CLASS      GROUP      DISK      VOLUME      STATUS
-----
slice    Class1     Group1     Disk1     Volume1     ACTIVE
slice    Class1     Group1     Disk2     Volume1     ACTIVE
```

7) Creating a shadow volume for backup

On backup server Node3, create a volume for backup (shadow volume) on BCV device emcpower3. The following settings are necessary on backup server Node3.



Application volume data may be damaged if data is written into a shadow volume in incorrect configuration. Be sure to confirm that the shadow volume configuration is correct in step 7-4).

7-1) Registering a shadow disk

Register disk (BCV device) emcpower3 with shadow class Class2, and name it Disk1.

```
# sdxshadowdisk -M -c Class2 -d emcpower3=Disk1
```

Point

- The disk name must correspond to the disk name assigned in step 1-1) to standard device emcpower1 that is the copy source of BCV device emcpower3. The disk names assigned in 1-1) can be viewed in the NAME field for disk information displayed with the sdxinfo command in step 6).
- The class can be assigned any name. However, if Node3 resides in the same domain as Node1 and Node2, it must be assigned a name different from the name of a class created in step 1-1).

7-2) Creating a shadow group

Connect shadow disk Disk1 to mirror type shadow group Group1.

```
# sdxshadowdisk -C -c Class2 -g Group1 -d Disk1
```

7-3) Creating a shadow volume

Create a shadow volume Volume1 to shadow group Group1.

```
# sdxshadowvolume -M -c Class2 -g Group1 -v Volume1 -s 1048576
```

Point

- The volume must be created in the size corresponding to the volume size in step 1-3). The size of a volume created in step 1-3) can be viewed in the BLOCKS field for volume information displayed with the sdxinfo command in step 6).
- If there are multiple volumes, the corresponding shadow volumes must be created in the order of ascending values (first block numbers) in the 1STBLK field for volume information displayed with the sdxinfo command in step 6).
- The volume can be assigned any name.

7-4) Viewing the configuration of the shadow volume

Using the sdxinfo command, confirm that the group configuration and the volume configuration are correct based on group information in the DISKS field, volume information in the 1STBLK field and in the BLOCKS field and so on.

```
# sdxinfo -c Class2
OBJ  NAME  TYPE  SCOPE  SPARE
-----
class Class2 local Node3 0

OBJ  NAME  TYPE  CLASS  GROUP  DEVNAM  DEVBLKS  DEVCONNECT  STATUS
-----
disk  Disk1 mirror Class2  Group1  emcpower3  8380800  Node3  ENABLE

OBJ  NAME  CLASS  DISKS  BLKS  FREEBLKS  SPARE
-----
group Group1 Class2 Disk1 8290304 7176192 0

OBJ  NAME  CLASS  GROUP  SKIP  JRM  1STBLK  LASTBLK  BLOCKS  STATUS
-----
```


volume	*	Class2	Group1	*	*	0	65535	65536	PRIVATE
volume	Volume1	Class2	Group1	off	off	65536	1114111	1048576	ACTIVE
volume	*	Class2	Group1	*	*	1114112	8290303	7176192	FREE
OBJ	CLASS	GROUP	DISK	VOLUME	STATUS				
-----	-----	-----	-----	-----	-----				
slice	Class2	Group1	Disk1	Volume1	ACTIVE				

8) Backing up to tape

On backup server Node3, back up data in the shadow volume to tape. In the following examples, back up data in shadow volume Volume1 to a tape medium of tape device /dev/rmt/0.



See

For details on the backup method, see the manuals of file systems to be backed up and used commands.



Information

In a GFS Shared File System

Back up through the method as described in step 8a).

8a) When backing up data held in a raw device with the dd(1M) command

```
# dd if=/dev/sfdsk/Class2/rdsk/Volume1 of=/dev/rmt/0 bs=32768
```

8b) When backing up a ufs file system with the tar(1) command

8b-1) Activate shadow volume Volume1 in the read and write access mode (rw).

```
# sdxshadowvolume -F -c Class2 -v Volume1
# sdxshadowvolume -N -c Class2 -v Volume1 -e mode=rw
```

8b-2) Check and repair consistency of the ufs file system on shadow volume Volume1.

If the file system was unmounted when the BCV was split in step 4), this step can be skipped.

```
# fsck -F ufs -y /dev/sfdsk/Class2/rdsk/Volume1
```

8b-3) Mount the ufs file system on shadow volume Volume1 on /mnt1, a temporary mount point, in the read only mode.

```
# mkdir /mnt1
# mount -F ufs -o ro /dev/sfdsk/Class2/rdsk/Volume1 /mnt1
```

8b-4) Back up data held in the file system to tape.

```
# cd /mnt1
# tar cvf /dev/rmt/0 .
```

8b-5) Unmount the file system mounted in step 8b-3).

```
# cd /  
# umount /mnt1  
# rmdir /mnt1
```

8c) When backing up a ufs file system with the ufsdump(1M) command

8c-1) Activate shadow volume Volume1 in the read and write access mode (rw).

```
# sdxshadowvolume -F -c Class2 -v Volume1  
# sdxshadowvolume -N -c Class2 -v Volume1 -e mode=rw
```

8c-2) Check and repair consistency of the ufs file system on shadow volume Volume1.

If the file system was unmounted when the BCV was split in step 3), this step can be skipped.

```
# fsck -F ufs -y /dev/sfdsk/Class2/rdsk/Volume1
```

8c-3) Back up data held in the file system to tape.

```
# ufsdump 0ucf /dev/rmt/0 /dev/sfdsk/Class2/rdsk/Volume1
```

8d) When backing up the sfxfs file system (GFS Local File System) with the tar(1) command

8d-1) Activate shadow volume Volume1 in the read and write access mode (rw).

```
# sdxshadowvolume -F -c Class2 -v Volume1  
# sdxshadowvolume -N -c Class2 -v Volume1 -e mode=rw
```

8d-2) Reconfigure the partition information of the sfxfs file system on shadow volume Volume1.

```
# sfxadm /dev/sfdsk/Class2/rdsk/Volume1
```

8d-3) Check and repair consistency of the sfxfs file system on shadow volume Volume1.

If the file system was unmounted when the BCV was split in step 3), this step can be skipped.

```
# fsck -F sfxfs -y /dev/sfdsk/Class2/rdsk/Volume1
```

8d-4) Mount the sfxfs file system on shadow volume Volume1 on /mnt1, a temporary mount point, in the read only mode.

```
# mkdir /mnt1  
# mount -F sfxfs -o ro /dev/sfdsk/Class2/dsk/Volume1 /mnt1
```

8d-5) Back up data held in the file system to tape.

```
# cd /mnt1  
# tar cvf /dev/rmt/0 .
```

8d-6) Unmount the file system mounted in step 6d-4).

```
# cd /  
# umount /mnt1  
# rmdir /mnt1
```

8e) When backing up the sfxfs file system (GFS Local File System) with the sfxdump(1M) command

8e-1) Activate shadow volume Volume1 in the read and write access mode (rw).

```
# sdxshadowvolume -F -c Class2 -v Volume1
# sdxshadowvolume -N -c Class2 -v Volume1 -e mode=rw
```

8e-2) Reconfigure the partition information of the sfxfs file system on shadow volume Volume1.

```
# sfxadm /dev/sfdsk/Class2/rdsk/Volume1
```

8e-3) Check and repair consistency of the sfxfs file system on shadow volume Volume1.

If the file system was unmounted when the BCV was split in step 3), this step can be skipped.

```
# fsck -F sfxfs -y /dev/sfdsk/Class2/rdsk/Volume1
```

8e-4) Back up data held in the file system to tape.

```
# sfxdump 0ucf /dev/rmt/0 /dev/sfdsk/Class2/rdsk/Volume1
```

9) Removing the shadow volume

After the backup process is complete, remove the shadow volume to prevent improper access to it. The following procedure must be performed on backup server Node3.

9-1) Stopping the shadow volume

Stop shadow volume Volume1.

```
# sdxshadowvolume -F -c Class2 -v Volume1
```

9-2) Removing the shadow volume

Remove shadow volume Volume1.

```
# sdxshadowvolume -R -c Class2 -v Volume1
```

9-3) Removing the shadow group

Remove shadow group Group1.

```
# sdxshadowgroup -R -c Class2 -g Group1
```

9-4) Removing the shadow disk

Remove shadow disk Disk1.

```
# sdxshadowdisk -R -c Class2 -d Disk1
```

10) Resynchronizing the BCV

Resynchronize standard device STD001 and BCV device BCV001 for the following backup. Execute the following command on Node1 or Node2 in the primary domain.

```
# symmir -g DevGroup establish STD001 bcv ld BCV001
```

To back up again, follow the procedure from step 4).

6.6.4.8 Restoring form a BCV

11) Stopping the services

Exit all applications using application volume Volume1 on Node1 and Node2 in the primary domain.

When Volume1 is used as a file system, unmount it.

12) Stopping the application volume

Stop application volume Volume1 on Node1 and Node2 in the primary domain. Execute the following command on Node1 or Node2.

```
# sdxvolume -F -c Class1 -v Volume1 -e allnodes
```

13) Disconnecting any nonrelevant disk from the application volume

In the primary domain, disconnect any disk (Disk2) other than Disk1 that is the restore target from Class1, among disks connected to Class1 to which application volume Volume1 belongs. Execute the following command on Node1 or Node2 in the primary domain.

```
# sdxswap -O -c Class1 -d Disk2
```

14) Restoring from the BCV

Restore data held in standard device STD001 from BCV device BCV001 in the primary domain. The following procedure must be performed on Node1 or Node2 in the primary domain.

14-1) Restoring from the BCV

Restore data held in standard device STD001 from BCV device BCV001.

```
# symmir -g DevGroup restore STD001 BCV 1d BCV001
```

14-2) Viewing the status of restore

When restore is in process, a BCV pair of standard device STD001 and BCV device BCV001 is in the RestInProg status. Confirm that restore is complete and the BCV pair is made the Restored status.

```
# symmir -g DevGroup query
Device Group (DG) Name: DevGroup
DG's Type           : REGULAR
DG's Symmetrix ID   : 000285502123

-----
Standard Device          BCV Device          State
-----
                Inv.                Inv.
-----
Logical   Sym Tracks Logical   Sym   Tracks STD <=> BCV
-----
STD001    005    0 BCV001    073 *    0 Restored
-----
```

15) Resuming the services and reconnecting the disk back to the application volume

Resume the services in the primary domain. The following settings are necessary on the node that runs the service.

Information

In the following example resuming the service is put above resynchronizing the application volume. Through this procedure the services are resumed first and then resynchronization of the volume is secured during the services operation. If resynchronizing the volume should be put above resuming the services, the procedure should be followed in the order of steps 15-1), 15-2), 15-4), 15-5) (confirming that the synchronization copying is complete), and 15-3).

15-1) Activating the application volume

Activate application volume Volume1.

```
# sdxvolume -N -c Class1 -v Volume1
```

15-2) Reconfiguring the partition information of the sfxfs file system

When application volume Volume1 is used as the sfxfs file system, reconfigure the partition information.

```
# sfxadm /dev/sfdsk/Class1/rdisk/Volume1
```

15-3) Resuming the services

When the file system on application volume Volume1 was unmounted in step 11), mount it again.

Start the applications using Volume1.

15-4) Reconnecting the disk to the application volume

Reconnect Disk2 disconnected from Class1 to which application volume Volume1 belongs in step 13).

```
# sdxswap -I -c Class1 -d Disk2 -e nowaitsync
```

After returning from the command, synchronization copying from the slice on Disk1 of volume Volume1 to the slice on Disk2 is executed.

15-5) Viewing the copy status

The status of synchronization copying can be viewed using the `sdxinfo -S` command. The copy destination slice is in COPY status if copying is in progress and it will be in ACTIVE status after the copy process ends normally (note, however, that it will be in STOP status when Volume1 is in STOP status).

```
# sdxinfo -S -c Class1 -o Volume1
OBJ    CLASS  GROUP  DISK    VOLUME  STATUS
-----
slice  Class1 Group1  Disk1   Volume1 ACTIVE
slice  Class1 Group1  Disk2   Volume1 COPY
```

6.6.4.9 Restoring from Tape

16) Stopping the services

Exit all applications using application volume Volume1 on Node1 and Node2 in the primary domain.

When Volume1 is used as a file system, unmount it.

17) Stopping the services application

To write-lock volume Volume1, inactivate Volume1 on Node1 and Node2 in the primary domain. Execute the following command on Node1 or Node2.

```
# sdxvolume -F -c Class1 -v Volume1 -e allnodes
```

18) Detaching any nonrelevant slice from the application volume

Temporarily detach the slice on any disk (Disk2) other than Disk1 that is the restore target from Volume1, among slices in application volume Volume1. Execute the following command on Node1 or Node2 in the primary domain.

```
# sdxslice -M -c Class1 -d Disk2 -v Volume1 -a jrm=off
```

 Point

On backup server Node3, after Disk1 data is restored from tape, the slice on Disk2 is supposed to be reattached to application volume Volume1 in the primary domain. At this point the entire volume data must be copied to the attached slice. For this reason, the JRM mode of a detached slice must be set to off in advance.

 Information

If access can be gained from backup server Node3 to Disk2, data can be restored from tape back to both Disk1 and Disk2 on Node3. Under these circumstances, this procedure (detaching a slice) should not be performed.

19) Viewing the configuration and status of the application volume

On Node1 and Node2 in the primary domain, see the configuration and status of application volume Volume1 that is restore target. Confirm that Volume1 is in STOP status and that only restore target slice Volume1.Disk1 is in STOP status among the slices constituting the volume and the other slices are in TEMP or TEMP-STOP status. If the volume or slice status is invalid, repair it referring to "F.1.3 Volume Status Abnormality" or "F.1.1 Slice Status Abnormality."

```
# sdxinfo -c Class1
OBJ      NAME      TYPE      SCOPE      SPARE
-----
class    Class1    shared    Node1:Node2 0

OBJ      NAME      TYPE      CLASS      GROUP      DEVDNAM      DEVBLKS      DEVCONNECT      STATUS
-----
disk     Disk1     mirror    Class1     Group1     emcpower1    8380800      Node1:Node2     ENABLE
disk     Disk2     mirror    Class1     Group1     emcpower2    8380800      Node1:Node2     ENABLE

OBJ      NAME      CLASS      DISKS              BLKS      FREEBLKS      SPARE
-----
group    Group1    Class1    Disk1:Disk2        8290304    7176192      0

OBJ      NAME      CLASS      GROUP      SKIP  JRM  1STBLK      LASTBLK      BLOCKS      STATUS
-----
volume *      Class1    Group1    *      *      0      65535      65536      PRIVATE
volume Volume1 Class1    Group1    off   on   65536    1114111    1048576    STOP
volume *      Class1    Group1    *      *      1114112  8290303    7176192    FREE

OBJ      CLASS      GROUP      DISK      VOLUME      STATUS
-----
slice   Class1    Group1    Disk1     Volume1     STOP
slice   Class1    Group1    Disk2     Volume1     TEMP
```

Information

When access can be gained from backup server Node3 to all of the disks constituting Volume1 (Disk1 and Disk2), you must confirm that all of the slices of Volume1 are in the STOP status.

20) Creating a shadow volume for restore

On backup server Node3, create a volume for restore (shadow volume) in disk emcpower1. The following settings are necessary on backup server Node3.

20-1) Registering a shadow disk

Register disk emcpower1 with shadow class Class3, and name it Disk1.

```
# sdxshadowdisk -M -c Class3 -d emcpower1=Disk1
```

Point

- The disk name must correspond to the disk name assigned to emcpower1 in step 1-1). The disk names assigned in 1-1) can be viewed in the NAME field displayed with the sdxinfo command in step 19).
- The class can be assigned any name. However, if Node3 resides in the same domain as Node1 and Node2, it must be assigned a name different from the name of the class created in step 1-1).

Information

When access can be gained from backup server Node3 to all of the disks constituting Volume1 (emcpower1 and emcpower2), you must register all of the disks constituting Volume1 (emcpower1 and emcpower2) with a shadow class.

20-2) Creating a shadow group

Connect shadow disk Disk1 to mirror type shadow group Group1.

```
# sdxshadowdisk -C -c Class3 -g Group1 -d Disk1
```

Information

When access can be gained from backup server Node3 to all of the disks constituting Volume1 (emcpower1 and emcpower2), you must connect all of the disks constituting Volume1 (emcpower1 and emcpower2) to a shadow group.

20-3) Creating a shadow volume

Create shadow volume Volume1 to shadow group Group1.

```
# sdxshadowvolume -M -c Class3 -g Group1 -v Volume1 -s 1048576
```

Point

- The volume must be created in the size corresponding to the volume size in step 1-3). The size of a volume created in step 1-3) can be viewed in the BLOCKS field for volume information displayed with the sdxinfo command in step 19).
- If there are multiple volumes, the corresponding shadow volumes must be created in the order of ascending values (first block numbers) in the 1STBLK field for volume information displayed with the sdxinfo command in step 19).

- The volume can be assigned any name.

20-4) Setting the access mode of the shadow volume

Activate shadow volume Volume1 in the read and write access mode (rw).

```
# sdxshadowvolume -F -c Class3 -v Volume1
# sdxshadowvolume -N -c Class3 -v Volume1 -e mode=rw
```

20-5) Viewing the shadow volume configuration

Using the `sdxinfo` command, confirm that the group configuration and the volume configuration are correct based on group information in the DISKS field, volume information in the 1STBLK field and in the BLOCKS field and so on.

```
# sdxinfo -c Class3
OBJ      NAME      TYPE      SCOPE      SPARE
-----
class    Class3    local     Node3      0

OBJ      NAME      TYPE      CLASS      GROUP      DEVDNAM      DEVBLKS      DEVCONNECT      STATUS
-----
disk     Disk1     mirror    Class3     Group1     emcpower1    8380800      Node3            ENABLE

OBJ      NAME      CLASS      DISKS      BLKS      FREEBLKS      SPARE
-----
group    Group1    Class3     Disk1      8290304    7176192      0

OBJ      NAME      CLASS      GROUP      SKIP      JRM      1STBLK      LASTBLK      BLOCKS      STATUS
-----
volume   *         Class3     Group1     *         *         0           65535        65536      PRIVATE
volume   Volume1   Class3     Group1     off      off      65536      1114111      1048576    ACTIVE
volume   *         Class3     Group1     *         *         1114112    8290303      7176192    FREE

OBJ      CLASS      GROUP      DISK      VOLUME      STATUS
-----
slice    Class3     Group1     Disk1     Volume1     ACTIVE
```

21) Restoring from tape

On backup server Node3, restore shadow volume data from tape to which it was backed up in step 8). In the following examples, restore data held in shadow volume Volume1 from a tape medium of tape device `/dev/rmt/0`.



See

For details on the restore method, see the manuals of file systems to be restored and used commands.



Information

In a GFS Shared File System

Restore through the method as described in step 21a).

21a) When restoring data held in a raw device with the `dd(1M)` command


```
# dd if=/dev/rmt/0 of=/dev/sfdsk/Class3/rdisk/Volume1 bs=32768
```

21b) When restoring a ufs file system with the tar(1) command

21b-1) Create a ufs file system to shadow volume Volume1.

```
# newfs /dev/sfdsk/Class3/rdisk/Volume1
```

21b-2) Mount the ufs file system on shadow volume Volume1 on /mnt2, a temporary mount point.

```
# mkdir /mnt2
# mount -F ufs /dev/sfdsk/Class3/dsk/Volume1 /mnt2
```

21b-3) Restore data held in the file system from tape.

```
# cd /mnt2
# tar xvf /dev/rmt/0
```

21b-4) Unmount the file system mounted in step 21b-2).

```
# cd /
# umount /mnt2
# rmdir /mnt2
```

21c) When restoring a ufs file system with the ufsrestore(1M) command

21c-1) Create a ufs file system to shadow volume Volume1.

```
# newfs /dev/sfdsk/Class3/rdisk/Volume1
```

21c-2) Mount the ufs file system on shadow volume Volume1 on /mnt2, a temporary mount point.

```
# mkdir /mnt2
# mount -F ufs /dev/sfdsk/Class3/dsk/Volume1 /mnt2
```

21c-3) Restore data held in the file system from tape.

```
# cd /mnt2
# ufsrestore rvf /dev/rmt/0
```

21c-4) Delete the temporary file created by the ufsrestore(1M) command.

```
# rm /mnt2/restoresymtable
```

21c-5) Unmount the file system mounted in step 21c-2).

```
# cd /
# umount /mnt2
# rmdir /mnt2
```

21d) When restoring the sfxfs file system (GFS Local File System or SynfinityFile file system) with the tar(1) command.

21d-1) Create the sfxfs file system to shadow volume Volume1.

```
# sfxnewfs /dev/sfdsk/Class3/rdisk/Volume1
```

21d-2) Mount the sfxfs file system on shadow volume Volume1 on /mnt2, a temporary mount point.

```
# mkdir /mnt2
# mount -F sfxfs /dev/sfdsk/Class3/dsk/Volume1 /mnt2
```

21d-3) Restore data held in the file system from tape.

```
# cd /mnt2
# tar xvf /dev/rmt/0
```

21d-4) Unmount the file system mounted in step 21d-2).

```
# cd /
# umount /mnt2
# rmdir /mnt2
```

21e) When restoring the sfxfs file system (GFS Local File System or SynfinityFile file system) with the sfxrestore(1M) command

21e-1) Create the sfxfs file system to shadow volume Volume1.

```
# sfxnewfs /dev/sfdsk/Class3/rdsk/Volume1
```

21e-2) Mount the sfxfs file system on shadow volume Volume1 on /mnt2, a temporary mount point.

```
# mkdir /mnt2
# mount -F sfxfs /dev/sfdsk/Class3/dsk/Volume1 /mnt2
```

21e-3) Restore data held in the file system from tape.

```
# cd /mnt2
# sfxrestore rvf /dev/rmt/0
```

21e-4) Delete the temporary file created by the sfxrestore(1M) command.

```
# rm /mnt2/restoresymtable
```

21e-5) Unmount the file system mounted in step 21e-2).

```
# cd /
# umount /mnt2
# rmdir /mnt2
```

22) Removing the shadow volume

After the restore process is complete, remove the shadow volume to prevent improper access to it. The following procedure must be performed on backup server Node3.

22-1) Stopping the shadow volume

Stop shadow volume Volume1.

```
# sdxshadowvolume -F -c Class3 -v Volume1
```

22-2) Removing the shadow volume

Remove shadow volume Volume1.

```
# sdxshadowvolume -R -c Class3 -v Volume1
```

22-3) Removing the shadow group

Remove shadow group Group1.

```
# sdxshadowgroup -R -c Class3 -g Group1
```

22-4) Removing the shadow disk

Remove shadow disk Disk1.

```
# sdxshadowdisk -R -c Class3 -d Disk1
```



Information

When access can be gained from backup server Node3 to all of the disks constituting Volume1 (emcpower1 and emcpower2), you must remove all of the disks registered with shadow class Class3 in step 20) (emcpower1 and emcpower2).

23) Resuming the services and reattaching the slice to the application volume

Resume the services in the primary domain. The following settings are necessary on the node that runs the services.



Information

In the following example resuming the services is put above resynchronizing the application volume. Through this procedure the services are resumed first and then resynchronization of the volume is secured during the services operation. If resynchronizing the volume should be put above resuming the services, the procedure should be followed in the order of steps 23-1), 23-2), 23-4), 23-5) (confirming that the synchronization copying is complete), and 23-3).

23-1) Activating the application volume

Activate application volume Volume1.

```
# sdxvolume -N -c Class1 -v Volume1
```

23-2) Reconfiguring the partition information of the sfxfs file system

When application volume Volume1 is used as the sfxfs file system, reconfigure the partition information.

```
# sfxadm /dev/sfdsk/Class1/rdisk/Volume1
```

23-3) Resuming the services

When the file system on application volume Volume1 was unmounted in step 16), mount it again.

Start the applications using Volume1.

23-4) Reattaching the slice to the application volume

Reattach slice Volume1.Disk2 temporarily detached from application volume Volume1 in step 18).

```
# sdxslice -R -c Class1 -d Disk2 -v Volume1
```

After returning from the command, synchronization copying from the slice on Disk1 of volume Volume1 to the slice on Disk2 is executed.

23-5) Viewing the copy status

The status of synchronization copying can be viewed using the `sdxinfo -S` command. The copy destination slice is in `COPY` status if copying is in progress and it will be in `ACTIVE` status after the copy process ends normally (note, however, that it will be in `STOP` status when `Volume1` is in `STOP` status).

```
# sdxinfo -S -c Class1 -o Volume1
OBJ     CLASS   GROUP   DISK     VOLUME   STATUS
-----
slice   Class1  Group1  Disk1    Volume1  ACTIVE
slice   Class1  Group1  Disk2    Volume1  COPY
```

6.7 Backing Up and Restoring Object Configurations

Due to multiple disk failures, the valid configuration database of a class may be lost, resulting in loss of configuration information of objects within the class. If that happens, after recovering the failed disks, objects such as volumes must be re-created. By backing up object configuration information in advance, efforts to restore object configurations can be reduced.

This section describes the procedures for backing up and restoring object configurations of classes.



Note

Classes Unavailable for Object Configuration Backup

Object configurations of the following classes cannot be backed up.

- Root class
- Shared class that include a switch group
- Class that include a proxy object
- Shadow class



Note

Systems Available for Object Configuration Restore

To restore object configuration according to backed up object configuration information, the system for restoration must be connecting disks that are equivalent in size to physical disks registered with the backed up class.

6.7.1 Backing Up

This subsection describes the procedures for backing up configuration information of objects within class `Class1`.

1) Saving configuration information

Save outputs of the `sdxinfo` command to a file. In this example, the path to a file is `"/var/tmp/Class1.info"`.

```
# sdxinfo -c Class1 -e long > /var/tmp/Class1.info
```

2) Creating a configuration file

Output the object configuration within `Class1` to a file in configuration table format. In this example, the path to a file is `"/var/tmp/Class1.conf"`.

```
# sdxconfig Backup -c Class1 -o /var/tmp/Class1.conf
```

3) Backing up the configuration file and configuration information

Save the files created in steps 1) and 2) to tape and so on.

6.7.2 Restoring

This subsection describes the procedures for restoring the object configuration within class Class1 according to the configuration file and configuration information saved in advance as shown in "6.7.1 Backing Up" in the event of loss of the Class1 object configuration caused by a problem of some kind.

1) Checking the class scope

With a cluster system, check names of nodes sharing the class. For node names that belong to the class scope, check the SCOPE field of class information output by the `sdxinfo` command and saved as shown in step 1) of "6.7.1 Backing Up."

2) Placing the configuration file

On a node where the object configuration of the class is restored (with a cluster system, on a node that belongs to the class scope), place the configuration file created in step 2) of "6.7.1 Backing Up" In this example, the path to a destination file is `/var/tmp/Class1.conf`.

3) Restoring the object configuration of the class

Execute the following command on the node where the configuration file was placed in step 2) to restore the object configuration of class Class1 according to descriptions in the configuration file `/var/tmp/Class1.conf`. Class1 is restored as a local class of that node.

After restoring the object configuration, reboot the node.

```
# sdxconfig Restore -c Class1 -i /var/tmp/Class1.conf
# shutdown -y -i6 -g0
```

Information

If the Physical Disk Configurations Are Different

If the physical disk configuration of the system for restoration is different from that of the backed up system, use the `sdxconfig Convert` command and change physical disk names in the configuration file.

(Example 1)

Change a physical disk described in the configuration file `/var/tmp/Class1.conf` from `c3t4d5` to `c2t3d4`.

```
# sdxconfig Convert -e replace -c Class1 -p c3t4d5=c2t3d4 -i /var/tmp/
Class1.conf -o /var/tmp/Class1.conf -e update
```

(Example 2)

Change the physical disk of Disk1 described in the configuration file `/var/tmp/Class1.conf` to `c2t3d4`.

```
# sdxconfig Convert -e replace -c Class1 -d Disk1=c2t3d4 -i /var/tmp/
Class1.conf -o /var/tmp/Class1.conf -e update
```

4) Changing the class type and expanding the class scope

If the backed up class, Class1, is a shared class, change the type and scope attributes of Class1. In this example, the scope of the backed up class is `node1:node2`.

4-1) Stop the volume in the class.

```
# sdxvolume -F -c Class1
```

4-2) Change the class type and expand the class scope.

```
# sdxattr -C -c Class1 -a type=shared,scope=node1:node2
```

Appendix A General Notes

A.1 Rules

A.1.1 Object Name

Users can name the following objects:

- Classes
- Disks (excluding shadow disk)
- Groups
- Volumes

The object name can contain a maximum of 32 alphanumeric characters, including the hyphen (-) and the underscore character (_). However, in the event of the single disk name, the length limit is a maximum of 28 alphanumeric characters.

The object name cannot start with the hyphen (-) or the underscore character (_). Be sure to assign an alphanumeric character to the first character in the object name.

The class name must be unique within the entire system (for a cluster system, within the entire cluster system). For this reason, an error occurs if you try to create more than one disk class with the same name.

The other object names are unique within the class and an attempt to create more than one object with the same name in a class will result in an error.

GDS assigns slice names by combining the names of disks, groups, and volumes to which the slices belong. The slice naming conventions are as follows.

- When the slice is a mirror slice and exists in a disk that is connected directly to the highest level mirror group:
disk_name.volume_name
- When the slice is a mirror slice and exists in a lower level group that is connected directly to the highest level mirror group:
lower_level_group_name.volume_name
- When the slice belongs to a stripe volume:
the_highest_level_stripe_group_name.volume_name
- When the slice belongs to a volume created within the highest level concatenation group:
the_highest_level_concatenation_group_name.volume_name
- When the slice belongs to a switch volume:
active_disk_name.volume_name
- When the slice is a single slice:
single_disk_name.volume_name

You can designate each object uniquely in the entire system by the object name and the class name to which it belongs.



Note

Same Class Names

Multiple single nodes on which classes with the same name exist can be changed over to a cluster system through installation of the cluster control facility. For details, see "[A.2.26 Changing Over from Single Nodes to a Cluster System.](#)"



Note

Shadow Disk Name

Name shadow disks according to the following rules:

- When the shadow disk is already registered with a class in another domain and managed as an SDX disk, assign the SDX disk name in said domain.
- When the shadow disk contains data copied from an SDX disk with the disk unit's copy function, assign the copy source SDX disk name.

A.1.2 Number of Classes

The number of root classes you can create for one node is limited to one.

There is no limit to the number of local classes and shared classes.

Separate classes conforming to the following rules.

- Register system disks with a root class.
- It is recommended to register local disks other than system disks (disks used on one node) with local classes, but not with a root class, to differentiate the local disks from the system disks in management.
- Register shared disks in a cluster system (disks used from multiple nodes in the cluster) with shared classes.
- In a cluster system, register shared disks whose scopes (groups of sharing nodes) are different with separate shared classes.
- In a cluster system, for applications to use shared disks, create one or more shared classes with respect to each cluster application.
- Divide the class if the number of created disks or volumes in a class exceeds the limit. For the numbers of disks and volumes, see "[A.1.3 Number of Disks](#)" and "[A.1.5 Number of Volumes](#)."
- In a large-scale system to which numerous disk units are connected, separating classes based on physical configurations and data contents of disks may bring higher manageability.
- When a disk unit is expanded, unless the rules above apply, register expanded disks to existing classes. Creating new classes is not required.

Do not separate classes more than necessary. Keeping the number of classes to a minimum will offer the following advantages.

- If a class includes more disks, the probability that GDS configuration information stored on disks is lost due to disk failure will be lower. Therefore, not separating classes more than necessary to increase the number of disks within one class will raise system reliability.
- In a cluster system, if there are less shared classes, it takes shorter time to switch nodes.
- If there are fewer classes, less memory resources will be required.

A.1.3 Number of Disks

The number of disks you can register with one class has the following limitations:

- To root class, you can register up to 100 disks.
- To local or shared class, you can register up to 1024 disks.

There are the following limits to the number of disks that can be connected to one group:

- To a mirror group, a maximum of 8 disks and lower level groups can be connected collectively. In other words, a maximum of eight-way multiplex mirroring is supported. However, be aware that the spare disk that will be automatically connected when a disk failure occurs is also included in the count.
- To a stripe group, a maximum of 64 disks and lower level groups can be connected collectively. In other words, a maximum of 64-column striping is supported.
- To a concatenation group, a maximum of 64 disks can be connected. In other words, a maximum of 64 disks can be concatenated.

- To a switch group, a maximum of 2 disks can be connected.

A.1.4 Number of Groups

The number of groups you can create within one class has the following limitations:

- Within root class, you can create up to 100 groups.
- Within local or shared class, you can create up to 1024 groups.

A.1.5 Number of Volumes

There are the following limits to the number of volumes that can be created within a group in the root class:

- The number of volumes with the physical slice attributes set to "on" is limited to a maximum of 6.
- Volumes with the physical slice attributes set to "off" cannot be created.

There are the following limits to the number of volumes that can be created within a group or a single disk in the local class or shared class:

- You can create a maximum of 6 volumes with their physical slice attribute set to "on."
- You can create a total of 1024 volumes with their physical slice attribute set to "on" or "off."
- You cannot create a volume with the physical slice attribute set to "on" in a stripe group or a concatenation group.
- You cannot create a volume with the physical slice attribute set to "on" in a shadow class regardless whether the physical slice is registered with VTOC.

In addition, there are the following limits to the number of volumes that can be created within a class:

- For the root class, the number is limited to a maximum of 256.
- For a local class, the number is limited to a maximum of 6144.
However, when groups are nested, the nested group can contain a maximum of 6144 volumes and lower level groups collectively.

A.1.6 Number of Keep Disks

With root class, you can register up to 100 keep disks.

A.1.7 Creating Group Hierarchy

The following nine kinds of group hierarchical structures, including a nonhierarchical structure, are available. However, groups that can be created in the root class are only non-nested mirror groups.

```
higher level group <-----> lower level group
mirror group (*1)
mirror group (*1) - stripe group (*3)
mirror group (*1) - stripe group (*3) - concatenation group (*7)
mirror group (*1) - concatenation group (*5)
stripe group (*2)
stripe group (*2) - concatenation group (*6)
concatenation group (*4)
concatenation group (*4) - switch group (*9)
switch group (*8)
```

Possible operations on groups at each hierarchical level that change the structure are as follows.

(*1) The highest level mirror group

- Disks, lower level stripe groups, and lower level concatenation groups can be connected or disconnected. However, disconnection is impossible if it can change the volume configuration or status.
- If no volume exists, the group itself can be removed.

- Volumes can be created or removed.

(*2) The highest level stripe group

- If no volume exists, disks and lower level concatenation groups can be connected or disconnected.
- If no volume exists, this type group can be connected to a mirror group.
- If no volume exists, the group itself can be removed.
- If more than one disks or lower level concatenation groups are connected to, volumes can be created or removed.

(*3) The lower level stripe group

- This type group can be disconnected from the highest level mirror group. However, disconnection is impossible if it can change the volume configuration or status.

(*4) The highest level concatenation group

- If no switch group is connected to, disks can be connected.
- If no disk is connected to, lower level switch groups can be connected.
- If no volume area exists on the disk that was connected last, that disk can be disconnected.
- If no volume area exists in the group that was connected last, that lower switch group can be disconnected.
- If no volume exists and if no lower level switch group is connected to, this group can be connected to a mirror group or a stripe group.
- If no volume exists, the group itself can be removed.
- Volumes can be created or removed.

(*5) Lower level concatenation group connected to the highest level mirror group

- Disks can be connected.
- If more than one disk is connected and no volume area exists on the disk that was connected last, that disk can be disconnected.
- This type group can be disconnected from the highest level mirror group. However, disconnection is impossible if it can change the volume configuration or status.

(*6) Lower level concatenation group connected to the highest level stripe group

- If more than one disk is connected and no volume exists within the highest level group, the disk that was connected last can be disconnected.

(*7) Lower level concatenation group connected to a lower level stripe group

- None.

(*8) The highest level switch group

- Disks can be connected.
- If no volume or inactive disk exists, the active disk can be disconnected.
- The inactive disk can be disconnected.
- If no volume exists, the group itself can be removed.
- Volumes can be created or removed.

(*9) Lower level switch group

- Disks can be connected.

- The inactive disk can be disconnected.
- If no volume area exists and if this group is the switch group that was last connected to a higher level concatenation group, this group can be disconnected from that concatenation group.

A.1.8 Proxy Configuration Preconditions

General Preconditions

- The master and proxy belong to the same class (excepting shadow classes).
- The master and proxy belong to different groups or single disks.
- The master is not related to any object as its proxy.
- To the proxy, no other object is related as its proxy.
- The type of master and proxy is mirror or single.

Synchronization Snapshot Preconditions

- To stripe the master and proxy, connect a stripe group to a mirror group and use them as the master and proxy. In this situation, the disk unit's copy function is unavailable.
- To concatenate disks to create large master and proxy objects, connect a concatenation group to a mirror group and use them as the master and proxy. In this situation, the disk unit's copy function is unavailable.

OPC Snapshot Preconditions

- Between the master and proxy, OPC or ROPC is available.
- To the master and proxy groups, no lower level group is connected.
- To create snapshots of volumes in the specified group at a time, the layout (offsets and sizes) of volumes within a master group and a proxy group must be consistent.
- See also "[A.2.18 Using the Advanced Copy Function in a Proxy Configuration](#)" and "[A.2.19 Instant Snapshot by OPC](#)."

A.1.9 Number of Proxy Volumes

To one master volume, multiple proxy volumes can be related.

The number of proxy volumes that can be related is limited to meet the following conditions:

- The total number of slices composing one master volume and slices composing any proxy volumes that are related to the master volume cannot exceed 32.

For example, if all master and proxy volumes consist of single volumes, a maximum of 31 proxy volumes can be related to one master volume.

A.1.10 Proxy Volume Size

The size of a proxy volume must be equal to that of the master volume to which it is related.



Note

System Volume's Proxy Volume

When relating each individual proxy volume to a system volume, the proxy volume must be created in the group to which the keep disk with the same cylinder size as the system disk's cylinder size belongs. This is because the system volume size is the multiple of the cylinder size.

When relating each individual proxy group to a group to which system volumes belong, it is not necessary to consider the cylinder size. This is because the proxy group's cylinder size changes to the same size as the joined master group's cylinder size.

A.1.11 Proxy Group Size

For a local class or a shared class, the size of a proxy group must be larger than the last block number of volumes within the master group to which the proxy is related.

For a root class, the size of the smallest physical disk that is directly connected to a proxy group must be larger than the last block number of volumes within the master group to which the proxy is related.

A.2 Important Points

A.2.1 Physical Disk Driver

GDS can only manage physical disks controlled by the following device drivers. In other words, physical disks controlled by other disk drivers cannot be registered with classes.

sd

SCSI disk driver
(Solaris standard equipment)

ssd

Device driver for disk arrays and Fibre Channel Arbitrated Loop (FC-AL) disks
(Solaris standard equipment)

dad

IDE disk driver
(Solaris standard equipment)

hddv

Fujitsu's disk array device driver

mphd

Multi-path driver from Fujitsu

mplb

Multi-path driver from Fujitsu

emcpower

Multi-path driver from EMC

emcp

Multi-path driver from EMC

vpathdd

Multi-path driver from Hitachi

dlimfdrv

Multi-path driver from Hitachi

mpd

Multi-path driver from Fujitsu Technology Solutions

Physical disks that can be registered with the root class are only those controlled by the following disk drivers.

sd, ssd, dad, mplb

When the above requirements are not met, the `sdxdisk -M` command and the `sdxshadowdisk -M` command will end abnormally and output one of the following error messages.

```
SDX: sdxdisk: ERROR: physical_disk_name: disk driver driver_name not supported
SDX: sdxdisk: ERROR: physical_disk_name: disk driver driver_name not supported for root class
SDX: sdxshadowdisk: ERROR: physical_disk_name: disk driver driver_name not supported
```

A.2.2 Managing System Disks

Volumes in local classes and shared classes cannot be used as:

`/` (root), `/usr`, `/var`, swap area, dump device, savecore directory

To manage disks for these uses with GDS, register the disks with the root class.



See

.....
For details on the dump device and the savecore directory, see the `dumpadm(1M)` command manual pages.
.....

A.2.3 Restraining Access to Physical Special File

After registration of disks with GDS is complete, accessing the disks using physical special files as below becomes impossible. GDS prevents unintentional access by the users to physical special files in order to protect mirroring statuses.

```
/dev/[r]dsk/cCtTdDs[0-7]          (for normal hard disks)
/dev/FJFSVmphd/[r]dsk/mphd/ s[0-7] (for MPHD disks)
/dev/FJFSVmplb/[r]dsk/mplb/ s[0-7] (for MPLB disks)
```

C indicates the controller; *T*, the target ID; *D*, the disk number; *I*, the MPHD instance number or the MPLB instance number; [0-7], the slice number with 1 digit from 0 to 7.

For example, if you execute the `dd` command to write to the disk with physical special file, the following error message will be displayed.

```
# dd if=/dev/zero of=/dev/rdsk/c1t1d0s0
dd: /dev/rdsk/c1t1d0s0: open: Device busy
```

Disk list output by the `format(1M)` command does not display the disk registered with GDS.

This access protection is cancelled when a disk is removed from a class. It will also be cancelled when an object is used for disk swap. The access protection function will turn back on if disk recovery operation is performed on the object.

However, when using EMC PowerPath 2.x.x, this function is not effective for `emcpower` devices. You must exercise care not to access the following special physical files, since access is not restrained.

```
/dev/[r]dsk/emcpowerN{a-h}
```

N indicates the `emcpower` device number; {a-h}, the slice number with one alphabet from a to h.

Access to physical special files is prevented on the following nodes:

- For disks that are registered with disk classes
 - When registered with the PRIMECLUSTER's resource database
 - All of the nodes in the relevant domains
 - When not registered with the PRIMECLUSTER's resource database
 - Nodes where the relevant disk classes (root classes or local classes) reside
- For disks that are registered with shadow classes
 - Nodes where the relevant shadow classes reside

For the following disks, access to physical special files is not prevented. Be careful not to access these physical special files.

- Disks that are registered with classes in other domains
- Disks that are registered with root classes or local classes on other nodes and for which disk resources are not registered with the resource database
- Disks that are registered with shadow classes on other nodes

A.2.4 Booting from a CD-ROM Device (or Network)

If the system cannot be booted, for example, it may be required recovering the system booting from a CD-ROM device (or network).

However, there are the following concerns when booting from a CD-ROM device (or network).

- An operation mistake can easily be caused since circumstances when booting with this method and when booting from a boot disk may alter the correlation between the disk and the device special file (`/dev/[r]disk/cTdS`).
- Only partial mirrored disks can be written in since access to the device special file of a physical slice is not restrained. Therefore, the mirroring state could collapse.

For these reasons, avoid booting from a CD-ROM device (or network) unless directed in this manual.

When mounting a file system by booting from a CD-ROM device (or network) in any procedure not described in this manual from necessity, it should be mounted as read-only.

A.2.5 Initializing Disk

When physical disks are registered with classes other than shadow classes, all data contained in the physical disks will be lost since GDS automatically reformats the disks (excepting when the disks are registered with the root class as keep disks). Therefore, when registering physical disks that contain data with classes, back up disk data in advance, and restore the data after creating volumes.

If the device special file name of the physical slice on a physical disk registered with a class is set to `/etc/vfstab` and so on, it is required to change it to the device special file name for the following volume.

Character device special file:

```
/dev/sfdisk/class_name/rdisk/volume_name
```

Block device special file:

```
/dev/sfdisk/class_name/disk/volume_name
```

If a disk restore operation is performed after swapping physical disks, data contained in the physical disk will also be lost.

When using GDS Management View, a confirmation window will be displayed if an operation that can cause data loss is attempted. However, when using a command, perform operations carefully because such a confirmation message is not displayed.

In order to register physical disks with classes other than shadow classes ensuring disk data, it is necessary to register the physical disks with the root class as keep disks.

A.2.6 Disk Size

Out of the physical disk area, the capacity available for creating volumes equals the physical disk size rounded down to the cylinder boundary, minus the private slice size. This size is called the available size of disk.

Private slice is an area that GDS reserves on each disk for storing data, such as configuration information and the JRM log.

When a physical disk is registered with a class, a private slice is reserved. Disks within the same class will all have the same private slice size. The size is determined by the size of the disk that is first registered with the class.

The size of a private slice can be estimated as below. The private slice size never exceeds the following estimated values.

- When the size of the disk first registered with the class is 10GB and below:
32MB

- When the size of the disk first registered with the class is over 10GB:
32MB + (0.1% of the disk capacity to the cylinder boundary, rounded-up)

The size of a log for JRM is determined by the size of the volume. As a result, the log area for JRM may not be reserved because the private slice size becomes insufficient when a mirror volume is created registering a disk larger in size than the first registered disk with the class. In such a case, the volume is incapable of using JRM. Therefore you are recommended to first register a disk of the maximum size with the class.

When physical disks are registered with shadow classes, the sizes of private slices are determined based on those saved in the VTOC of the physical disks. For this reason, regardless of the order of registering physical disks with shadow classes, the sizes of private slices of the shadow classes become consistent with those of the corresponding disk classes. In other words, the order of registering physical disks with shadow classes is not a cause for concern.



Note

Cylinder Size

The cylinder size of a disk registered with a local class or a shared class is 32,768 blocks (= 16 MB). To calculate the sizes of disks to register with local classes or shared classes, assume that the cylinder size is 16 MB.



Note

Large Disks

1 TB or larger disks cannot be managed with GDS.

A.2.7 Volume Size

The size of a volume is automatically adjusted conforming to the following conditions.

- When creating a volume into the group that a keep disk belongs to

A volume is created in the size calculated by rounding-up the size specified when creating the volume to a multiple of the keep disk's cylinder size.

(Example)

If the keep disk's cylinder size is 8 MB (= 16,384 blocks) and 20 MB (= 40,960 blocks) is specified as the size of a volume when it is created, the size is rounded to a multiple of 16,384 blocks, and a volume with 49,152 blocks (= 24 MB) is created.

- When creating a volume into a stripe group

A volume is created in the size calculated by rounding-up the size specified when creating the volume to a common multiple of (stripe width) * (stripe column number) and the cylinder size (32,768 blocks = 16 MB).

(Example)

If the stripe width is 32, the number of the stripe columns is 3 and 20 MB (= 40,960 blocks) is specified as the size of a volume when it is created, the size is rounded to a common multiple of 96 (= 32 * 3) blocks and 32,768 blocks, and a volume with 98,304 blocks (= 48 MB) is created.

- When creating other volumes

A volume is created in the size calculate by rounding-up the size specified when creating the volume to a multiple of the cylinder size (32,768 blocks = 16 MB).

(Example)

If 20 MB (= 40,960 blocks) is specified as the size of a volume when it is created, the size is rounded to a multiple of 32,768 blocks, and a volume with 65,536 blocks (= 32 MB) is created.

A.2.8 Hot Spare

- Hot Spare for Hierarchized Mirror Groups


If an I/O error occurs in a disk that is connected to a lower level group, a spare disk is automatically connected to the highest level mirror group, but not to the lower level group. For example, if a disk and a concatenation group are mirroring each other, and an I/O error occurs in one of the disks connected to the concatenation group, a spare disk is automatically connected in place of the concatenation group for mirroring.

A spare disk is selected independently of the disk case or controller number of a disk where an I/O error occurred.

In external hot spare mode (default), spare disks are selected randomly.

In internal hot spare mode, spare disks whose controller number is 0 and that do not belong to disk array units are selected.

Figure A.1 Hot Spare for Hierarchized Mirror Groups



- Number of Spare Disks

There is no limit to the number of spare disks that can be registered with one class. Although there is no general rule in deciding the number of spare disks, it is recommended to assign 10% of disks and lower level groups for spare disks. In other words, one spare disk for every 10 disks or lower level groups combined is a good rule of thumb.

- Spare Disk Size

Spare disk automatic connection is restrained if there is not sufficient space on a spare disk to copy volumes in a mirror group. It is recommended to assign the largest size disk within the class for the spare disk.

- Hot Spare for Proxy Volumes

Spare disks are not connected to groups that include proxy volumes. It is recommended to create proxy volumes in groups other than those that include volumes used by primary services, or on single disks.

- Shadow Classes

Spare disks cannot be registered with shadow classes.

- Disk Array Unit's Hot Spare Function

If disk array units with hot spare functions are mirrored, it is recommended to use their own hot spare functions.

- Spare Disk Failure

If an I/O error occurs in a spare disk that was automatically connected to a mirror group, another spare disk will not automatically be connected in place of the failed spare disk.

- Synchronization Copying Invoked by Hot Spare

The synchronization copying with hot spare is run at lower speed as compared to similar copying with other events (such as volume creation and disk creation) in order to suppress the load imposed on the system. By default, delay time by 50 milliseconds is set. To change this delay time, use the `sdxparam` command.

- Required time for Synchronization Copying Invoked by Hot Spare

The required time for synchronization copying invoked by hot spare depends on the performance of the CPU or disk. You can estimate the indication of the required time by the following formula:

Total volume size (GB) x 2 (minute) + (Total volume size (Block) / 128) x spare_copy_delay (milliseconds)

You can check the value of `spare_copy_delay` (delay time for synchronization copying invoked by hot spare) with the `sdxparam -G`.



For details, see "D.12 `sdxparam` - Configuration parameter operations."

- Spare Disk Manual Connection

In the within-case hot spare mode, if a disk case totally becomes inaccessible due to an I/O cable coming out or disk case going down, a spare disk that belongs another disk case is not automatically connected. For example, if disk case 1 shown in [Figure 1.10 Hot Spare in Internal Mode] of "1.2.2 Hot Spare" is down, a spare disk (disk 4) is not automatically connected in place of disk 1.

In such an event, follow the procedures below to manually recover the mirroring status by using a spare disk.

1. Change the spare disk to an undefined disk.



For disk type changing methods, see "Changing the Class Attributes" in "5.4.1 Class Configuration" when using the GDS Management View, or "D.7 `sdxattr` - Set objects attributes" when using the command.

2. Connect the disk in step 1. to the mirror group where the I/O error occurred.



For the disk connection methods, see "5.4.2 Group Configuration" when using the GDS Management View, or the description about the `-C` option in "D.2 `sdxdisk` - Disk operations" when using the command.

A.2.9 System Disk Mirroring

You can mirror the system disk containing the operating root file system by registering the system disk with root class. When registering a system disk, you must specify "keep" for its disk type.

You cannot register system disks with local class and shared class.

When mirroring a system disk, GDS additionally sets the `OpenBoot nvramrc` parameter. It will also set `use-nvramrc?` parameter to "true" so that the command registered with `NVRAMRC` will automatically be executed when booting the system.

However, when a command is already registered with `NVRAMRC` and `use-nvramrc?` is set to "false", a following warning message will be displayed, and `use-nvramrc?` will not be changed to "true."

```
WARNING: cannot set use-nvramrc? parameter to true, NVRAMRC not empty
```

When `use-nvramrc?` is set to "false", removing the command registered with `NVRAMRC` is recommended.

When the warning message is displayed, follow the procedures below and set `use-nvramrc?` to "true."

1. Confirm the contents of `NVRAMRC`. When the setting is correct, change the `use-nvramrc?` parameter to "true" with `eprom(1M)`.

2. If the contents of NVRAMRC are invalid, remove the command registered with NVRAMRC, and reboot the system. use-nvramrc? will be automatically changed to "true."

Information

- How to confirm use-nvramrc? parameter:

```
# eeprom use-nvramrc?  
use-nvramrc?=false
```

- How to change and confirm use-nvramrc? parameter:

```
# eeprom use-nvramrc?=true  
# eeprom use-nvramrc?  
use-nvramrc?=true
```

- How to confirm the command registered with NVRAMRC:

```
# eeprom nvramrc
```

(Registered command will be displayed.)

- How to remove and confirm the command registered to NVRAMRC:

```
# eeprom nvramrc=  
# eeprom nvramrc  
nvramrc: data not available.
```

Note

- In the root class, the logical partitioning, disk concatenation, and disk striping functions are not available.
- When mirroring a system disk, GDS additionally sets the following parameters in addition to the use-nvramrc? and nvramrc parameters. After configuring mirroring, do not edit these parameters:
 - boot-device
 - bd-ix? (only for PRIMEPOWER)
 - multipath-boot? (only for SPARC Enterprise)

A.2.10 Keep Disk

By specifying a disk as "keep" type and registering it with root class, you can retain the disk data while configuring mirroring. However, you cannot register the disk under the following conditions. In such case, change the disk configuration and try registering again.

- The disk's slice numbers are all in use. In other words, the number of slices created on the disk has reached the upper limit (seven).
- There is not sufficient free space area or swap area on the disk for private slice.
- In 32-bit Solaris, there is a swap device with capacity equal to or larger than 2 GB.

For details on free space area (or swap area), see "[A.2.6 Disk Size](#)."

Additionally, when registering physical disks other than system disks (with running /, /usr, /var, or swap areas) as keep disks, all slices on the physical disks must be in open status. For example, if the slice is mounted as a file system, unmount it in advance.

You cannot register a keep disk with local class or shared class.

A.2.11 GFS Partition on a Keep Disk

Registering a keep disk with the root class causes the device special files of file systems within the keep disk to be renamed. For this reason, when a keep disk registered with the root class contains a partition used by the sfxfs file system, you must make the following modifications for the relevant partition prior to rebooting the system after executing [System Disk Settings] in GDS Management View or the `sdxroot -M` command.

- Changing the partition information settings

For details on the method of changing the settings, see "PRIMECLUSTER Global File Services Configuration and Administration Guide" or the manual page of the `sfxadm(1M)` command.

- Modifying the `/etc/vfstab` file

You must change the device special file name from that for the physical slice to that for the logical volume.

In addition, a GFS Local File System cannot be used as `/(root)`, `/usr`, `/var`, or `/opt`. Also a partition on the system disk cannot be used in a GFS Shared File System. For details, see "PRIMECLUSTER Global File Services Configuration and Administration Guide."

A.2.12 Creating a Snapshot by Slice Detachment

A slice can be detached only from a mirror volume with a physical slice. Therefore, if a disk is not connected directly to a mirror group, snapshot creation by slice detachment is impossible. In addition, a slice cannot be detached from a shadow volume.

A.2.13 The Difference between a Mirror Slice and a Proxy Volume

Although data matches on mirrored slices or synchronized master volumes and proxy volumes, the purposes of use are different.

Mirrored slices are equal one another, and their purpose is to maintain data redundancy in order to provide continuous access as long as any normal slice remains even if an error occurs in one of the slices.

However, even if the master volume and the proxy volume are synchronized, they are separate volumes and not equals. You may consider the master the primary volume, and the proxy the secondary volume. This means that you cannot continue accessing a master volume where all slices comprising the master volume are abnormal even if proxy volumes are normal. The purpose of proxy volumes is to create snapshots (saved copies from the master volume at certain moment) for a different service running concurrently with the primary service but not to improve the data redundancy of the master volume used in the primary service.

While the function of GDS Snapshot creating snapshots by detaching slices is a by-product of mirroring, the function itself is the primary purpose of GDS Snapshot by proxy volumes. Therefore, the application of proxy volumes provides more flexible disk configurations and service styles for snapshot management.



.....
See "Figure 1.33 Difference between a Mirrored Slice and Synchronized Proxy Volume" in "1.5.1 Snapshot by Synchronization."
.....

A.2.14 Just Resynchronization Mechanism (JRM)

There are three types of Just Resynchronization Mechanism (JRM): for volumes, for slices and for proxies.

JRM for Volumes

JRM for volumes speeds up the resynchronization process when booting the system after a system panic or the like. GDS records the changed portion in the private slice. The resynchronization copy performed at rebooting after an unexpected system failure copies the only portion that was written during the system down to realize high-speed resynchronization and minimize the load of copy processing.



-
- For details on the setting methods, see "5.2.3.4 Volume Configuration" and "D.4 sdxvolume - Volume operations."
 - For details on the changing methods, see "5.4.3 Volume Configuration" and "D.7 sdxattr - Set objects attributes."

- When using GDS Management View, the mode ("on" or "off") of JRM for volumes can be checked using the volume information field in the Main Screen. For details, see "[5.3.1.1 Confirming SDX Object Configuration](#)."
- When using a command, the mode of JRM for volumes can be checked using the JRM field of the volume information displayed with the `sdxinfo` command. For details, see "[D.6 sdxinfo - Display object configuration and status information](#)."

JRM for Slices

JRM for slices speeds up the resynchronization process when reattaching a detached slice to the volume. GDS records the changes made on the volume and slice in the memory while the slice is being detached. The resynchronization copy performed when the detached slice is reattached copies the updated portions only to realize high-speed resynchronization.

JRM for slices becomes effective when a slice is detached while the `jrm` attribute of the slices is on. However, if a system is stopped or if the slice is taken over by the `sdxslice -T` command while the slice is detached, just resynchronization is not conducted when the temporarily detached slice is attached again. Resynchronization is performed by copying the entire data, not only the updated portions.

Therefore, if you plan to shut down the system, or have a slice taken over, attaching the slice to the volume in advance is highly recommended.



See

- For details on the setting methods, see "[D.5 sdxslice - Slice operations](#)."
- For details on the changing methods, see "[D.7 sdxattr - Set objects attributes](#)."
- The mode of JRM for slices can be checked using the JRM field of the slice information displayed with the `sdxinfo` command with the `-e` long option. For details, see "[D.6 sdxinfo - Display object configuration and status information](#)."

JRM for Proxies

JRM for proxies speeds up the just resynchronization process when joining a parted proxy again to the master and when the master data is restored from the proxy. GDS records the changes made on the master and the proxy on the memory while the proxy is parted. The just resynchronization conducted when rejoining or restoring copies only the updated portions to realize high-speed synchronization.

JRM for proxies is enabled when the `pjrm` attribute of a proxy volume is set to "on" and the proxy volume is parted. However, if any node that is included in the scope of the class is stopped while the proxy is parted, just resynchronization is not put in operation. In other words, the entire data, not only the updated portions, is copied.

Therefore, if you plan to shut down the system, joining the proxy to the master in advance is highly recommended.

Reference to these matters is not necessary when you are using the copy function of a disk unit.



See

- For details on the setting methods, see "[5.3.2.2 Backup \(by Synchronization\)](#)" and "[D.14 sdxproxy - Proxy object operations](#)."
- For details on the changing methods, see "[D.7 sdxattr - Set objects attributes](#)."
- When using GDS Management View, the mode ("on" or "off") of JRM for proxies can be checked using the proxy volume information field in the Main Screen. For details, see "[5.3.1.2 Viewing Proxy Object Configurations](#)."
- When using a command, the mode of JRM for proxies can be checked using the PJRM field of the volume information displayed with the `sdxinfo` command with the `-e` long option. For details, see "[D.6 sdxinfo - Display object configuration and status information](#)."

A.2.15 Online Volume Expansion

- Volume Configuration Limitations

Online volume expansion is available for volumes in the following configurations.

- Single volume
- Mirror volume
 - Any mirroring multiplicity is supported.
 - Hierarchized groups are supported.

- Online Mirror Volume Expansion

For mirror volumes with mirroring multiplicity of two and higher, change the mirroring multiplicity to one, expand the volumes, and then execute synchronization copying for re-mirroring. See the outline of the operating procedures below. These operations are executable without stopping applications using the volumes.

1. Disconnect disks and lower level groups from the mirror group to change the mirroring multiplicity to one.
2. Expand the volume size with the `sdxvolume -S` command.
3. Reconnect the disks and lower level groups disconnected in step 1. with the mirror group.

If the mirror volumes are active, resynchronization copying is automatically performed after step 3. is done. If the mirror volumes are inactive, similar copying is automatically performed when they are started.

- Stripe Type Volume and Concatenation Type Volume Expansion

The capacity of stripe type volumes and concatenation type volumes cannot be expanded. To expand one of these types, back up data, recreate the volume, and then restore the data back to the volume. In configurations where a striping group or a concatenation group is connected to a mirror group (with any multiplicity), volume expansion applied using the striping or concatenation feature is possible.



See

.....
 For the methods of backing up and restoring, see "[Chapter 6 Backing Up and Restoring.](#)"

- Concatenation and Online Volume Expansion

Even if there is no sufficient continuous free space after the last block of a volume, by concatenating unused disks, online volume expansion will be available. Online volume expansion is available for volumes that meet all of the following conditions:

- Volumes belong to a mirror group;
- To the mirror group, one or more concatenation groups are connected;
- To each of the concatenation groups, one or more disks are connected.

To use this function, create volumes in configuration conforming to these conditions in advance. For example, if there is only one available disk, connect only the disk to a concatenation group, connect the concatenation group to a mirror group, and create volumes in the mirror group.

- Expansion of Areas Used by Applications

After volumes are expanded, applications such as file systems and databases need to recognize the expanded areas with methods specific to the applications. For example, after a volume that contains a ufs file system is expanded in the Solaris 9 operating System, the file system area must also be expanded with the `growfs(1M)` command.

If an application using a volume cannot recognize an expanded volume area, do not expand the volume. If such a volume is expanded, the application may no longer operate normally or volume data may be unavailable. A volume that contains the GFS Local File System or the GFS Shared File System cannot be expanded.

A.2.16 Swapping Physical Disks

For the procedures for swapping disks, see "[5.3.4 Disk Swap](#)" when using GDS Management View, or see "[D.8 sdxswap - Swap disk](#)" when using a command.

This sub-section describes important points about disk swap.

Physical Disk Size

Physical disk swap cannot be performed using a physical disk whose size is smaller than the original disk size.

Physical Disks That Cannot Be Swapped

Physical disk swap cannot be performed for a disk where the only valid slice (in ACTIVE or STOP status) within the volume exists.

For example, it is impossible to perform disk swap if:

- The volume is in a group to which only one disk is connected.
- The volume is on a single disk that is not connected to a group.

In these situations, it will be possible to perform disk swap by making any one of the following configuration changes.

- a. If the disk to be swapped is connected to a mirror group, add a new disk to the mirror group and complete volume synchronization copying normally.
- b. When performing disk swap for a single disk, add the single disk and another unused disk to a new mirror group, and complete volume synchronization copying.
- c. Remove the existing volume from the disk to be swapped.
Before removing the volume, back up volume data if necessary.

When data on a disk to be swapped is valid, for example when conducting preventive maintenance, a. or b. is recommended. Here, if the disk unit supports hot swap, disks can be swapped without stopping active applications.

Swapping Physical Disks in a Cluster System

If physical disks that are registered with the resource database were swapped, after physical disk restoration (or the `sdxswap-0` command) is completed, perform automatic resource registration regardless of the class types. This is required for changing VTOC volume names of original disks to those of replacing disks that are registered with the PRIMECLUSTER resource database. If automatic resource registration is not performed, a message indicating a disk unit connection error may be output or failover may occur when nodes are started.

Swapping Physical Disk Registered with Shadow Classes

When swapping disks of shadow disks registered with GDS Snapshot shadow classes, the shadow disks must be removed with the relevant GDS Snapshot command first. For details on GDS Snapshot commands, see "[Appendix D Command Reference](#)."

The subsequent operations vary depending on whether the disk units to be swapped are registered with disk classes.

- If the disk is registered with a disk class

After removing the shadow disk, perform disk swap in the domain managing the disk class using the relevant procedure.

An error in a disk may not cause failures on both the related SDX and shadow objects. Even if only either of the objects fails, shadow disk removal and disk swap using the relevant procedure are both necessary.

- If the disk is not registered with a disk class

In this situation, the disk to be swapped is a copy destination of the disk unit's copy function.

It is not necessary to perform GDS operations described in "[5.3.4 Disk Swap](#)" and "[D.8 sdxswap - Swap disk](#)." After removing the shadow disk, perform disk swap referring to the manual of the disk unit's copy function.

A.2.17 Object Operation When Using Proxy

If classes, groups, and slices are related to master volumes or proxy volumes, operations that change the class attributes, change the group configurations or attributes, or handle the slices cannot be performed. To perform such operations, cancel the relationship between the master and the proxy once.

For objects that are related to the master and the proxy:

- Master volumes can be started or stopped.
- Proxy volumes that are parted from the master can be started or stopped.
- New volumes can be created within a group other than the master group or the proxy group and existing volumes other than the master or the proxy can be removed.
- The hot spare function is available for groups with master volumes. Spare disks are not connected to groups with proxy volumes.

The following operations can be performed, but the operation will result in an error if copying is in process, or there is an EC session, a BCV pair, or an SRDF pair, between the master and the proxy.

- Attributes of master and proxy volumes can be changed with the `sdxattr -V` command.
- A disk connected to a group with master volumes and proxy volumes can be made exchangeable with the `sdxswap -O` command, and after swapping disks, the swapped disk can be returned to a usable state with the `sdxswap -I` command.
- Synchronization copying of the master volume or the proxy volume separated from the master volume can be started, cancelled, interrupted and resumed, and the parameters of the synchronization copying changed. However, synchronization copying between volumes or its parameters cannot be changed.
- Master volume can be restored with the `sdxfix -V` command.

For EC sessions, BCV pairs, and SRDF pairs between the master and the proxy, check the FUNC field displayed with the `sdxinfo` command. If the master and the proxy are in parted status, a copy session can be canceled with the `sdxproxy Cancel` command.

A.2.18 Using the Advanced Copy Function in a Proxy Configuration

In a proxy configuration, by working in cooperation with the Advanced Copy functions ((R)EC/(R)OPC) of Fujitsu Storage Systems ETERNUS, the copying between master and proxy can be performed without imposing loads on primary servers or a SAN. In this situation, disk array units carry out copying processes and the processes will continue running even if the server is rebooted.

When the Advanced Copy function is available, it is used for copying between the master and the proxy. However, in the following situations, a soft copy function (a copy function of a GDS driver operating on a server) is used.

- The Advanced Copy function is not used when:
 - Use of soft copy functions was specified explicitly using the `sdxproxy` command with the `-e softcopy` option
 - The copy destination volumes are in a mirroring configuration.
Note, however, that even if the master as the copy destination is in a mirroring configuration, OPC is available for copying from the proxy when:
 - Executing [Operation]:[Proxy Operation]:[Restore] in GDS Management View and selecting "No" to "Rejoin" in the [Restore Master] dialog box. For details see "[5.3.3 Restore](#)."
 - Executing the `sdxproxy Restore` command. For details see "[D.14 sdxproxy - Proxy object operations](#)."
 - A lower level group is connected to a group to which master volumes or proxy volumes belong.
 - The number of concurrent EC or OPC sessions has reached the upper limit defined by the disk array unit.



See

.....
The number of allowed concurrent sessions is either the upper limit within one physical disk (LU) or the upper limit within one disk array unit. For details, see the handbook of the relevant disk array.
.....

- Multiple proxy volumes are related to a master volume and the number of proxy volumes with EC sessions has reached the upper limit (16 volumes).

- Disks that constitute the master or the proxy were registered with a class before installing the Advanced Copy function on the disk array unit.

In these situations, operations based on OPC functions are impossible. For such operations see "A.2.19 Instant Snapshot by OPC."

EC is used for synchronization copying from a master to a proxy after joining or rejoining them, copying for maintaining synchronization, and recording the portion updated while a master and a proxy are parted.

OPC is used for synchronization copying, instant snapshot processing, and restoring a master using the proxy data.

If any EC sessions exist between a master and a proxy, OPC cannot be initiated between the master and another proxy.

When both EC and OPC are available as with ETERNUS6000 series, EC has precedence over OPC. Once an EC session is stopped, you cannot use the EC function afterwards. EC sessions are stopped in the following situations.

- EC sessions are stopped when:
 - Executing [Operation]:[Proxy Operation]:[Part] in GDS Management View and selecting "Yes" to "Instant Snapshot" in the [Part Proxy] dialog box
 - Canceling the EC sessions with the `sdxproxy Cancel` command
 - Joining a master and a proxy and using a soft copy function with the `sdxproxy Join -e softcopy` command
 - Creating instant snapshots with the `sdxproxy Part -e instant` command
 - Rejoining a master and a proxy and using a soft copy function with the `sdxproxy Rejoin -e softcopy` command

To make the EC function available after conducting these operations, break the relationship between the master and the proxy once and rejoin them.

To check the modes of the copying in execution, use either:

- The [Copy Type] field of the slice information field in GDS Management View
- The CPTYPE field displayed with the `sdxinfo` command

Additionally, the types and the statuses of sessions between the master and the proxy can be viewed in the FUNC field and the CPSTAT field displayed with the `sdxinfo` command.



Note

Advanced Copy Control

When the Advanced Copy function is available, executing the `sdxproxy` command directs GDS to control the Advanced Copy on the master and the proxy. Do not use any other methods other than the `sdxproxy` command to apply such control on a master and a proxy.

A.2.19 Instant Snapshot by OPC

The following functions are based on OPC functions of Fujitsu Storage Systems ETERNUS.

- Instant snapshot by OPC
 - [Operation]:[Proxy Operation]:[Update] in GDS Management View
 - "Yes" to "Instant Snapshot" in the [Part Proxy] dialog box invoked through [Operation]:[Proxy Operation]:[Part] in GDS Management View
 - `sdxproxy Update` command
 - `sdxproxy Part -e instant` command
- Master restoration by OPC
 - "No" to "Rejoin" in the [Restore Master] dialog box invoked through [Operation]:[Proxy Operation]:[Restore] in GDS Management View
 - `sdxproxy Restore` command

These functions are unavailable with disk units other than the following types.

- ETERNUS DX60 / DX80 / DX90 (*1)
- ETERNUS DX80 S2 / DX90 S2 (*1)
- ETERNUS DX410 / DX440 (*1)
- ETERNUS DX410 S2 / DX440 S2 (*1)
- ETERNUS DX8100 / DX8400 / DX8700 (*1)
- ETERNUS8000
- ETERNUS6000
- ETERNUS4000 (Model 100 or higher) (*2)
- ETERNUS3000 (Model 100 or higher)
- ETERNUS2000 (Model 100 or higher) (*3)
- ETERNUS GR720
- ETERNUS GR730
- ETERNUS GR740

*1) GDS patch 914423-10 or later and GDS Snapshot patch 914457-03 or later are required.

*2) GDS patch 914423-01 or later is required.

*3) GDS patch 914423-05 or later is required.

These functions are also unavailable under such conditions that prevent the use of Advanced Copy functions as described in "The Advanced Copy function is not used when:" in ["A.2.18 Using the Advanced Copy Function in a Proxy Configuration."](#)



Note

Rebooting a Server While OPC Running

Even if a server is rebooted while the copying between master and proxy is being processed by OPC, the OPC copying will continue running. However, if OPC sessions are not present when the server is up again, GDS will assume that the copying failed and copy destination volumes will be INVALID. For recovering this status, see ["\(4\) Master volume is in INVALID status."](#) and ["\(5\) Proxy volume is in INVALID status."](#) in ["F.1.3 Volume Status Abnormality."](#)

A.2.20 To Use EMC Symmetrix

For using an EMC's Symmetrix storage unit, EMC's PowerPath is essential. GDS is capable of managing EMC Symmetrix emcpower devices and native devices that compose the emcpower devices.

Typically, emcpower devices are the targets of management. However, in a system where native devices are already the management targets, native devices can still be used.

The following devices are exceptions to devices that can be managed by GDS.

- BCV (Business Continuance Volume) devices
- SRDF target (R2) devices
- GateKeeper devices
- CKD (Count Key Data) devices
- VCMDB (Volume Configuration Management Data Base) used by EMC's SAN management software (Volume Logix, ESN Manager, SAN Manager and so on).

After completing the configuration of these devices and EMC software, follow the procedure below and describe a list of devices excluded from disk management by GDS in the `/etc/opt/FJSVsdx/lib/exdevtab` file (referred to as the Excluded Device List). The Excluded Device List must include all disks that cannot be managed by GDS in addition to the devices above.

[Procedure]

1. The `syminq` command provided by `SYMCLI` is available for checking respective BCV, R2, GateKeeper, and CKD devices. Execute the `syminq` command, and describe all devices indicated as BCV, R2, GK, and CKD(`cCtTdD`, `emcpowerN`) in the Excluded Device List.
2. The `syminq` command is unavailable for checking VCMDB devices. When using EMC's SAN management software (Volume Logix, ESN Manager, SAN Manager and so on), ask your EMC engineer or systems administrator who configured that SAN management software about the names of VCMDB devices and describe them in the Excluded Device List.
3. For managing native devices, describe all `emcpowerN` and additionally native devices (`cCtTdD`) excluded from management in the Excluded Device List.

```
emcpower0 (Nontarget)
c2t0d0 (Target)
c3t0d0 (Nontarget)
```

4. For managing `emcpower` devices, describe all native devices (`cCtTdD`) in the Excluded Device List.

```
emcpower0 (Target)
c2t0d0 (Nontarget)
c3t0d0 (Nontarget)
```

5. In addition to the devices specified in steps 1. through to 3., describe any other devices to be excluded from GDS management in the Excluded Device List.

You are recommended to suffix such tags as "PP", "BCV", "R2", "GK", "CKD" and "VCMDB" to the device names for Excluded Device List management efficiency. A device name and a tag must be separated by one or more spaces.

The Excluded Device List should appear as follows.

- For managing native devices

```
# cat /etc/opt/FJSVsdx/lib/exdevtab
# device(full path) type
/dev/rdisk/emcpower0c PP

# /dev/rdisk/c2t0d0s2 PP
/dev/rdisk/c3t0d0s2 PP
...
/dev/rdisk/emcpower7c BCV
/dev/rdisk/c2t0d7s2 BCV
/dev/rdisk/c3t0d7s2 BCV

/dev/rdisk/emcpower8c BCV
/dev/rdisk/c2t0d8s2 BCV
/dev/rdisk/c3t0d8s2 BCV

/dev/rdisk/emcpower9c GK
/dev/rdisk/c2t0d9s2 GK
/dev/rdisk/c3t0d9s2 GK

/dev/rdisk/emcpower10c GK
/dev/rdisk/c2t0d10s2 GK
/dev/rdisk/c3t0d10s2 GK

/dev/rdisk/emcpower11c CKD
/dev/rdisk/c2t0d11s2 CKD
/dev/rdisk/c3t0d11s2 CKD

/dev/rdisk/emcpower12c R2
/dev/rdisk/c2t0d12s2 R2
/dev/rdisk/c3t0d12s2 R2
#
```

- For managing emcpower devices

```
# cat /etc/opt/FJSVsdx/lib/exdevtab
# device(full path) type
/dev/rdsd/emcpower7c BCV
/dev/rdsd/emcpower8c BCV
/dev/rdsd/emcpower9c GK
/dev/rdsd/emcpower10c GK
/dev/rdsd/emcpower11c CKD
/dev/rdsd/emcpower12c R2
/dev/rdsd/c2t0d0s2 PP
/dev/rdsd/c3t0d0s2 PP
...
/dev/rdsd/c2t0d7s2 BCV
/dev/rdsd/c3t0d7s2 BCV
/dev/rdsd/c2t0d8s2 BCV
/dev/rdsd/c3t0d8s2 BCV
/dev/rdsd/c2t0d9s2 GK
/dev/rdsd/c3t0d9s2 GK
/dev/rdsd/c2t0d10s2 GK
/dev/rdsd/c3t0d10s2 GK
/dev/rdsd/c2t0d11s2 CKD
/dev/rdsd/c3t0d11s2 CKD
/dev/rdsd/c2t0d12s2 R2
/dev/rdsd/c3t0d12s2 R2
#
```

Information

sdsexdevtab and exdevtab.sh

The script samples "/etc/opt/FJSVsdx/bin/sdsexdevtab.sample" and "exdevtab.sample.sh" for simply creating the Excluded Device List "/etc/opt/FJSVsdx/lib/exdevtab" are provided.

To use one of these scripts, open the script with an editor and modify the following two parameters (the syminq command and powermt command paths) according to the execution environment.

```
SYMINQ=/usr/symcli/bin/syminq
POWERMT=/etc/powermt
```

For managing native devices, execute sdsexdevtab. By executing sdsexdevtab, emcpower devices, native devices excluded from management, and all BCV, R2, GateKeeper and CKD devices are included in the Excluded Device List. VC MDB devices will not be included in the list. If necessary, edit sdsexdevtab in advance or add any other disks to be excluded from GDS management to the list using the above steps 1. through to 5.

```
# cp /etc/opt/FJSVsdx/bin/sdsexdevtab.sample
  /etc/opt/FJSVsdx/bin/sdsexdevtab
# chmod u+x /etc/opt/FJSVsdx/bin/sdsexdevtab

# /etc/opt/FJSVsdx/bin/sdsexdevtab -M > /etc/opt/FJSVsdx/lib/exdevtab
# vi /etc/opt/FJSVsdx/lib/exdevtab
...
```

For managing emcpower devices, execute exdevtab.sh. By executing exdevtab.sh, the native devices of emcpower devices and all BCV, GateKeeper, and CKD devices are included in the Excluded Device List. R2 devices and VC MDB devices will not be included in the list. If necessary, edit exdevtab.sh in advance or add any other disks to be excluded from GDS management to the list using the above steps

1. through to 5.

```
# cp /etc/opt/FJSVsdX/bin/exdevtab.sample.sh
  /etc/opt/FJSVsdX/bin/exdevtab.sh
# chmod u+x /etc/opt/FJSVsdX/bin/exdevtab.sh

# /etc/opt/FJSVsdX/bin/exdevtab.sh
# vi /etc/opt/FJSVsdX/lib/exdevtab
...

```

Note

Using EMC TimeFinder or EMC SRDF in a Proxy Configuration

For using EMC TimeFinder or EMC SRDF in a proxy configuration, do not describe BCV and R2 devices connected with proxy groups, and among their native devices, devices to be included in GDS management in the Excluded Device List. For details, see "[A.2.21 Using EMC TimeFinder or EMC SRDF in a Proxy Configuration](#)."

Note

For PRIMECLUSTER Systems

- In a PRIMECLUSTER system, create Excluded Device Lists on all nodes that constitute the cluster.
- Devices that cannot be managed with GDS are also nontargets of automatic resource registration (automatic configuration), and before performing automatic resource registration it is necessary to describe these devices also in the Excluded Device List for automatic resource registration. For details, see the "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide."
- For registering BCV devices as resources in automatic resource registration, it is necessary to split or cancel BCV pairs before executing the registration.
- For registering R2 devices of SRDF pairs as resources in automatic resource registration, it is necessary to split the SRDF pairs before executing the registration.

Note

Kernel Parameters

- When an emcpower device is managed by GDS using EMC PowerPath 1.5.0.x, add the following line in the /etc/system file if the file does not include it.

```
forceload: drv/emcpower
```

- When an emcpower device is managed by GDS using EMC PowerPath 2.0.2 or after, add the following lines in the /etc/system file if the file does not include them.

```
set lwp_default_stksize=0x6000
set rpcmod:svc_default_stksize=0x6000
```

- When an emcpower device is managed by GDS using EMC PowerPath 2.1.1 or after, add the following lines in the /etc/system file if the file does not include them.

```
set lwp_default_stksize=0x6000
set rpcmod:svc_default_stksize=0x6000
set emcp:bPowerIgnoreFEXCL=1
```

- When using EMC PowerPath 3.x, no kernel parameter setting is required.

Information

Devices Described in exdevtab

A disk device described in the exdevtab file cannot be registered with a class. If an attempt to register the disk device included in the exdevtab file is made, an error message "ERROR: *physical disk name*: no such device" is output.

A.2.21 Using EMC TimeFinder or EMC SRDF in a Proxy Configuration

In a proxy configuration, by using EMC's TimeFinder and SRDF, the copy functions of EMC's Symmetrix storage units can be used and the synchronization copying between master and proxy can be performed without imposing loads on primary servers or a SAN. In this situation, the storage units carry out copying processes and the processes will continue running even if the server is rebooted.

TimeFinder and SRDF are available to copy data of an entire physical disk to another physical disk, but they are not available to copy data of a disk slice to the other disk area. Therefore, to utilize TimeFinder or SRDF with a proxy configuration, a pair of groups instead of a pair of volumes must be related as a master and a proxy. If a proxy volume instead of a proxy group is specified as a target of the operation to part, rejoin, or restore, the operation results in an error, while using TimeFinder or SRDF.

Configuring to meet the following requirements should be done before joining a master group and proxy group.

- To utilize TimeFinder:
 1. One of the disks connected to the master group must be the standard device that has been registered with the device group on all of the nodes within the class scope.
 2. A disk connected to the proxy group must be a BCV device that is associated with the same device group as described in 1. on all of the nodes within the class scope.
 3. If the standard device in 1. and the BCV device in 2. are established, the BCV pair must be canceled in advance.
- To utilize SRDF:
 1. One of the disks connected to the master group must be the source (R1) device that has been registered with the device group on all of the nodes within the class scope.
 2. A disk connected to the proxy group must be the target (R2) device that is paired with the source (R1) device as above on all of the nodes within the class scope.
 3. The source (R1) device in 1. and the target (R2) device in 2. must be in split status.

Moreover, ensure the following points about management of BCV devices, source (R1) devices and target (R2) devices, which are used with a proxy configuration.

- GDS configuration databases cannot be stored in BCV devices and target (R2) devices since the devices are overwritten by data in copy source disks. Therefore, GDS does not regard BCV devices and target (R2) devices as "disks that can be accessed normally" described in "[Explanation]" of "(1) Class becomes closed status during operation." in "F.1.4 Class Status Abnormality."
- In the Excluded Device List, do not describe BCV devices and target (R2) devices connected with proxy groups, and among their native devices, devices to be excluded from GDS management. For details on the Excluded Device List, see "A.2.20 To Use EMC Symmetrix."
- When the target (R2) device of the SRDF pair is not described in the Excluded Device List, such an SRDF pair must be split prior to conducting Automatic resource registration. For details on the procedure for Automatic resource registration, see "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide."
- When the target (R2) device of the SRDF pair is registered with a GDS class, such an SRDF pair must be split in advance.

- The BCV, source (R1), and target (R2) devices used in a proxy configuration should not be operated using the SYMCLI commands.
- If the master and proxy are parted forcibly while copying by TimeFinder or SRDF is in process, a message informing device abnormality is submitted to EMC's customer support center.

When a disk unit's copy function is available, synchronization copying from a master to a proxy is performed using that function. However, the soft copy function (copy function of a GDS driver running on the server) will be used in the following cases.

- The soft copy function is specified to be used.
- A pair of volumes instead of a pair of groups is related as a master and a proxy.
- The configuration of the proxy group to which the data are copied is a mirroring configuration.
- The master and the proxy belong to a root class.
- A proxy volume with a different physical slice attribute from the master volume is created into a proxy group.
- A lower level group is connected to the master group or the proxy group.
- A disk of the same size as the disk to which the data are copied is not connected to the master group from which the data are copied.
- Disks that constitute the master or the proxy were registered with a class before installing GDS Snapshot.

TimeFinder and SRDF are used for synchronization copying from a master to a proxy after joining or rejoining them, copying for maintaining synchronization, and recording the portion updated while a master and a proxy are parted.

If both TimeFinder and SRDF are available, TimeFinder has precedence over SRDF.

Once the BCV pair or SRDF pair is cancelled, the TimeFinder function or SRDF function is not used. BCV pairs and SRDF pairs are cancelled when:

- Executing [Operation]:[Proxy Operation]:[Restore] in GDS Management View
- Using the `sdxproxy Cancel` command to cancel (break) the BCV or SRDF pair
- Using a soft copy function when joining a master and a proxy with the `sdxproxy Join -e softcopy` command
- Using a soft copy function when rejoining a master and a proxy with the `sdxproxy Rejoin -e softcopy` command
- Restoring data from a proxy back to a master with the `sdxproxy RejoinRestore` command

To make the TimeFinder and SRDF functions available after conducting these operations, break the relationship between the master and the proxy once, remove proxy volumes, and then rejoin the master and the proxy.

To check the modes of the copying in execution, use either:

- The [Copy Type] field of the slice information field in GDS Management View
- The CPTYPE field displayed with the `sdxinfo` command

Additionally, the types and the statuses of BCV pairs and SRDF pairs between the master and the proxy can be viewed in the FUNC field and the CPSTAT field displayed with the `sdxinfo` command.

A.2.22 Ensuring Consistency of Snapshot Data

If snapshots are created while an application is accessing the volume, the snapshots may result from incomplete volume data and the data consistency may not be ensured.

To ensure the consistency of your snapshot data, you must stop the application that is accessing the volume in advance. After creating the snapshot, start the application again.

For example, when using the volume (master volume) as a file system such as Global File Services or ufs, unmount the volume with the `umount(1M)` command before creating a snapshot, and mount it with the `mount(1M)` command afterwards. This way, you can ensure the consistency of the snapshot data.

To create a snapshot while running the application, the file system or database system you are using to manage the data must be able to ensure data integrity.

For an example, see ["6.4 Online Backup and Instant Restore through Proxy Volume."](#)

A.2.23 Data Consistency at the time of Simultaneous Access

When the same block within a volume is accessed simultaneously from multiple nodes, data consistency is maintained by access exclusion control performed by the application that accesses the shared disk simultaneously.

A.2.24 Volume Access Mode

There are two types of volume access modes: "Default Access Mode" which is set by default as an access mode attribute and "Current Access Mode" which is set for a volume that is activated.

"Current Access Mode" is valid only while the volume is activated and will become invalid when the volume is stopped. When the volume is restarted, it will start in "Default Access Mode," except for when the access mode is specified at the time of restart.

For example, if you wish to normally use the volume in the read and write mode, and temporarily switch to the read only mode, set the access mode attribute to "rw", and use the `sdxvolume -N` command specifying the `-e mode=ro` option to activate the volume in the read only mode temporarily.

The "default access mode" for a shadow volume is ro (read only) and it cannot be changed, but the "current access mode" can be set to either ro (read only) or rw (read and write).



See

- For the method for setting the "default access modes" (access mode attribute values) of logical volumes, see ["D.7 sdxattr - Set objects attributes."](#)
- For the method for setting the "current access modes" of logical volumes, see ["D.4 sdxvolume - Volume operations."](#)
- For the method for setting the "current access mode" for shadow volumes, see ["D.17 sdxshadowvolume - Shadow volume operations."](#)
- The "default access modes" (access mode attribute values) and the "current access modes" of logical volumes and shadow volumes can be checked using the MODE field and the CMODE field displayed with the `sdxinfo -V` command respectively. For details, see ["D.6 sdxinfo - Display object configuration and status information."](#)

A.2.25 Operation in Cluster System

Understand and pay attention to the following points when changing the configuration in a cluster system.

- Before registering disks with a class, perform automatic resource registration and register disk units with the PRIMECLUSTER resource database.



See

For details on the automatic resource registration, see "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide."

- Disks that have not been registered with the resource database yet cannot be registered with a shared class.
- Disks that have not been registered with the resource database yet can be registered with a root class or a local class.
- When expanding the scope of a local class to which disks not registered with the resource database belong into a shared class, perform automatic resource registration in advance, and register all disks that belong to the local class with the resource database.
- For disks to be registered with a shadow class, disk resource creation is not required.
- Do not register certain disks with the resource database in multiple cluster domains. Before executing automatic resource registration, describe those disks in the Excluded Device List.
- After automatic resource registration, to register MPHD disks or MPLB disks with a class from GDS Management View, execute [Update Physical Disk Information] in the [Operation] menu before executing the registration.

- To perform object operations in a cluster system, enable cluster control. If cluster control is "off", it is impossible to perform shared object operations. Additionally, root and local object operations may cause errors or inconsistency.

A.2.26 Changing Over from Single Nodes to a Cluster System

The following two methods are available for installing the cluster control facility on one or more single nodes where classes exist and changing over to a cluster system. Normally, the method "a" is used, but the method "b" can be used to skip backup and restore.

a. Deleting classes and re-creating classes after changing over to a cluster system

Back up volume data if necessary, and delete the root class and local classes. After installing and configuring the cluster control facility, re-create the classes and volumes, and restore volume data as needed.

b. Changing over to a cluster system without deleting classes

Change over from single nodes to a cluster system without deleting classes. However, you must understand and pay attention to the following points.

Important Point 1

A local class created on a single node cannot be used directly in a cluster system. When the cluster control facility is activated, the following error message is output to the system log and the GDS daemon log file, and the operation of the local class becomes unavailable.

ERROR: *class*: cannot operate in cluster environment, created when cluster control facility not ready

For details on resolution, see "(1) The error message "ERROR: class: cannot operate in cluster environment, ..." is output, and the operation cannot be conducted on the class class." in "F.1.9 Cluster System Related Error."

Important Point 2

Expanding the class scope after changing over from multiple single nodes to a cluster system, may output the following messages.

ERROR: *class*: class names must be unique within a domain

This error occurs when the name of a class created on a single node is the duplicate name of a class on another node. If this error occurs, rename either of the classes, and expand the class scope.

ERROR: *class*: volume minor numbers must be unique within a domain

This error occurs when the minor number of a volume created on a single node is the duplicate number of a volume on another node. If this error occurs, re-create either of the volumes, and expand the class scope.

The minor number of a volume can be viewed in the following manner.

```
# cd /dev/sfdsk/class/dsk
# ls -l
brw----- 1 root    root      253, 33 May  6 09:00 volume1
                ^^
```

Additionally, this error may occur when any lower level group exists in a class created on a single node or a class on another node. In this event, duplicate minor numbers cannot be checked with the method as shown above. Re-create all volumes and lower level groups in the class created on the single node and then expand the class scope.

A.2.27 Disk Switch

The disk switch function is available only when using an application that controls the disk switch function. Unless the application's manual instructs to creation of switch groups and switch volumes, do not create switch groups.

A.2.28 Shadow Volume

Rebooting a Node

The configuration information of a shadow volume is not saved on the private slice but managed in the memory. For this reason, the shadow volume configuration definitions are cleared when the node on which the shadow volume is defined is rebooted. However, the device special file remains. If such a device special file is left not deleted, issues as described below may occur.

Before intentional shutdowns, it is recommended to remove shadow volumes. If a shadow volume is removed with the `sdxshadowvolume -R` command, the device special file is also deleted. For details on the `sdxshadowvolume -R` command, see "[D.17 sdxshadowvolume - Shadow volume operations.](#)"

When a node is shut down leaving the relevant shadow volume not removed, or if a node on which a shadow volume is defined is rebooted unexpectedly because of an event such as a panic and a power cutoff, the device special file for the shadow volume must be deleted in the following procedure.

[Procedure]

1. Check the system for existing classes.

In the following example, there are `RootClass`, `Class1`, and `Class2`.

```
# sdxinfo -C
OBJ      NAME      TYPE      SCOPE      SPARE
-----
class    RootClass  root      (local)    0
class    Class1     local     node1      0
class    Class2     shared    node1:node2 0
```

2. Find the directories containing the device special files of classes.

In the following example, `RootClass`, `Class1`, and `Class2` are the directories for the device special files of those existing classes, and `_adm` and `_diag` are the special files used by GDS. `Class3`, other than those directories, is the directory for the device special file of the extinct shadow class.

```
# cd /dev/sfdsk
# ls -F
Class1/      Class2/      Class3/      RootClass/   _adm@
_diag@
```

3. Delete the directory for the device special file of the extinct shadow class.

```
# rm -r Class3
```

Even if the device special file of an extinct shadow volume remains, no problem will arise if a shadow volume in the same configuration, of the same class name, and with the same volume name is re-created.

Otherwise, the following issues will occur. If a logical volume or a shadow volume is created in the situation that the device special file, `/dev/sfdsk/Shadow Class Name[r]dsk/Shadow Volume Name`, of an extinct shadow volume remains, the minor number of the created volume may become the same as the minor number of `/dev/sfdsk/Shadow Class Name[r]dsk/Shadow Volume Name`. In this situation, if `/dev/sfdsk/Shadow Class Name[r]dsk/Shadow Volume Name` is accessed without recognition of extinction of the shadow volume, the newly created volume is accessed, and it can cause an application error and corruption of data on the newly created volume.

Accessing a Shadow Volume

Shadow volumes and the corresponding logical volumes are managed independently. For example, the change of the slice status in one volume is not updated in the slice status in the other volume. For this reason, you must note the following operational particulars when using shadow volumes.

Synchronization of Shadow Volumes

When a shadow volume is created in another domain (domain beta) for the disk area managed as a mirror volume in a certain domain (domain alpha), the mirror volume in domain alpha and the shadow volume in domain beta cannot be accessed simultaneously. If they are accessed simultaneously, the following issues arise.

- If an I/O error occurs in the slice configuring the mirror volume in domain alpha, that slice becomes INVALID and is detached from the mirror volume. However, GDS in domain beta does not detect this I/O error, and consequently the shadow slice is not made INVALID and is not detached from the shadow volume. Here, synchronization of the shadow volume is not ensured.
- Likewise, if an I/O error occurs in the shadow slice in domain beta, the slice in the corresponding mirror volume in domain alpha is not made INVALID and is not detached from the mirror volume. Here, synchronization of the mirror volume is not ensured. If an I/O error occurs on the shadow slice, working around, such as swapping the disks and resynchronization copying of the mirror volume, is required in domain alpha.

These particulars apply when the disk area for a mirror volume and a shadow volume are identical. A mirror volume and a shadow volume corresponding to a replica of the mirror volume (a temporarily detached slice, a proxy volume or a copy destination disk area for a disk unit's copy function) can be accessed simultaneously.

Just Resynchronization Mechanism (JRM) for Volumes

When a shadow volume is created in another domain (domain beta) for the disk area managed as a mirror volume in a certain domain (domain alpha) and accessed, the following must be set up for the mirror volume in domain alpha.

- Mirror volumes must be inactivated to prevent access to the mirror volume corresponding to the shadow volume.
- JRM for volumes must be enabled ("on") for the mirror volume corresponding to the shadow volume.

These settings are necessary for the following reasons.

If a node in domain alpha panics and resynchronization copying is conducted on the mirror volume in domain alpha while the shadow volume is accessed in domain beta, synchronization between the shadow volume and the mirror volume is no longer ensured. Though the settings as above, resynchronization copying is no longer conducted on the mirror volume in domain alpha even if a node in domain alpha panics.

The settings as above are necessary only for a mirror volume created for the disk area identical to the shadow volume's disk area. When a shadow volume corresponding to a replica of a mirror volume (a temporarily detached slice, a proxy volume or a copy destination disk area for a disk unit's copy function) is created, these settings are not necessary for the copy source mirror volume.



Information

Resynchronization Copying after Panic

Resynchronization copying is not conducted after panic when JRM for volumes is enabled ("on") and that volume is not written in. Resynchronization copying occurs after panic when JRM for volumes is disabled ("off") and that volume is active.

Just Resynchronization Mechanism (JRM) for Slices

When a slice is temporarily detached from a mirror volume in a certain domain (domain alpha) and data is written from a shadow volume in another domain (domain beta) to the area of this volume or slice, JRM for slices must be disabled ("off") prior to reattaching the slice.

If JRM for slices is enabled ("on"), the following issue arises.

When JRM for slices is enabled ("on"), only the difference between the volume and the slice is copied by reattaching the slice. The difference information for the volume and the slice is managed by JRM for slices in domain alpha. However, JRM for slices in domain alpha does not recognize write events from domain beta, and the difference resulting from data being written from domain beta are not updated in the difference information. The difference resulting from write events from domain beta, therefore, are not copied when the slice is reattached while JRM for slices is "on" in domain alpha. As a result, synchronization of the volume is no longer ensured.

Just Resynchronization Mechanism (JRM) for Proxies

If a proxy volume is parted from the master in a certain domain (domain alpha) and data is written from a shadow volume in another domain (domain beta) to the area of this master or proxy, JRM for proxies must be disabled ("off") prior to rejoining the proxy. In addition, JRM for proxies must be disabled ("off") prior to restoring the master using the proxy.

If JRM for proxies is enabled ("on"), the following issues arise.

When JRM for proxies is enabled ("on"), only the difference between the master and the proxy is copied by rejoining or restoring. The difference information for the master and the proxy is managed by JRM for slices in domain alpha. However, JRM for proxies in domain alpha does not recognize write events from domain beta, and the difference resulting from data being written from domain beta are not updated in the difference information. The difference resulting from write events from domain beta, therefore, are not copied when the proxy is rejoined or the master is restored while JRM for proxies is "on" in domain alpha. As a result, synchronization between the master and the proxy is no longer ensured.

When one of a disk unit's copy function (EC, REC, TimeFinder, and SRDF) with a resynchronization feature based on equivalent copy capability is used for master-to-proxy copy processes, data written from domain beta is also updated in the difference information managed by these disk unit's copy functions. Under these circumstances, JRM for proxies do not have to be disabled ("off") prior to rejoining. Note, however, that JRM for proxies must be disabled ("off") prior to restoring since necessity of synchronization copying is determined based on the difference information managed by JRM for proxies. To ensure the data integrity, it is recommended to disable JRM for proxies prior to rejoining even when a disk unit's copy function is used.



Information

A Copy Function of a Disk Unit with Resynchronization Feature Based on Equivalent Copy

When just resynchronization from a master to a proxy is conducted with one of a disk unit's copy functions (EC, REC, TimeFinder, SRDF) with a resynchronization feature based on equivalent copy capability, this feature is used regardless of whether JRM for proxies is "on" or "off."

Writing into Shadow Volumes

Data may be written to a shadow volume even if the operation for writing is not especially intended. For example, executing mount(1M) (excluding when using the -o ro option), fsck(1M), newfs(1M), sfxnewfs(1M), or sfxadm(1M) results in the write operation.

When a proxy is rejoined, a master is restored, or a slice is reattached once a shadow volume is created, it is recommended to disable the just resynchronization mechanism mode (JRM) regardless of whether or not data is written into the shadow volume in order to ensure the data integrity.

A.2.29 Configuration Tables

Do not use editors such as vi(1) and sed(1) to edit configuration tables created with the sdxconfig Backup command or those saved in configuration files. To edit configuration tables, use the sdxconfig Convert command.

A.2.30 GDS Management View

Physical Disk Recognition

When any operation that changes the physical disk configuration, such as addition or deletion of disk units, is conducted during system operation, update physical disk information with new information. Execute [Update Physical Disk Information] on the [Operation] menu when:

- The power of the disk unit was turned on after the system was booted.
- The disk unit became unavailable for some kind of problem during system operation, but was recovered without system reboot.
- The configuration of devices was changed.

Object Status Monitoring

Objects indicated as failed in GDS Management View are only those in which GDS detected errors. Even if a hardware error occurs in a disk unit, the disk unit status is indicated as normal until the disk unit is accessed and the error is detected.

A.2.31 Swapping Internal Disks (FC-AL)

If FC-AL internal disks are swapped, the device numbers of the replacement disks may change. For this reason, FC-AL internal disks cannot be swapped through the procedures for physical disk swapping with [Swap Physical Disk] and [Restore Physical Disk] in the GDS Management View or with the `sdxswap` command. For FC-AL internal disk swapping, remove the relevant disk from the class once and register the replacement with the class again.

The following example shows the procedures for swapping an FC-AL internal disk connected to a mirror group for another. In this example, the swapped disk is a mirrored disk and removal of volumes from the disk prior to swapping is unnecessary. If the swapped disk is not a mirrored disk, such as a single disk and a disk whose highest level group is not a mirror group, volumes must be removed from the disk in advance. If volume removal is required, back up volume data according to need, remove the volumes, remove the disk from the class, register the replacement with the class, then re-create volumes and restore the data.

1. Check GDS I/O error messages output in the `/var/adm/messages` file for the device name (devname) of the failed disk.

To search the GDS I/O error messages, use the string "sfdsk" as a key. In the following example, the device name is `ssd@w210000203716ee10`.

```
WARNING: sfdsk: write error on mirror slice:
  volume info:      devno(maj,min)=0x03f4002e(253,46)
  device info:      devno(maj,min)=0x01d8001e(118,30)
                    devname=/sbus@2,0/SUNW,socal@d,10000/sf@1,0/
                    ssd@w210000203716ee10,0 (ssd3)
  error info:       blk in disk=17308981, blk in slice=5232448
                    length=8192, resid=8192, errno=5
```

2. Check the physical disk name of the swapped disk (failed disk).

For the argument of the `grep` command, specify the device name confirmed in step 1.

In the following example, the physical disk name is `c2t4d0`.

```
# ls -l /dev/rdisk | grep ssd@w210000203716ee10
lrwxrwxrwx 1 root root      78 Feb  3 18:49 c2t4d0s0 -> ../../
devices/sbus@2,0/SUNW,socal@d,10000/sf@1,0/
ssd@w210000203716ee10,0:a,raw
lrwxrwxrwx 1 root root      78 Feb  3 18:49 c2t4d0s1 -> ../../
devices/sbus@2,0/SUNW,socal@d,10000/sf@1,0/
ssd@w210000203716ee10,0:b,raw
lrwxrwxrwx 1 root root      78 Feb  3 18:49 c2t4d0s2 -> ../../
devices/sbus@2,0/SUNW,socal@d,10000/sf@1,0/
ssd@w210000203716ee10,0:c,raw
lrwxrwxrwx 1 root root      78 Feb  3 18:49 c2t4d0s3 -> ../../
devices/sbus@2,0/SUNW,socal@d,10000/sf@1,0/
ssd@w210000203716ee10,0:d,raw
lrwxrwxrwx 1 root root      78 Feb  3 18:49 c2t4d0s4 -> ../../
devices/sbus@2,0/SUNW,socal@d,10000/sf@1,0/
ssd@w210000203716ee10,0:e,raw
lrwxrwxrwx 1 root root      78 Feb  3 18:49 c2t4d0s5 -> ../../
devices/sbus@2,0/SUNW,socal@d,10000/sf@1,0/
ssd@w210000203716ee10,0:f,raw
lrwxrwxrwx 1 root root      78 Feb  3 18:49 c2t4d0s6 -> ../../
devices/sbus@2,0/SUNW,socal@d,10000/sf@1,0/
ssd@w210000203716ee10,0:g,raw
lrwxrwxrwx 1 root root      78 Feb  3 18:49 c2t4d0s7 -> ../../
devices/sbus@2,0/SUNW,socal@d,10000/sf@1,0/
ssd@w210000203716ee10,0:h,raw
```

3. Check the disk name, class name, and group name of the swapped disk.

For these names, execute the `sdxinfo -D` command and check the NAME, CLASS, GROUP fields of the row that shows the physical disk name confirmed in step 2. in the DEVNAM field.

In the following example, for physical disk `c2t4d0`, the disk name is `Disk2`, the class name is `Class1`, and the group name is `Group1`.

```
# sdxinfo -D
OBJ      NAME      TYPE      CLASS     GROUP     DEVNAM     DEVBLKS
DEVCONNECT STATUS
-----
-----
disk     Disk1     mirror   Class1    Group1    c0t0d0     17682084
Node1                   ENABLE
disk     Disk2     mirror   Class1    Group1    c2t4d0     17682084
Node1                   ENABLE
```

4. Disconnect the swapped disk from the group.

With the `-c`, `-g`, and `-d` options of the `sdxdisk -D` command, specify the class name, the group name, and the disk name confirmed in step 3.

```
# sdxdisk -D -c Class1 -g Group1 -d Disk2
```

5. Check if the swapped disk was disconnected from the group.

For checking, execute the `sdxinfo -D` command and check if the asterisk (*) is displayed in the GROUP field of the row that shows the physical disk name of the disconnected disk in the DEVNAM field.

```
# sdxinfo -D
OBJ      NAME      TYPE      CLASS     GROUP     DEVNAM     DEVBLKS
DEVCONNECT STATUS
-----
-----
disk     Disk1     mirror   Class1    Group1    c0t0d0     17682084
Node1                   ENABLE
disk     Disk2     mirror   Class1    *         c2t4d0     17682084
Node1                   ENABLE
```

6. Remove the swapped disk from the class.

With the `-c` and `-d` options of the `sdxdisk -R` command, specify the class name and the disk name confirmed in step 3.

```
# sdxdisk -R -c Class1 -d Disk2
```

7. Check if the swapped disk was removed from the class.

For checking, execute the `sdxinfo -D` command and check if information of the removed disk is not displayed.

```
# sdxinfo -D
OBJ      NAME      TYPE      CLASS     GROUP     DEVNAM     DEVBLKS
DEVCONNECT STATUS
-----
-----
disk     Disk1     mirror   Class1    Group1    c0t0d0     17682084
Node1                   ENABLE
```

8. Stop the system.

```
# shutdown -y -i0 -g0
```

9. Swap disks.

Ask your local customer support (CE) for disk unit swapping.

The CE will swap disk units according to the following procedures.

1. Take a copy of the WWN (World Wide Name) of the swapped disk.
2. Check the auto-boot? variable setting of OpenBoot, and if the auto-boot? variable value is "true", change the value to "false."
3. Turn off power of the main unit.
4. Swap disks.
5. Turn on power of the main unit.
6. When the auto-boot? value was changed to "false" in step 2., set it back to "true."

10. Restart the system.

Note

An FC-AL disk has been assigned a unique WWN (World Wide Name) and Solaris uses this WWN as management information. For Solaris to recognize the WWN of the replacement disk, the system must be rebooted with the boot -r command.

```
ok boot -r
```

11. Together with the CE, check if the replacement is recognized by Solaris and check its physical disk name.

In the following example, the WWN is 0020370e35e5 and the physical disk name is c2t4d0.

```
# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
    0. c2t4d0 <sfdisk type1.0; SUN9.0G cyl 4924 alt 2 hd 27 sec
133> 00100004
        /sbus@2,0/SUNW,socal@d,10000/sf@1,0/
ssd@w21000020370e35e5,0

Specify disk (enter its number):
```

Information

Disks managed by GDS are not displayed with the format command.

12. Check if the device special files of the replacement exist.

For the argument of the ls command, specify the physical disk name confirmed in step 11. attaching the asterisk (*) to the end.

```
# cd /dev/rdisk
# ls -lL c2t4d0*
crw-r----- 1 root    sys      118, 32 Feb  3 17:09 c2t4d0s0
crw-r----- 1 root    sys      118, 33 Feb  3 13:21 c2t4d0s1
crw-r----- 1 root    sys      118, 34 Feb  3 17:09 c2t4d0s2
crw-r----- 1 root    sys      118, 35 Feb  3 13:21 c2t4d0s3
crw-r----- 1 root    sys      118, 36 Feb  3 13:21 c2t4d0s4
crw-r----- 1 root    sys      118, 37 Feb  3 13:21 c2t4d0s5
```

```
crw-r----- 1 root sys 118, 38 Feb 3 13:21 c2t4d0s6
crw-r----- 1 root sys 118, 39 Feb 3 15:35 c2t4d0s7
```

13. If the relevant system is a cluster system, re-create the resource of the replacement according to the following procedures.

If not, go on to step 14.

1. Check the resource of the original disk.

Check if the resource of the original disk c2t4d0 connected to node Node1, confirmed in step 2., is displayed with the clgettree command.

```
# /etc/opt/FJSVcluster/bin/clgettree
Cluster 1 cluster

        Domain 2 Domain0
...
                Node 3 Node1 ON
...
                        DISK 33 c2t4d0 UNKNOWN
...
```

2. Delete the original disk resource.

```
# /etc/opt/FJSVcluster/bin/cldeldevice -n c2t4d0
```

3. Check if the original disk resource was deleted.

For checking, check if the resource confirmed in step 1. is not displayed with the clgettree command.

```
# /etc/opt/FJSVcluster/bin/clgettree
Cluster 1 cluster

        Domain 2 Domain0
...
                Node 3 Node1 ON
...
```

4. Create the resource of the replacement.

Execute the following command on any node that composes the cluster system.

```
# /etc/opt/FJSVcluster/bin/clautoconfig -r
```

5. Confirm that the resource of the replacement was created.

For checking, check if the resource of replacement c2t4d0 confirmed in step 11. is displayed with the clgettree command.

```
# /etc/opt/FJSVcluster/bin/clgettree
Cluster 1 cluster

        Domain 2 Domain0
...
                Node 3 Node1 ON
...
                        DISK 44 c2t4d0 UNKNOWN
...
```

14. Register the replacement with the class.

With the -c option of the sdxdisk -M command, specify the class name confirmed in step 3. With the -d option, specify the physical disk name confirmed in step 11. and the disk name confirmed in step 3.

```
# sdxdisk -M -c Class1 -d c2t4d0=Disk2
```

15. Confirm that the replacement was registered with the class.

```
# sdxinfo -D
OBJ      NAME    TYPE    CLASS   GROUP   DEVNAM  DEVBLKS  DEVCONNECT
STATUS
-----
-----
disk     Disk1   mirror  Class1  Group1  c0t0d0  17682084 Node1
ENABLE
disk     Disk2   mirror  Class1  *       c2t4d0  17682084 Node1
ENABLE
```

16. Connect the replacement to the group.

With the -c, -g, and -d options of the sdxdisk -D command, specify the class name, the group name, and the disk name confirmed in step 3.

```
# sdxdisk -C -c Class1 -g Group1 -d Disk2
```

17. Confirm that the replacement was connected to the group.

```
# sdxinfo -D
OBJ      NAME    TYPE    CLASS   GROUP   DEVNAM  DEVBLKS  DEVCONNECT
STATUS
-----
-----
disk     Disk1   mirror  Class1  Group1  c0t0d0  17682084 Node1
ENABLE
disk     Disk2   mirror  Class1  Group1  c2t4d0  17682084 Node1
ENABLE
```

18. When the replacement is connected to the group, resynchronization copying for the disk is performed.

If the resynchronization copying is in process, with the sdxinfo -S command, COPY is displayed in the STATUS field of the row that shows the disk name of the replacement. If the resynchronization copying is complete, ACTIVE is displayed in the STATUS field.

```
# sdxinfo -S
OBJ      CLASS   GROUP   DISK     VOLUME  STATUS
-----
-----
slice    Class1  Group1  Disk1    Volume1 ACTIVE
slice    Class1  Group1  Disk1    Volume2 ACTIVE
slice    Class1  Group1  Disk2    Volume1 COPY
slice    Class1  Group1  Disk2    Volume2 COPY
```

A.2.32 Swapping Internal Disks (SPARC T3-1/T3-2/T3-4)

If internal disks (SPARC T3-1/T3-2/T3-4) are swapped, the device numbers of the new disks may change. For this reason, internal disks (SPARC T3-1/T3-2/T3-4) cannot be swapped through the procedures for physical disk swapping with [Swap Physical Disk] and [Restore Physical Disk] in the GDS Management View or with the sdxswap command. For internal disks (SPARC T3-1/T3-2/T3-4) swapping, remove the relevant disk from the class once and register the new one with the class again.

The following example shows the procedures for swapping internal disks (SPARC T3-1/T3-2/T3-4) connected to a mirror group. In this example, the disk is mirrored and you do not need to remove volumes from the disk prior to swapping. If the disk to be swapped is not a mirrored disk, such as a single disk and a disk whose highest level group is not a mirror group, volumes must be removed from the disk in advance.

In this case, back up volume data as necessary, remove the volumes, remove the disk from the class, and register the new disk with the class before re-creating volumes and restoring the data.

1. Identify a raw device name (devname) of the disk to be swapped by referring to the GDS I/O error messages output in the `/var/adm/messages` file.

To search the GDS I/O error messages, use the string "sfdsk" as a key. In the following example, the raw device name is "disk@g5000c5001d4809ff."

```
WARNING: sfdsk: write error on slice:
        volume info:   devno(maj,min)=0x13300000002(307,2)
        device info:   devno(maj,min)=0x20000000038(32,56)
        devname=/scsi_vhci/disk@g5000c5001d4809ff (sd7)
        error info:    blk in disk=17308981, blk in slice=5232448
                       length=8192, resid=8192, errno=5
```

2. Check the device name of the disk to be swapped.

Specify the raw device name identified in step 1 as an argument of the `grep` command.

In the following example, the device name of the disk to be swapped is "c0t5000C5001D4809FFd0."

```
# ls -l /dev/rdisk | grep disk@g5000c5001d4809ff
lrwxrwxrwx 1 root root 52 Apr 25 13:46 c0t5000C5001D4809FFd0s0 -> ../../devices/
scsi_vhci/disk@g5000c5001d4809ff:a,raw
lrwxrwxrwx 1 root root 52 Apr 25 13:46 c0t5000C5001D4809FFd0s1 -> ../../devices/
scsi_vhci/disk@g5000c5001d4809ff:b,raw
lrwxrwxrwx 1 root root 52 Apr 25 13:46 c0t5000C5001D4809FFd0s2 -> ../../devices/
scsi_vhci/disk@g5000c5001d4809ff:c,raw
lrwxrwxrwx 1 root root 52 Apr 25 13:46 c0t5000C5001D4809FFd0s3 -> ../../devices/
scsi_vhci/disk@g5000c5001d4809ff:d,raw
lrwxrwxrwx 1 root root 52 Apr 25 13:46 c0t5000C5001D4809FFd0s4 -> ../../devices/
scsi_vhci/disk@g5000c5001d4809ff:e,raw
lrwxrwxrwx 1 root root 52 Apr 25 13:46 c0t5000C5001D4809FFd0s5 -> ../../devices/
scsi_vhci/disk@g5000c5001d4809ff:f,raw
lrwxrwxrwx 1 root root 52 Apr 25 13:46 c0t5000C5001D4809FFd0s6 -> ../../devices/
scsi_vhci/disk@g5000c5001d4809ff:g,raw
lrwxrwxrwx 1 root root 52 Apr 25 13:46 c0t5000C5001D4809FFd0s7 -> ../../devices/
scsi_vhci/disk@g5000c5001d4809ff:h,raw
```

3. Check the disk name, class name, and group name of the disk to be swapped.

Execute the `sdxinfo -D` command and find the line where the device name confirmed in step 2 is shown in the `DEVNAM` field.

In the following example, the disk name, class name, and the group name for the device name "c0t5000C5001D4809FFd0" are "rootDisk0002," "RootClass," and "rootGroup" respectively.

```
# sdxinfo -D
OBJ      NAME          TYPE  CLASS    GROUP    DEVNAM          DEVBKLS  DEVCONNECT  STATUS
-----
disk     rootDisk0001 mirror RootClass rootGroup c0t5000CCA0150FEA10d0 585912500 *           ENABLE
disk     rootDisk0002 mirror RootClass rootGroup c0t5000C5001D4809FFd0 585912500 *           ENABLE
disk     dataDisk0001 mirror RootClass dataGroup c0t5000CCA00AC1C874d0 585912500 *           ENABLE
disk     dataDisk0002 mirror RootClass dataGroup c0t5000CCA0150F96F0d0 585912500 *           ENABLE
```

4. Disconnect the disk to be swapped from the group.

Specify the class name, the group name, and the disk name confirmed in step 3 for the `-c`, `-g`, and `-d` options of the `sdxdisk -D` command.

```
# sdxdisk -D -c RootClass -g rootGroup -d rootDisk0002
```



If a spare disk is connected

When you try to disconnect the disk to be swapped while the spare disk is connected, the following error will be displayed.

```
SDX:sdxdisk: ERROR: spare disk connected for disk_name
```

In this case, you should disconnect the spare disk first.

The following example shows the procedure to disconnect the spare disk (disk0001) before disconnecting the disk to be swapped (rootDisk0002).

```
# sdxinfo -D
OBJ      NAME           TYPE  CLASS    GROUP    DEVDNAM                DEVBLKS  DEVCONNECT  STATUS
-----
disk     rootDisk0001  mirror RootClass rootGroup c0t5000CCA0150FEA10d0 585912500 *          ENABLE
disk     rootDisk0002  mirror RootClass rootGroup c0t5000C5001D4809FFd0 585912500 *          ENABLE
disk     disk0001      spare  RootClass rootGroup c0t5000CCA00AC1C874d0 585912500 *          ENABLE
```

```
# sdxdisk -D -c RootClass -g rootGroup -d disk0001
SDX:sdxdisk: INFO: disk0001: disconnected disk from group rootGroup
# sdxdisk -D -c RootClass -g rootGroup -d rootDisk0002
SDX:sdxdisk: INFO: rootDisk0002: disconnected disk from group rootGroup
```

-
5. Check if the disk to be swapped is disconnected from the group.

Execute the `sdxinfo -D` command and check if an asterisk (*) is displayed in the GROUP field of the disconnected disk (rootDisk0002).

```
# sdxinfo -D
OBJ      NAME           TYPE  CLASS    GROUP    DEVDNAM                DEVBLKS  DEVCONNECT  STATUS
-----
disk     rootDisk0001  mirror RootClass rootGroup c0t5000CCA0150FEA10d0 585912500 *          ENABLE
disk     dataDisk0001  mirror RootClass dataGroup c0t5000CCA00AC1C874d0 585912500 *          ENABLE
disk     dataDisk0002  mirror RootClass dataGroup c0t5000CCA0150F96F0d0 585912500 *          ENABLE
disk     rootDisk0002  undef  RootClass *          c0t5000C5001D4809FFd0 585912500 *          ENABLE
```

6. Remove the disk to be swapped from the class.

Specify the class name and the disk name confirmed in step 3 for the `-c` and `-d` options of the `sdxdisk -R` command.

```
# sdxdisk -R -c RootClass -d rootDisk0002
```

7. Check if the disk to be swapped has been removed from the class.

Execute the `sdxinfo -D` command to make sure the information of the deleted disk (rootDisk0002) is not displayed.

```
# sdxinfo -D
OBJ      NAME           TYPE  CLASS    GROUP    DEVDNAM                DEVBLKS  DEVCONNECT  STATUS
-----
disk     rootDisk0001  mirror RootClass rootGroup c0t5000CCA0150FEA10d0 585912500 *          ENABLE
disk     dataDisk0001  mirror RootClass dataGroup c0t5000CCA00AC1C874d0 585912500 *          ENABLE
disk     dataDisk0002  mirror RootClass dataGroup c0t5000CCA0150F96F0d0 585912500 *          ENABLE
```

8. Identify the port device name to which the disk to be swapped is connected.

Check the "Phys_Id" field to find the port device name for the disk to be swapped.

In the following example, the port device name of the physical disk "c0t5000C5001D4809FFd0" is `"/devices/pci@400/pci@2/pci@0/pci@e/scsi@0/iport@4:scsi."`

```
# cfgadm -av
Ap_Id                Receptacle  Occupant      Condition  Information
```

```

When          Type          Busy          Phys_Id
...
c5::w5000c5001d4809fd,0  connected      configured    unknown \
    Client Device: /dev/dsk/c0t5000C5001D4809FFd0s0(sd7)
unavailable disk-path  n \
    /devices/pci@400/pci@2/pci@0/pci@e/scsi@0/iport@4:scsi::w5000c5001d4809fd,0
...

```

9. Swap disks.

Ask your local customer support (CE) to swap your disk unit.

Both hot swap and non-hot swap are available.

10. Check the device name of the new disk.

The device name of the new disk is the name of the device connected to the port device identified in step 8 and its "Occupant" is displayed as "configured."

In the following example, the device name of the new disk is "c0t5000C5001D4806BFd0."

```

# cfgadm -av
Ap_Id          Receptacle  Occupant      Condition  Information
When          Type          Busy          Phys_Id
...
c5::w5000c5001d4809fd,0  connected    unconfigured  unknown   (sd7)
unavailable disk-path  n \
    /devices/pci@400/pci@2/pci@0/pci@e/scsi@0/iport@4:scsi::w5000c5001d4809fd,0
c5:: w5000c5001d4806bd,0  connected    configured     unknown \
    Client Device: /dev/dsk/c0t5000C5001D4806BFd0s0(sd8)
unavailable disk-path  n \
    /devices/pci@400/pci@2/pci@0/pci@e/scsi@0/iport@4:scsi:: w5000c5001d4806bd,0
...

```

11. Confirm that the device special file (cCtWWNdDs{0-7}) of the new disk exists.

Append an asterisk (*) to the end of the device name confirmed in step 10 and specify it as an argument of the ls command.

In the following example, you can see that the device special files from "c0t5000C5001D4806BFd0s0" to "c0t5000C5001D4806BFd0s7" exist.

```

# ls -lL /dev/rdisk/c0t5000C5001D4806BFd0*
crw-r----- 1 root    sys   32, 16 May 16 17:26 /dev/rdisk/c0t5000C5001D4806BFd0s0
crw-r----- 1 root    sys   32, 17 May 17 12:10 /dev/rdisk/c0t5000C5001D4806BFd0s1
crw-r----- 1 root    sys   32, 18 May 17 09:44 /dev/rdisk/c0t5000C5001D4806BFd0s2
crw-r----- 1 root    sys   32, 19 May 17 12:10 /dev/rdisk/c0t5000C5001D4806BFd0s3
crw-r----- 1 root    sys   32, 20 May 17 12:10 /dev/rdisk/c0t5000C5001D4806BFd0s4
crw-r----- 1 root    sys   32, 21 May 17 12:10 /dev/rdisk/c0t5000C5001D4806BFd0s5
crw-r----- 1 root    sys   32, 22 May 17 12:10 /dev/rdisk/c0t5000C5001D4806BFd0s6
crw-r----- 1 root    sys   32, 23 May 17 12:10 /dev/rdisk/c0t5000C5001D4806BFd0s7

```

12. If the relevant system is a cluster system, register the new disk resources with PRIMECLUSTER according to the following steps. If not, skip to step 13.

1. Identify the node ID of the cluster managed by PRIMECLUSTER.

In the following example, the node ID is "0."

```

# /etc/opt/FJSVcluster/bin/clgetnode
RID 3
KEY TRC89
RNAME TRC89
NODEID 0

```


13. Reload the GDS disk information.

```
# sdxinfo -x Refresh
```

14. Register the new disk with GDS.

Specify the class name of the original disk for the `-c` option of the `sdxdisk -M` command. Specify the device name of the new disk and the disk name of the original disk for the `-d` option.

```
# sdxdisk -M -c RootClass -d c0t5000C5001D4806BFd0=rootDisk0002
```

15. Confirm that the new disk is registered with GDS.

```
# sdxinfo -D
OBJ      NAME          TYPE   CLASS      GROUP      DEVDNAM          DEVBKLS  DEVCONNECT  STATUS
-----
disk     rootDisk0001  mirror RootClass  rootGroup   c0t5000CCA0150FEA10d0  585912500  *           ENABLE
disk     dataDisk0001  mirror RootClass  dataGroup   c0t5000CCA00AC1C874d0  585912500  *           ENABLE
disk     dataDisk0002  mirror RootClass  dataGroup   c0t5000CCA0150F96F0d0  585912500  *           ENABLE
disk     rootDisk0002  undef  RootClass  *           c0t5000C5001D4806BFd0  585912500  *           ENABLE
```

16. Connect the new disk to the group.

Specify the class name, group name, and disk name confirmed in step 3 for the `-c`, `-g`, and `-d` options of the `sdxdisk -C` command.

```
# sdxdisk -C -c RootClass -g rootGroup -d rootDisk0002
```

17. Confirm that the new disk is connected to the group.

```
# sdxinfo -D
OBJ      NAME          TYPE   CLASS      GROUP      DEVDNAM          DEVBKLS  DEVCONNECT  STATUS
-----
disk     rootDisk0001  mirror RootClass  rootGroup   c0t5000CCA0150FEA10d0  585912500  *           ENABLE
disk     rootDisk0002  mirror RootClass  rootGroup   c0t5000C5001D4806BFd0  585912500  *           ENABLE
disk     dataDisk0001  mirror RootClass  dataGroup   c0t5000CCA00AC1C874d0  585912500  *           ENABLE
disk     dataDisk0002  mirror RootClass  dataGroup   c0t5000CCA0150F96F0d0  585912500  *           ENABLE
```

18. When the new disk is connected to the group, Just Resynchronization of the new disk will be performed.

```
# sdxinfo -S
OBJ      CLASS      GROUP      DISK          VOLUME      STATUS
-----
slice    RootClass  rootGroup  rootDisk0001  rootVolume  ACTIVE
slice    RootClass  rootGroup  rootDisk0002  rootVolume  COPY
slice    RootClass  rootGroup  rootDisk0001  swapVolume  ACTIVE
slice    RootClass  rootGroup  rootDisk0002  swapVolume  COPY
slice    RootClass  rootGroup  rootDisk0001  homeVolume  ACTIVE
slice    RootClass  rootGroup  rootDisk0002  homeVolume  COPY
slice    RootClass  group0001  disk000001    volume0001  ACTIVE
slice    RootClass  group0001  disk000002    volume0001  ACTIVE
```

A.2.33 Use of GDS in Solaris 10 Environments

When using Solaris Zones, GDS installation, command execution, and volume creation can be performed only in the global zone. By importing GDS volumes created in the global zone to non-global zones, it becomes possible to use the GDS volumes in the non-global zones.

Understand and pay attention to the following points when using Solaris Zones.

- Install GDS in the global zone. GDS cannot be installed in a non-global zone.

- GDS commands cannot be executed in non-global zones. If commands of GDS installed in the /usr/sbin directory in the global zone are executed in non-global zones, the following error occurs.

```
ERROR: cannot be executed in non-global zone
```

If a GDS command is used in a non-global zone, the ld.so.1(1) error message may be output depending on the non-global zone creation method.

- When the same block within a GDS volume is accessed simultaneously from multiple zones, data consistency is not ensured, as in access to slices that are not managed with GDS. Access exclusion control for maintaining data consistency should be performed by the application that accesses the same block from multiple zones.



For the method for importing devices in the global zone to non-global zones, see the Solaris manual.

A.2.34 Emergency Startup Disk for Restoration

It is recommended that an emergency startup disk for restoration where the system for restoration is installed be excluded from GDS management.

The emergency startup disk for restoration cannot be managed as a local or shared class object.

It is possible to register the emergency startup disk for restoration with the root class and configure mirroring, but the following points should be taken into consideration.

- Even in a mirroring environment configured with GDS, the system cannot be booted from the mirror destination disk of the emergency startup disk for restoration.
- Because the system is booted for restoration work only from the mirror source disk, only data on the mirror source disk may be updated during the work and the data equivalence between the mirror source disk and the mirror destination disk is no longer guaranteed. Therefore, after the restoration work, unconfigure the mirroring before booting from the restored system disk. This operation procedure is different from that for "5.1.9 Unmirroring the System Disk." For details, contact your local customer support.



PRIMECLUSTER GDS Snapshot can be used to create replicas (snapshots) of system volumes to create an alternative boot environment in which the system is booted from the replication volume.

This alternative boot environment can serve as the system for restoration, instead of an emergency startup disk.

It is easier to create an alternative boot environment than to create an emergency startup disk, and additionally, it is possible to make the alternative boot environment redundant by configuring mirroring.

For details, see "1.5.6 Creating an Alternative Boot Environment."

A.3 General Points

A.3.1 Guidelines for Mirroring

Pay attention to the following guidelines when constructing mirroring configurations.

- Connecting disks and lower level groups with the same available sizes to mirror group is recommended.

The available size of a mirror group (the capacity available as volume) is the same as the available size of the smallest disk or lower level group that is connected.

When connecting disks or lower level groups with different available sizes to a mirror group, you will only be able to use the capacity of the smallest disk or lower level group. For example, if a 4 GB disk and a 9 GB disk is connected to one mirror group, only 4 GB out of the 9 GB disk will be accessible.

- Mirroring disks with similar performance specifications, or groups with the same configuration (including the performance specifications of disks that is connected to the group) is recommended.

When mirroring disks with different performance specifications such as revolution speed, the read performance becomes unbalanced and the write performance will depend on the slower disk performance.

The same applies when mirroring disks and groups, or when mirroring groups with different configuration.

A.3.2 Guidelines for Striping

Pay attention to the following guidelines when constructing striping configurations.

- In order to improve I/O performance with striping, it is necessary to adjust the stripe width and the number of stripe columns depending on the way an application accesses the disk.

If the striping configuration is not appropriate, you cannot gain much performance improvement. And, depending on the way an application accesses the disk, the performance may not improve even after adjusting the stripe width or the number of stripe columns.

- Do not make the stripe widths too large.

The sizes of stripe groups and stripe volumes are rounded to the common multiple of stripe width times the number of stripe columns and cylinder size. Therefore, if the stripe width is too large, use of the disk area may be inefficient or a volume with the intended size may be created.

- Where possible, connect disks and lower level groups with the same available sizes to the same stripe group.

The available size of the stripe group (available capacity as volumes) equals the available size of the smallest disk or the lower level group connected to the stripe group multiplied by the number of stripe columns and rounded down to the common multiple of stripe width times the number of stripe columns and cylinder size.

When connecting disks or lower level groups with different available sizes to a stripe group, the larger disk or lower level group will only be able to use the capacity of the smaller disk or lower level group. For example, if a 4 GB disk and a 9 GB disk are connected to one stripe group, the 9 GB disk will only be able to use approximately 4 GB. Which means, the available size of stripe group will be approximately 8 GB (4 GB x 2).

- Where possible, striping across disks with similar performance specifications is recommended.

When striping disks with different performance specifications such as revolution speed, the performance becomes unbalanced and will depend on the slower disk performance.

- Using striping in combination with mirroring is recommended.

In a striping configuration, the risk of losing data from a disk failure increases as more disks are involved compared to a usual disk configuration.

By mirroring stripe groups, both the I/O load balancing and data redundancy can be achieved at the same time.

A.3.3 Guidelines for Concatenation

Pay attentions to the following guidelines when constructing concatenation configurations.

- The available size of the concatenation group (available capacity as volumes) equals the total of the available size of disks connected to the concatenation group.

- Where possible, concatenating disks with similar performance specifications is recommended.

When concatenating disks with different performance specifications such as the revolution speed, the performance becomes unbalanced.

- Concatenation in combination with mirroring is recommended.

When concatenating disks, the risk of losing data from a disk failure increases as more disks are involved compared to a usual disk configuration.

By mirroring concatenation groups, large capacity and data redundancy can be achieved at the same time.

A.3.4 Guidelines for Combining Striping with Mirroring

Pay attention to the following guidelines when striping and mirroring simultaneously.

- Where possible, mirroring stripe groups with similar configuration is recommended.

See "[A.3.1 Guidelines for Mirroring](#)" and "[A.3.2 Guidelines for Striping](#)" also.

Appendix B Log Viewing with Web-Based Admin View

For details, see the supplementary "Web-Based Admin View Operation Guide."

Appendix C Web-Based Admin View Operating Environment Setting

For details, see the supplementary "Web-Based Admin View Operation Guide."

Appendix D Command Reference

This appendix discusses the commands provided by GDS and GDS Snapshot.

This appendix explains the format and facility of commands, specifiable options, and return values.

GDS provides the following commands.

Command	Function
sdxcass	Class operations
sdxdisk	Disk operations
sdxgroup	Group operations
sdxvolume	Volume operations
sdxslice	Slice operations
sdxinfo	Display object configuration and status information
sdxattr	Change attribute values of an object
sdxswap	Swap disk
sdxfix	Restore a crashed object
sdxcopy	Synchronization copy operations
sdxroot	Root file system mirroring definition and cancellation
sdxparam	Configuration parameter operations
sdxconfig	Object configuration operations

GDS Snapshot provides the following commands.

Command	Function
sdxproxy	Proxy object operations
sdxshadowdisk	Shadow disk operations
sdxshadowgroup	Shadow group operations
sdxshadowvolume	Shadow volume operations



Information

Commands That Operate Multiple Objects

When an error occurs in operation for part of the objects, the command may either continue operation for the other objects or terminate the process.

In either situation, referring to "[Appendix E GDS Messages](#)" check the meaning of the error message and take necessary action.

D.1 sdxcass - Class operations

SYNOPSIS

```
sdxcass -R -c class
```

DESCRIPTION

Use sdxcass to perform operations on class objects (excluding shadow class) specified by *class*. You must be superuser to use this command.

PRIMARY OPTIONS

You can use the following option.

-R

Remove

Removes the class definition specified by *class*.

If *class* is a shared class, the definition is removed from all nodes.

A disk registered with *class* will be removed with the *class*. However, if there is a group or a volume, the class will not be removed.

To place a removed disk under GDS management again, you need to re-register the physical disk in the class.

For further details, refer to "[D.2 sdxdisk - Disk operations](#)."

SUB OPTIONS

Sub options are as follows:

-c class

The *class* indicates the class name that is the target of the operation.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.

Otherwise, a non-zero value is returned.

D.2 sdxdisk - Disk operations

SYNOPSIS

```
sdxdisk -c -c class -g group -d disk, ...  
    [-v num=volume:jrm[, num=volume:jrm, ...]]  
    [-a attribute=value[, attribute=value]] [-e delay=msec ]
```

```
sdxdisk -D -c class -g group -d disk
```

```
sdxdisk -M -c class [-a attribute=value[, attribute=value, ...]]  
    -d device=disk[:type] [, device=disk[:type] , ...] [-e chkps]
```

```
sdxdisk -R -c class -d disk
```

DESCRIPTION

Use sdxdisk to perform operations on disk objects (excluding shadow disk) specified by *disk*.

You must be superuser to use this command.

PRIMARY OPTIONS

You can use either of the following options.

-C

Connect

Connects one or more disks (keep disks, single disks or undefined disks) specified by *disk*,... to a group specified by *group*. The *class* indicates the class name with which *disk* is registered.

To connect *disk* to a group other than a switch group, specify a disk connected to all the nodes that belong to the scope of *class*.

If no group with the same name as *group* exists, a new group is created.

You cannot connect a spare disk to a group. You cannot connect a keep disk or a single disk to an existing group. Also, a keep disk and a single disk, or multiple keep disks, or multiple single disks cannot be connected to the same group simultaneously.

The disk attribute will be changed to match the attribute of the group (mirror, stripe, concatenation, or switch) it is connected to. Disks and lower level groups connected to a group are mirrored, striped, concatenated or made exchangeable according to the type attribute.

Details about connecting disks to a mirror group, a stripe group, a concatenating group and a switch group are explained below.

When connecting to a mirror group

Disks and lower level groups connected to the same mirror group will mirror each other. When only one disk or one lower level group is connected, the volume created within that mirror group will not be mirrored. When configuring a mirroring environment with "n"-way multiplexing, "n" numbers of disks or lower level groups must be connected. A maximum of eight-way multiplex mirroring is supported.

When one or more volumes already exist within the mirror group specified by *group*, the slice configuration of disks or lower level groups that are already connected to *group* will be automatically copied to the newly connected *disks*.

Also, when there is an activated volume within *group*, volume data in addition to the slice configuration will be automatically copied upon returning from the *sdxdisk* command, therefore increasing the mirroring multiplicity.

By connecting a single disk with single volumes to a mirror group, single volumes will be changed to mirror volumes.

The available size of the mirror group (available capacity as volumes) will be the same as the available size of the smallest disk or the lower level group connected to the mirror group. When a keep disk is connected, the available size of the group will be the same as the available size of the keep disk. If connecting *disk* results in decreasing the available size of *group*, a warning message will be sent to standard error output.

When a keep disk is connected, a volume will be created for each physical slice within the keep disk. If the physical slice is not open, the created volume will be activated and synchronization copying will automatically begin after returning from the *sdxdisk* command.

In classes that include switch groups, mirror groups cannot be created.

When connecting to a stripe group

Disks specified by *disk*,... will be connected to *group* in the order they are listed. Disks and lower level groups connected to the same stripe group will each configure a stripe column, and will be striped in the order they were connected. When only one disk or one lower level group is connected, a volume cannot be created within that stripe group. When striping "n" number of columns, "n" number of disks or lower level groups must be connected. A minimum of two columns and a maximum of 64 columns are supported.

When a stripe group specified by *group* already exists, stripe columns will be added after the existing stripe columns in *group*, in the order they are specified by *disk*,.... However, you cannot increase stripe columns by connecting disks to stripe groups with volumes, or to stripe groups connected to a higher level group.

The available size of the stripe group (available capacity as volumes) equals the available size of the smallest disk or the lower

level group connected to the stripe group multiplied by the number of stripe columns, and rounded down to the common multiple of stripe width times stripe columns and cylinder size. If connecting *disk* decreases the available size of *group*, a warning message will be sent to standard error output.

You cannot connect a single disk to a stripe group.

In classes that include switch groups, stripe groups cannot be created.

When connecting to a concatenation group

Disks connected to the same concatenation group specified by *disk,...* will be concatenated in the order they are listed. A maximum of 64 disks can be concatenated.

The available size of the concatenation group (available capacity as volumes) equals the total of the available size of disks connected to the concatenation group.

The available size of the group can be increased by connecting disks to the existing concatenation group. When a concatenation group specified by *group* already exists, disks will be added in the order they are specified by *disk,...* after the disk that was last concatenated in *group*. However, if the concatenation group specified by *group* is connected to a stripe group that is connected to a mirror group, disks cannot be added.

You cannot connect a single disk to a concatenation group.

In classes that include switch groups, concatenation groups cannot be created with this command. Additionally, to concatenation groups to which lower level switch groups are connected, disks cannot be added.

When connecting to a switch group

One of two disks connected to a switch group becomes the active disk and the other one becomes the inactive disk. Use the `sdxattr -G` command to switch the disks from active to inactive and vice versa. A switch group can connect a maximum of two disks. If only one disk is connected, the disk acts as the active disk and an active disk switch cannot be performed.

To create a switch group specified by *group*, specify the active disk in the `-a actdisk` option. The other disk not specified by the `-a actdisk` option becomes the inactive disk. When connecting *disk* to an existing switch group specified by *group*, the *disk* becomes the inactive disk.

Unlike mirror groups, even if a switch group specified by *group* already includes volumes, synchronization copying to the newly connected *disk* is not performed. To perform a disk switch for continuous services in the event of an error in the active disk, previously create copies of data from the active disk to the inactive disk with a disk unit's copy function and so on.

The available size of a switch group (capacity available for volumes) conforms to the available size of the smallest disk connected to the switch group. If the available size of *group* decreases as a result of *disk* connection, a warning message is sent to standard error output.

In classes that include any one of the following objects, switch groups cannot be created.

- Disk other than an undefined disk
- Mirror group
- Stripe group
- Concatenation group to which no lower switch group is connected

class must be a shared class of which scope includes 2 nodes. The physical scopes of the active disk and the inactive disk must also meet either the following conditions.

- The active disk and the inactive disk are connected to both the nodes included in the scope of *class* and are not connected to nodes not included in that scope.
- The active disk is connected to only one of the nodes included in the scope of *class* and the inactive disk is connected to the other node included in that scope.

-D

Disconnect

Disconnects a disk (including a spare disk) specified by *disk* from a group specified by *group*.

The *class* indicates the class name with which the *disk* is registered, and the *group* indicates the group name to which *disk* is connected.

The disconnected disk will return to its original type attributes (keep disk, single disk or undefined disk).

If only *disk* is connected to *group*, *group* will automatically be removed upon disconnecting *disk*. However, when *disk* is the only object connected to *group* and *group* is connected to a higher level group, disconnection will result in an error. In such case, disconnect *group* from the higher level group using the `sdsgroup -D` command, and then disconnect *disk*.

You cannot disconnect *disk* if the disconnection will result in a change in the status of any of the existing volumes within *group*.

Conditions on when you cannot disconnect a disk from a mirror group, a stripe group, a concatenation group or a switch group are explained below.

When disconnecting from a mirror group

For example, *disk* cannot be disconnected from a mirror group if one or more volumes exist within the mirror group specified by *group* and the disk specified by *disk* is the only object connected to *group*.

When disconnecting from a stripe group

A disk cannot be disconnected from a stripe group with one or more existing volumes, or from a stripe group connected to a higher level group.

When disconnecting from a concatenation group

The only disk that can be disconnected from a concatenation group is the disk that was concatenated last.

Disks containing volume areas cannot be disconnected from a concatenation group.

If the concatenation group specified by *group* is connected to a stripe group that is connected to a mirror group, disks cannot be disconnected.

When disconnecting from a switch group

Inactive disks can be disconnected regardless whether or not volumes exist.

The active disk can be disconnected from a switch group if all the following conditions are satisfied.

- The switch group is not connected to a higher level concatenation group.
- The switch group includes no volume.
- The inactive disk is not connected to the switch group.

If the switch group includes volumes, before disconnecting the active disk, remove those volumes. If the switch group includes the inactive disk, switch the active disk to it with the `sdxattr -G` command and then disconnect the former active disk.

-M

Make

Registers one or more physical disks, specified by *device*, with *class*. The *class* gives the name of the destination class. Once physical disks have been registered, they can then be managed using GDS. A disk managed by GDS is called an SDX disk. Users will use the disk name specified by *disk* to perform operations on the disk.

If no class with the name specified by *class* already exists, then one is automatically created.

The "keep" type *device* may be registered with "root" type *class*. However, when registering multiple keep type devices together, you must register "undef" type *devices* of the same number or more at the same time.



Since the `sdxdisk` command initializes the registered physical disks (except for "keep" type *devices*), when registering a physical disk containing data, you must first create data backup.

-R

Remove

Removes a disk specified by *disk* from a class specified by *class*. The *class* indicates the class name with which the *disk* is registered.

Once the disk is removed, it can no longer be managed using GDS.

When the last disk is removed from a *class*, that class definition is automatically removed.

A disk cannot be removed when a volume exists within *disk*, or when *disk* is connected to a group.

If removal of *disk* will result in *class* closure, the *disk* cannot be removed. The *class* will be closed when it includes:

- less than 3 disks in ENABLE status and no disk normally accessible
- three to 5 disks in ENABLE status and less than 2 disks normally accessible
- six or more disks in ENABLE status and less than 3 disks normally accessible

SUB OPTIONS

Sub options are as follows:

-a *attribute=value[,attribute=value]* (when using **-C**)

When using the **-C** option and defining a new group name with the **-g** option, a new group is automatically created. Using this option sets the *attribute* of the created *group* to *value*.

The *attribute* indicates the attribute name, and *value* indicates the attribute value. Always separate *attribute* and *value* with an equal (=) sign. When indicating multiple attributes, specifiers should be combined using commas(,) as the delimiter.

If no group is created, indicating a different attribute value from the existing *group* will result in an error. To change the attribute value of an existing *group*, use the `sdxattr -G` command.

You can indicate the following combination to *attribute* and *value*.

If multiple attributes are indicated and an error results because of any part of them, the entire process is canceled.

type=mirror, type=stripe, type=concat or type=switch (default is **mirror**)

Sets the type attribute of *group*. If *class* is the root class, specifying "stripe" or "concat" will result in an error. When *class* is not a shared class of which scope includes 2 nodes, specifying "switch" will also result in an error.

mirror

Sets type attribute to "mirror".

stripe

Sets type attribute to "stripe".

concat

Sets type attribute to "concatenation".

switch

Sets type attribute to "switch".

width=blks (default is 32)

Sets the stripe width of *group*. The *blks* indicates the stripe width in blocks (base 10). One block is 512 bytes. For *blks*, you can indicate an integer (one and larger) that is two raised to the power, which is equal to or smaller than the available size of the smallest disk specified by *disk*,.... If *hgroup* is not a stripe group, this option will result in an error.

actdisk=disk

Sets the active disk of *group*. Specify a disk name of the active disk into *disk*. When *group* is an existing group, not specifying the -a type=switch option, or specifying a disk other than that specified by -d option into *disk*, will result in an error.

-a attribute=value[,attribute=value,...] (when using **-M**)

When using the -M option and defining a new class name with the -c option, a class is automatically created. Using this option sets the created *class attribute* to *value*.

The *attribute* indicates the attribute name, and the *value* indicates the attribute value. Always separate *attribute* and *value* with an equal (=) sign. Specifiers should be combined using commas(,) as the delimiter.

If no class is created, this option is ignored. To change the attributes of an existing class, use the `sdxattr -C` command.

You can use the following *attribute* and *value* pairs in any combination.

If multiple attributes are specified and an error results because of any part of them, the entire process is canceled.

type=root, type=local or type=shared (default is **local**)

Sets the attributes of *class* type.

root

Sets the type attribute to "root".

Objects within "root" type *class* can only be used by the current node.

Unlike local type class, a system disk including root file system may be registered with "root" type *class*.

Only one root type class can be created within a system. When a root type class already exists, indicating a new "root" type

class will result in an error.

For the scope attribute, the current node's node identifier will be automatically set.

local

Sets the type attribute to "local".

Objects within "local" type class can only be used by the current node.

For the scope attribute, the current node's node identifier will be automatically set.

shared

Sets the type attribute to "shared."

By combining this with the scope attribute, the objects in the *class* can be shared among multiple nodes including the current node.

A shared type class can include physical disks connected to all the nodes that belong to the scope. When the scope includes 2 nodes, disks connected to only one node in the scope can be registered as undefined disks.

scope=node [:node:...] (default is the current node only)

Sets the node set which share the *class* whose type attribute is specified to be "shared".

In *node*, indicates a node identifier that is defined by PRIMECLUSTER.

hs=on or **hs=off** (default is **on**)

It sets the operation of the hot spare.

on

Enables the hot spare.

off

Disables the hot spare. If the operation mode is set to off, spare disk automatic connection is restrained.

hsmode=exbox or **hsmode=bybox** (default is **exbox**)

Sets the spare disk selection mode for automatic connection by hot spare.

exbox

Sets the spare disk selection mode to the external mode.

If an I/O error occurs in a disk of a disk array unit, this method selects a spare disk that belongs to a different disk case from that of the failed disk. If an I/O error occurs in a disk irrelevant to a disk array unit (such as an internal disk), it selects a spare disk that is connected to a different controller from that of the failed disk. When no applicable unconnected spare disk is found there, a spare disk that belongs to the same disk case or is connected to the same controller as that of the disk with the I/O error, is selected.

bybox

Sets the spare disk selection mode to the internal mode.

If an I/O error occurs in a disk of a disk array unit, this method selects a spare disk that belongs to the same disk case as that of the failed disk. If an I/O error occurs in a disk irrelevant to a disk array unit (such as an internal disk), it selects a spare disk that is connected to the same controller as that of the failed disk. When no applicable unconnected spare disk is found there, spare disk automatic connection is restrained.

-c class

The *class* indicates the class name to which the disk is registered or is to be registered, where the disk is the target of the operation.

-d device=disk[:type] [,device=disk[:type],...] (when using -M)

The *device* indicates the name of the physical disk, the *disk*, the name of the disk, and *type*, the type attribute of the disk. The *device* must always be followed by an equal sign (=), and if a type is given, it is delimited from *disk* by a colon (:). To register multiple *devices*, combine multiple definitions with a comma (,) as the delimiter.

The physical disk name can be specified in either the following formats:

```
cCtTdD    (for normal hard disks)
mphd/I    (for MPHD disks)
mplb/I    (for MPLB disks)
emcpower/N (for emcpower disks)
```

Where *C* is the controller number, *T* is the target ID, *D* is the disk number, *I* is the MPHD instance number or the MPLB instance number, and *N* is the emcpower device number.

emcpower device number is either a 1 or 2 digit number using numbers from 0 to 9, and does not specify small letters corresponding to slice number.

The *type* can be indicated as one of the following. If no type is indicated, the default value of undef (undefined disk) is used. If *class* is not the root class, specifying "keep" will result in an error. If *device* is not connected to part of nodes included in the scope of *class*, or if a switch group exists in *class*, specifying a type other than "undef" will also result in an error.

keep

Keep disk. When it is registered with a class or connected to a group, the format and data of the disk will be retained.

single

Single disk. Single volumes can be created on it.

spare

Spare disk.

undef

Undefined disk.

When "spare" is specified for *type*, and the available size of *device* is smaller than the available size of the smallest mirror group in *class*, a warning message notifying that the hot spare function may not be available will be sent to standard error output.

When only one *device* is specified and "keep" is specified for *type*, the *device* must have an unused physical slice number and sufficient free space so that the private slice can be created. When multiple *devices* are specified, *device* of which *type* is specified as "keep" must have sufficient free space or a swap device with sufficient space.

The *sdxdisk* command on system disks indicated as "keep" by *type* will terminate normally, even if all or part of the physical slices are open.

However, if anything other than system disk is indicated as "keep", an open physical slice will result in an error.

For example, if a physical slice is used as a file system, you must first unmount the file system to regain the unused status of physical slice before executing the *sdxdisk* command.

When "single" is specified for *type*, *device* will be registered as a single disk. For a "single" disk, you can create single volumes on it using the *sdxvolume* command without connecting the disk to any group.

-d *disk* (when using **-D, -R**)

The *disk* indicates the disk name that is the target of the operation.

-d *disk*,... (when using **-C**)

The *disk* indicates the disk name that is the target of the operation. To indicate multiple disks, separate each disk name with a comma (,) as the delimiter.

-e *chkps* (when using **-M**)

Registers *device* with *class* even if the private slice exists in the *device*, as far as disk identification information (class and disk names) stored in the private slice matches identification information of a disk already registered with the *class*. For example, if *device* contains a copy of the private slice of a disk that is already registered with *class*, to register the *device* to the *class*, turn on this option.

If *class* is not a shared class, this command results in an error.

-e *delay=msec* (when using **-C**)

When a disk is connected to a mirror group, data contained in the volume will be copied as needed.

This option delays the issuing of the input/output request to the disk at the time of copying by milliseconds specified by *msec*, allowing adjustment for the influence on the application accessing the volume.

The value is set to 0 by default.

Values from 0 to 1000 may be specified for *msec*.

If *group* is not a mirror group, this option is ignored.

-g *group* (when using **-C,-D**)

The *group* indicates the group name to which the disk is connected, or is to be connected, where *disk* is the target of the operation.

-v *num=volume:jrm*[,*num=volume:jrm*,...] (when using **-C**)

When connecting a "keep" type *disk*, indicates the attribute value of the volume that is to be created.

If no "keep" type *disk* is specified, this option is ignored.

The *num* must always be followed by an equal sign (=), and *volume* and *jrm* must be delimited by a colon (:). To indicate multiple volume attributes, combine multiple definitions with a comma (,) as the delimiter.

The *num* indicates the physical slice number (integer between 0 and 7 inclusive, excluding 2) of the "keep" type *disk* where the volume data is stored, *volume* gives the name of the volume to be created, and *jrm* the just resynchronization mode setting of the created volume, which can be either on or off.

When there is more than one physical slice larger than 0 in the "keep" type *disk*, you must specify volume attributes corresponding to all physical slices.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.

Otherwise, a non-zero value is returned.

D.3 sdxgroup - Group operations

SYNOPSIS

```
sdxgroup -C -c class -h hgroup -l lgroup, ...  
          [-a attribute=value[,attribute=value]] [-e delay=msec]
```

```
sdxgroup -D -c class -h hgroup -l lgroup
```

```
sdxgroup -R -c class -g group
```

DESCRIPTION

Use sdxgroup to perform operations on group objects (excluding shadow groups). You must be superuser to use this command.

PRIMARY OPTIONS

You can use one of the following options.

-C

Connect

Connects one or more groups (stripe group, concatenation groups, or switch groups) specified by *lgroup*,... to a group (a mirror group, stripe group, or concatenation group) specified by *hgroup*. The *class* indicates the name of the class to which *lgroup* belongs. If *class* is the root class, this command results in an error.

When no group with the same name as *hgroup* exists, a group will be created automatically.

Groups specified by *hgroup* are referred to as higher level group, and groups specified by *lgroup* are referred to as lower level group.

Lower level groups and disks connected to the same higher level group are mirrored, striped, or concatenated according to the type attribute of the higher level group. Connecting a group to a higher level group does not change the type attribute of the lower level group.

You cannot connect groups when:

- *lgroup* is a mirror group
- *hgroup* is a switch group
- type attributes of *lgroup* and *hgroup* are the same

A group that already contains volumes cannot be connected to another group.

Details about connecting groups to a mirror group, a stripe group, and a concatenation group are explained below.

When connecting to a mirror group

You can connect one or more groups (stripe group or concatenation group) specified by *lgroup*,... to *hgroup* which is a mirror group.

Disks and lower level groups connected to the same mirror group will mirror each other. When only one disk or one lower level group is connected, volumes created within that mirror group will not be mirrored. When configuring a mirroring environment with "n"-way multiplexing, "n" numbers of disks or lower level groups must be connected. A maximum of eight-way multiplex mirroring is supported.

When one or more volumes already exist within the mirror group specified by *hgroup*, the slice configuration of disk or lower level

group that is already connected to *hgroup* will be automatically copied to the newly connected *lgroup*. Also, when there is an activated volume within *hgroup*, volume data in addition to the slice configuration will be automatically copied upon returning from the *sdxgroup* command, therefore increasing the mirroring multiplexity.

The available size of the mirror group (available capacity as volumes) will be the same as the available size of the smallest disk or the lower level group connected to the mirror group. If connecting *lgroup* decreases the available size of *hgroup*, a warning message will be sent to standard error output.

In classes that include switch groups, mirror groups cannot be created.

When connecting to a stripe group

You can connect one or more groups (concatenation group) specified by *lgroup,...* to *hgroup* which is a stripe group. Groups specified by *lgroup,...*, will be connected to *hgroup* in the order they are listed.

Disks and lower level groups connected to the same stripe group will each configure a stripe column, and will be striped in the order they are connected. When only one disk or one lower level group is connected, a volume cannot be created within that stripe group. When striping "n" number of columns, "n" number of disks or lower level groups must be connected. A minimum of two columns and a maximum of 64 columns are supported.

When a stripe group specified by *hgroup* already exists, stripe columns will be added after the stripe columns that already exist in *hgroup*, in the order they are specified by *lgroup,...*. However, you cannot increase the stripe columns by connecting groups to stripe groups with volumes, or to stripe groups connected to a higher level group.

The available size of a stripe group (available capacity as volumes) equals the available size of the smallest disk or the lower level group connected to the stripe group multiplied by the number of stripe columns, and rounded down to the common multiple of stripe width times stripe columns and cylinder size. If connecting *lgroup* decreases the available size of *hgroup*, a warning message will be sent to standard error output.

In classes that include switch groups, stripe groups cannot be created.

When connecting to a concatenation group

This command can connect one or more groups (switch groups) specified by *lgroup,...* to *hgroup* which is a concatenation group. Switch groups connected to the same concatenation group will be concatenated in the order they are specified in *lgroup,...*. Concatenation of a maximum of 64 groups is supported.

The available size (available capacity as volumes) of a concatenation group equals the total available size of lower level groups connected to the concatenation group.

By connecting lower level groups to an existing concatenation group, the available size of the concatenation group can increase. If the concatenation group specified by *hgroup* already exists, lower level groups are concatenated in the order they are specified in *lgroup,...* following the last concatenated lower level group in *hgroup*. However, to concatenation groups connected to higher level groups, lower level groups cannot be connected.

To concatenation groups to which disks are connected, switch groups cannot be connected.

-D

Disconnect

Disconnects group specified by *lgroup* from the higher level group *hgroup*. The *class* indicates class name to which *lgroup* belongs, and *hgroup* indicates the higher level group name to which *lgroup* is connected.

When *lgroup* is the only object connected to *hgroup*, *hgroup* will automatically be removed upon disconnecting *lgroup*. However, when *lgroup* is the only object connected to *hgroup*, and *hgroup* is connected to a higher level group, disconnection will result in an error. In such case, disconnect *hgroup* from its higher level group, and then disconnect *lgroup*.

You cannot disconnect *lgroup* if the disconnection may result in a change in the status of any existing volume within *hgroup*.

Restrictions that prevent group disconnection from a mirror group, a stripe group, and a concatenation group are explained below.

When disconnecting from a mirror group

For example, you cannot disconnect *lgroup* from a mirror group if one or more volumes exist within the mirror group specified by *hgroup*, and *lgroup* is the only object connected to *hgroup*.

When disconnecting from a stripe group

You cannot disconnect a lower level group from a stripe group with one or more existing volumes, or from a stripe group connected to a higher level group.

When disconnecting from a concatenation group

Only the last concatenated lower level group can be disconnected from a concatenation group.

Lower level groups that have volume areas cannot be disconnected from concatenation groups.

-R

Remove

Remove the group definition specified by *group*. The *class* indicates the class name to which *group* belongs.

Disks and lower level groups connected to *group* will be disconnected. The disconnected disk's attribute will return to its original setting (keep disk, single disk or undefined disk).

group cannot be removed when one or more volumes exist within *group*, or when *group* is connected to a higher level group.

SUB OPTIONS

Sub options are as follows:

-a *attribute=value[,attribute=value]* (when using **-C**)

When using the **-C** option and defining a new group name with the **-h** option, a new group *hgroup* is automatically created. Using this option sets the *attribute* of the created *hgroup* to *value*.

The *attribute* indicates the attribute name, and *value* indicates the attribute value. Always separate *attribute* and *value* with an equal (=) sign. When indicating multiple attributes, specifiers should be combined using commas(,) as the delimiter.

If no group is created, indicating a different attribute *value* from the existing *hgroup* will result in an error. You cannot change the attribute value of an existing *hgroup*.

You can indicate the following combination to *attribute* and *value*.

If multiple attributes are indicated and an error results because of any part of them, the entire process is canceled.

type=mirror, type=stripe or type=concat (default is **mirror**)

Sets the type attribute of *hgroup*.

mirror

Sets the type attribute to "mirror."

stripe

Sets the type attribute to "stripe."

concat

Sets the type attribute to "concatenation."

width=blks (default is 32)

Sets the stripe width of *hgroup*. The *blks* indicates the stripe width in block number (base 10). One block is 512 bytes. For *blks*, you can indicate an integer that is two raised to the power (one and larger), which is equal to or smaller than the available size of the smallest group specified by *lgroup*,... If *hgroup* is not a stripe group, this option will result in an error.

-c class

The *class* indicates the class name to which the group belongs, where *group* is the target of the operation.

-e delay=msec (when using **-C**)

When a group is connected to a mirror group, data contained in the volume will be copied as needed.

This option delays the issuing of the input/output request to the disk by milliseconds specified by *msec*, allowing adjustment for the effect on the application accessing the volume.

Default is 0.

Values from 0 to 1000 may be specified for *msec*.

If *hgroup* is not a mirror group, this option is ignored.

-g group (when using **-R**)

The *group* indicates the group name that is the target of the operation.

-h hgroup (when using **-C,-D**)

The *hgroup* indicates the higher level group name to which the lower level group is connected or is to be connected, where the lower level group is the target of the operation.

-l lgroup (when using **-D**)

The *lgroup* indicates the lower level group name that is the target of the operation.

-l lgroup,... (when using **-C**)

The *lgroup* indicates the lower level group name that is the target of the operation. To connect multiple groups, separate each group name with a comma (,) as the delimiter.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.
Otherwise, a non-zero value is returned.

D.4 sdxvolume - Volume operations

SYNOPSIS

```
sdxvolume -F -c class [-v volume,... ]  
          [-e { allnodes|node=node[:node:...]}]
```

```
sdxvolume -M -c class {-g group|-d disk} -v volume -s size  
          [-a attribute=value [, attribute=value]] [-e delay=msec]
```

```
sdxvolume -N -c class [-v volume,... ]  
          [-e [allnodes|node=node [:node:... ]], delay=msec, mode=val, nosync, unlock]
```

```
sdxvolume -R -c class -v volume
```

```
sdxvolume -S -c class -v volume -s size
```

DESCRIPTION

Use sdxvolume to perform operations on volume objects (excluding shadow volumes) specified by *volume*.
You must be superuser to use this command.

PRIMARY OPTIONS

You can use either of the following options.

-F

oFline

Stops one or more volumes specified by *volume*,... If the -v option is omitted, all volumes within the *class* are taken offline. Offline volumes cannot be accessed.

If synchronization copying of *volume* is in process, it cannot be stopped. You can stop the synchronization copying using the sdxcopy -C command.

The stopped volume will be activated when the node is rebooted. (Except for when the volume is locked.)

Attempting this operation while *volume* is in use results in an error.

-M

Make

Creates a volume specified by *volume*, within the highest level group specified by *group*, or within a single disk specified by *disk*. The *size* indicates the number of blocks on *volume*, *class* indicates the class name associated with the *group* or *disk*.

If *class* is the root class, a maximum of 6 volumes with their physical slice attribute set to "on" can be created within the same *group*.

If *class* is a local class or a shared class, the same *group* or *disk* can contain a maximum of 6 volumes with their physical slice attribute set to "on." A maximum of 1024 volumes can be created in total, including the volumes with physical slice attribute set to "off."

When `-a pslice=value` option is omitted, volumes with physical slice attribute set to "on," will be created. However, note that you cannot create a volume with physical slice attribute set to "on," if *group* is stripe group, concatenation group, or a mirror group where its only directly connected object is a lower level group. In such case, you must indicate the `-a pslice=off` option, and set the physical slice attribute to "off."

After volume creation is complete, the volumes are started on a node where the command was executed and become accessible through the following special files.

```
/dev/sfdsk/class/rdsk/volume
/dev/sfdsk/class/dsk/volume
```

If *group* is a mirror group, the system will automatically execute a synchronization copying upon returning from the `sdxvolume` command.

The features of volumes created when *group* is mirror group, stripe group and switch group are explained below.

When group is a mirror group

To ensure data availability, GDS restricts the mirroring on a single piece of disk unit. In the case of mirror groups, a mirror volume that consists of mirror-multiplexing equal to the number of connected disks or lower level groups is created (maximum of eight). When only one disk or one lower level group is connected, the volume created within that mirror group will not be mirrored.

If the last block number of the volume that is created within a mirror group is larger than the available size of any of the spare disks registered with *class*, a warning message is sent to standard error output informing you that the hot spare feature is disabled.

When group is a stripe group

In a stripe group, stripe volumes with columns equal to the number of connected disk or lower level groups are created. When only one disk or lower level group is connected, volume cannot be created.

When group is a switch group

In a switch group, switch volumes with redundancy equivalent to the number of connected disks (a maximum of 2) are created. If only one disk is connected, an active disk switch cannot be performed.

If the active disk is not connected to a node where the command was executed, the volumes are not started. To use the created switch volumes, perform an active disk switch with the `sdxattr -G` command, or move to a node to which the active disk is connected, and then start the volumes with the `sdxvolume -N` command.

-N

oNline

Activates one or more volumes specified by *volume*,.... If the `-v` option is omitted, all volumes within *class* are activated. Activated volumes can be accessed.

If there is a slice in TEMP status on the *volume*, a warning message is sent to standard error output.

If *volume* is a mirror volume, the system will determine whether synchronization has been lost upon returning from the `sdxvolume` command and automatically execute a synchronization copying as needed (except for when `-e nosync` is specified).

If *volume* is a switch volume, it cannot be started on nodes to which the active disk is not connected. If *volume* belongs to the highest level concatenation group to which lower level switch groups are connected, it also cannot be started on nodes to which the active disk of *volume* is not connected.

-R

Remove

Removes the volume specified by *volume* and releases the disk area used in the group or the single disk.

If the specified volume is active, this command results in an error.



Be aware that any data stored on *volume* will be lost.

-S

reSize

Expands the size of a volume specified by *volume* to *size* blocks.

class indicates the name of a class to which *volume* belongs.

volume must be a volume that belongs to any one of:

- A single disk
- A mirror group that consists of only one disk
- A mirror group that consists of only one lower level group

The size can be expanded even when the *volume* is active.

The first block of *volume* is not changed. If any area of a volume other than *volume* exists in the area of *size* blocks after the first block of *volume*, it results in an error.

SUB OPTIONS

Sub options are as follows:

-a *attribute=value[,attribute=value]* (when using **-M**)

Use this to set an *attribute* for the *volume*.

The *attribute* indicates the attribute name, and the *value* indicates the attribute value. Always separate *attribute* and *value* with an equal (=) sign. When indicating multiple attributes, specifiers should be combined using commas(,) as the delimiter.

You can indicate the following combination to *attribute* and *value*.

If multiple attributes are specified and an error results because of any part of them, the entire process is canceled.

jrm=on or **jrm=off** (default is **on**)

Sets the volume's JRM mode.

If a group other than a mirror group is specified in the **-g** option, this command results in an error.

on

JRM is "on".

off

JRM is "off".

pslice=on or pslice=off (default is **on**)

Sets the physical slice attribute value of volume. When using **-g** option to indicate a group in which a physical slice cannot be created (i.e. stripe group, concatenation group, or mirror group where the only object directly connected to mirror group is lower level group), this option cannot be set to "on." If *class* is the root type, this option cannot be set to "off."

on

Physical slice attribute value is set to "on". Among slices composing the volume, the slices within the single disk, the disks that are connected to the switch group or the disks that are directly connected to the mirror group are registered to the VTOC, and physical slices are created.

off

Physical slice attribute value is set to "off". None of the slices consisting the volume is registered to the VTOC, and physical slice will not be created. When physical slice attribute is set to "off", the slice cannot be detached.

-c class

The *class* indicates the class name to which the volume that is the target of the operation belongs, or the class name in which the volume is to be created.

-d disk (when using **-M**)

The *disk* indicates the single disk name in which the single volume will be created.

-e allnodes (when using **-F,-N**)

Stops or activates the volume on all nodes included in the scope of *class*. Stopped nodes are ignored. *class* must be a shared class.

When neither this option nor **-e node=node[:node:...]** option is specified, volume is stopped or started only on the self-node.

-e delay=msec (when using **-M,-N**)

If synchronization is not maintained when creating or activating a mirror volume, synchronization copying will take place automatically (except for when **-e nosync** is specified).

This option delays the issuing of the input/output request to the disk at the time of copying by milliseconds specified by *msec*, allowing adjustment for the effect on the application accessing the volume.

The value is set to 0 by default.

Values from 0 to 1000 may be specified for *msec*.

If mirror group is not specified with *group*, this option is ignored.

-e mode=val (when using **-N**)

Specifies the access mode for one or more volumes that will be activated.

val indicates either of the following options.

rw

Sets access mode for read and write.

ro

Sets access mode for read only. Opening a read-only volume in write mode will result in an error.

Although volume will be activated in the access mode specified by *val*, the access mode attribute for volume will remain unchanged. Access mode specified by *val* ("Current Access Mode") is valid only while the volume is activated, and will become invalid once the volume is stopped. When the volume is restarted, it will start in the mode set by access mode attribute ("Default Access Mode"), except for when the access mode is specified at the time of restart.

In order to start a volume that is already activated on the current node in a different access mode, you must first stop the volume.

-e **node**=*node*[:*node*,...] (when using -F,-N)

Stops or activates the volume on one or more specified nodes.

Stopped nodes are ignored. You must specify the node identifier of the node to stop or activate the volume to *node*. If a node not included in the scope of *class* is specified, the volume is not stopped or activated on any node. *class* must be a shared class.

If this option and the -e allnodes are both omitted, the volume is stopped or activated only on the current node.

-e **nosync** (when using -N)

Disables automatic synchronization copying after activating a mirror volume.

If mirror group is not specified with *group*, this option is ignored.



Volumes that are activated using this option will not be mirrored. In order to configure a mirroring environment, you must perform synchronization copying with the `sdxcopy -B` command.

-e **unlock** (when using -N)

The volume will be activated regardless to whether or not it is locked.

Lock mode will not be changed unless you change it with the `stxattr -V` command.

-g *group* (when using -M)

The *group* indicates the group name in which the volume will be created.

-s *size* (when using -M)

Specifies the size of the *volume* being created, in blocks (base 10). One block is 512 bytes.

When *group* indicates stripe group, the size of *volume* created will be *size* rounded up to a common multiple of stripe width multiplied by stripe columns and cylinder size. In other cases, the size of *volume* created will be *size* rounded up to the integer multiple of cylinder size.

-s *size* (when using -S)

Specifies the number of blocks (decimal number) to which the size of the specified *volume* is expanded. One block is 512 bytes.

The size of the expanded *volume* will be the size rounded up to the integer multiple of the cylinder size.

-v volume (when using **-M,-R**)

The *volume* indicates the volume name that is the target of operation.

-v volume,... (when using **-F,-N**)

The *volume,...* indicates one or more volume names that is the target of the operation. To indicate multiple volumes, separate each volume name with a comma (,) as the delimiter.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.

Otherwise, a non-zero value is returned.

D.5 sdxslice - Slice operations

SYNOPSIS

```
sdxslice -F -c class -s slice,...
```

```
sdxslice -M -c class -d disk -v volume  
[-a attribute=value[,attribute=value]]
```

```
sdxslice -N -c class -s slice,...
```

```
sdxslice -R -c class {-d disk -v volume|-s slice}  
[-e delay=msec,waitsync]
```

```
sdxslice -T -c class -s slice,...
```

DESCRIPTION

Use sdxslice to perform operations on slice objects (excluding shadow slices).

You must be superuser to use this command.

PRIMARY OPTIONS

You can use one of the following options.

-F

oFffline

Stops the slice or slices specified by *slice*,... Offline slices cannot be accessed. The *slice* indicates the mirror slice name detached from the mirror volume using **-M** option.

The *class* indicates the class name to which *slice* belongs.

Attempting this command while *slice* is in use results in an error.



.....
In the case of a shared class, even offline slices will be activated upon reboot.
.....

-M

Make

Temporarily detaches one of the mirror slices used in a copy of a mirror volume specified by *volume* with mirroring multiplicity of two and higher, which is a part of the disk specified by *disk*. The *class* indicates the class name to which *volume* belongs.

Only when the physical slice attribute value of *volume* is "on", you can detach the slice. When the physical slice attribute is set to "off", you must turn it to "on" using the `sdxattr -V` command before executing this command.

Once detached, a special file is placed on the system. The path name is given below.

```
/dev/sfdsk/class/rdsk/disk.volume  
/dev/sfdsk/class/dsk/disk.volume
```

Users can access the slice with this special file. You can use this slice to create a data backup of *volume*.

If *class* is a shared class, only the node that detached the slice can access the slice. Other nodes sharing the *class* cannot access it. If you need to access from other nodes, you can take over the access right with the `-T` option.

A slice can be detached even though the *volume* is active. You must ensure the integrity of backup data at the file-system layer or database layer. If you are handling the volume as the file-system, for instance, there will be situations where you must regain integrity using the `fsck` command.



Be aware that as long as a slice is not attached using the `-R` option, the degree of mirror multiplexing stays reduced.

-N

Nline

Activates the slice or slices specified by *slice*,... Activated slices can be accessed.

The *slice* indicates the mirror slice name detached from the mirror volume using `-M` option.

The *class* indicates the class name to which *slice* belongs.

-R

Remove

Reassembles the slice as part of the volume, where the slice is specified by *slice* or combination of *disk* and *volume*.

The *disk* and *volume* combination or *slice* indicates the mirror slice name disconnected from the mirror volume using `-M` option.

The *class* indicates the class name to which the slice belongs.

After (when using `-e waitsync` option, before) returning from the `sdxslice` command, the slice is automatically reassembled with the volume. If the volume is active at this time, a synchronization copy is executed.

Attempting this command while the slice is in use results in an error.

-T

Takeover

Takes over the slice or slices specified by *slice* from another node. When the takeover is complete, the slice will stop on the original node and be activated on the current node, allowing operation to the slice on the current node.

When a slice is attached to the volume after executing this command, the entire block will be copied regardless to the setting of JRM mode.

This option is valid only for shared class.

The *slice* indicates the mirror slice name disconnected from the mirror volume using -M option.

The *class* indicates the class name to which *slice* belongs.

Attempting this command while the slice is in use results in an error.

SUB OPTIONS

Sub options are as follows:

-a *attribute=value[,attribute=value]* (when using **-M**)

Sets the *attribute* attribute of the slice to be *value*.

Both attribute values become invalid at the point when the slice is assembled with the volume.

The *attribute* indicates the attribute name, and the *value* indicates the attribute value. Always separate *attribute* and *value* with an equal (=) sign. Specifiers should be combined using a comma (,) as the delimiter.

You can use the following *attribute* and *value* pairs in any combination.

If multiple attributes are specified and an error results because of any part of them, the entire process is canceled.

jrm=on or **jrm=off** (default is **on**)

Sets the slice's JRM mode.

on

JRM is on.

off

JRM is off.

mode=rw or **mode=ro** (**rw** when omitted)

Sets the slice access mode.

rw

Sets access mode for read and write.

ro

Sets access mode for read only. Opening a read-only slice in write mode will result in an error.

-c *class*

The *class* indicates the class name to which the slice belongs.

-d *disk* (when using **-M,-R**)

The *disk* indicates the disk name to which the slice belongs, where slice is the target of the operation.

-e delay=msec (when using **-R**)

This option delays the issuing of the input/output request to the disk at the time of synchronization copying of the slice detached from volume, by milliseconds specified by *msec*.

Always separate delay and *msec* with an equal (=) sign.

This option allows you to adjust the influence on the application accessing the volume.

The value is set to 0 by default.

Values from 0 to 1000 may be specified for *msec*.

-e waitsync (when using **-R**)

When synchronization copying is executed, returns the command after the copying process is complete.

-s slice (when using **-R**)

The *slice* indicates slice that is the target of operation.

Slice name should be indicated in *disk.volume* format.

-s slice,... (when using **-F,-N,-T**)

The *slice* indicates one or more slice names that are the target of the operation. To indicate multiple slices, separate each slice name with a comma (,) as the delimiter.

Slice name should be indicated in *disk.volume* format.

-v volume (when using **-M,-R**)

Specifies the name of volume comprising the slice that is the target of the operation.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.
Otherwise, a non-zero value is returned.

D.6 sdxinfo - Display object configuration and status information

SYNOPSIS

```
sdxinfo [-ACDGSV] [-c class] [-o object] [-e long]
```

DESCRIPTION

Use *sdxinfo* to display configuration and status information of the objects which the current node shares.
The *sdxinfo* command allows you to view configurations and status information of SDX objects, proxy objects and shadow objects.

PRIMARY OPTIONS

Primary options indicate the type of objects to display. If nothing is specified, only information on the pertinent object will be displayed.

Object types can be any combination of the following.

-A

All

Displays all specified objects, and information on all related objects. This is equivalent to -CDGSV. If any other options are combined with this, they are ignored.

-C

Class

Displays all specified objects, and class information of the related objects.

-D

Disk

Displays all specified objects, and disk information of the related objects.

-G

Group

Displays all specified objects, and group information of the related objects.

-S

Slice

Displays all specified objects, and slice information of the related objects.

-V

Volume

Displays all specified objects, and volume information of the related objects.

SUB OPTIONS

Sub options are used to specify the names of objects to display. If nothing is specified, the command is interpreted as if all objects at the current node had been specified.

-c class

The *class* indicates the class name whose information will be displayed. If this option is omitted, this command is interpreted as if all classes had been specified.

In case of combining this option with the -o option, the objects related to the specified *object* within the *class* is displayed.

-e long

Displays detailed object information.

-o object

The *object* indicates the object name (class name, disk name, group name or volume name) whose information will be displayed. If this option is omitted, this command is interpreted as if all object names had been specified.

In case of combining this option with the -c option, the objects related to the specified *object* within the *class* is displayed.

DISPLAYED INFORMATION

Information displayed in response to the `sdxinfo` command can be interpreted as follows.

Class information:

OBJ

Displays class as an object classification.

NAME

Displays the class name.

TYPE

Displays one of the following values.

root

root class

local

local class

shared

shared class

SCOPE

Displays the node names as scope attribute values. In a PRIMECLUSTER system, "(local)" is displayed for the root class, the node identifier (CF node name) is displayed for a local class, and node identifiers separated by colons (:) are displayed for a shared class.

HS

When the `-e` long option is used, this displays hot spare operation, which can be either of the following.

on

Enables the hot spare.

off

Disables the hot spare. Be aware that spare disk automatic connection is prevented.

For a shadow class it always displays "on", but the hot spare feature is practically invalid since a shadow class cannot include a spare disk.

SPARE

Displays the number of spare disks that are not connected to the group.

SHADOW

When the -e long option is specified, one of the following is displayed as the class description.

0

Class created with the sdxdisk -M command.

1

Shadow class created with the sdxshadowdisk -M command.

HSMODE

Displays one of the following values to indicate the spare disk selection method for automatic connection by hot spare when the -e long option is specified. For a shadow class, an asterisk (*) is displayed.

exbox

External mode. If an I/O error occurs in a disk of a disk array unit, this method selects a spare disk that belongs to a different disk case from that of the failed disk. If an I/O error occurs in a disk irrelevant to a disk array unit (such as an internal disk), it selects a spare disk that is connected to a different controller from that of the failed disk. When no applicable unconnected spare disk is found there, a spare disk that belongs to the same disk case or is connected to the same controller as that of the disk with the I/O error, is selected.

bybox

Internal mode. If an I/O error occurs in a disk of a disk array unit, this method selects a spare disk that belongs to the same disk case as that of the failed disk. If an I/O error occurs in a disk irrelevant to a disk array unit (such as an internal disk), it selects a spare disk that is connected to the same controller as that of the failed disk. When no applicable unconnected spare disk is found there, spare disk automatic connection is prevented.

Disk information:

OBJ

Displays disk as an object classification.

NAME

Displays the disk name.

TYPE

The type attribute value can be any of the following:

mirror

Mirror. It is connected to a mirror group.

stripe

Stripe. It is connected to a stripe group.

concat

Concatenation. It is connected to a concatenation group.

switch

Switch. It is connected to a switch group.

keep

Keep. When it is registered with a class or connected to a group, the format and data of the disk are retained.

single

Single. Single volumes can be created on it.

spare

Spare. "spare" is also displayed when it is connected to a group.

undef

Undefined. Its usage is not determined yet.

CLASS

Displays the class name to which the disk belongs.

GROUP

Displays the group name to which the disk is connected. If the disk is not connected to any group, an asterisk (*) is displayed.

DEVNAM

Displays the physical disk name in either the following formats. If the disk is not connected to the current node, an asterisk (*) is displayed.

```
cCt TdD    (for normal hard disks)
mphd/      (for MPHD disks)
mplb/      (for MPLB disks)
emcpowerN (for emcpower disks)
```

where *C* is the controller number, *T* is the target ID, *D* is the disk number, *I* is the MPHD instance number or the MPLB instance number, and *N* is the emcpower device number.

DEVBLKS

Displays the size of the physical disk. The size is given in blocks (sectors.) If the disk is not connected to the current node, 0 (zero) is displayed.

FREEBLKS

When the *-e* long option is used, displays the number of free allocatable blocks (sectors) for a new volume. If the disk is not a single disk, an asterisk (*) is displayed.

DEVCONNECT

Displays a list of node identifiers of the nodes to which the physical disk is connected, separated using colons (:) as delimiters. If PRIMECLUSTER has not been installed or the physical disk has not been registered in the PRIMECLUSTER resource database, an asterisk (*) is displayed.

STATUS

Disk status can be any of the following:

ENABLE

Available for work.

DISABLE

Not available for work.

SWAP

Not available for work, but available for disk exchanging.

E

When `-e` long option is used, displays error status of disk, which can be any of the following.

0

I/O error has not occurred on both the current node and the shared node.

1

I/O error has occurred on either the current node or the shared node.

Group information:

OBJ

Displays group as an object classification.

NAME

Displays the group name.

CLASS

Displays the class name to which the group belongs.

DISKS

Displays the name of disks or lower level groups that are connected to the group, separated using colons (:) as delimiters. In case of a stripe group, names are listed in the order they are striped. Likewise, in case of a concatenation group, names are listed in the order they are concatenated.

BLKS

Displays the size of the group, which is the total of available size of group (available capacity as volumes) plus one private slice size. Size is given in blocks (sectors).

FREEBLKS

The number of free allocatable blocks for a new volume. If the group is a lower level group, an asterisk (*) is displayed.

SPARE

Displays the number of spare disks that can be connected to the group. Unless it is a mirror group, an asterisk (*) is displayed.

MASTER

When the -e long option is used, it displays the group name of master group. When it is not a proxy group, an asterisk (*) is displayed.

TYPE

When -e long option is used, displays type attribute value, which can be any of the following.

mirror

Mirror group.

stripe

Stripe group.

concat

Concatenation group.

switch

Switch group

WIDTH

When the -e long option is used, displays stripe width in blocks (sectors). If the group is not a stripe group, an asterisk (*) is displayed.

ACTDISK

Displays the disk name of the active disk when the -e long option is specified. Unless it is a switch group, an asterisk (*) is displayed.

Volume information:**OBJ**

Displays volume as an object classification.

NAME

Displays the name of the volume. If it is an area that cannot be allocated (area for private slice) or that can be allocated but have not (unallocated area), an asterisk (*) is displayed.

TYPE

When the -e long option is used, displays type attribute, which can be any of the following.

mirror

Mirror. It belongs to a mirror group.

stripe

Stripe. It belongs to a stripe group.

concat

Concatenation. It belongs to a concatenation group.

switch

Switch. It belongs to a switch group.

single

Single. It belongs to a single disk.

CLASS

Displays the class name to which the volume belongs.

GROUP

Displays the highest level group name to which the volume belongs. When it belongs to a single disk, an asterisk (*) is displayed.

DISK

When the -e long option is used, displays the name of the single disk to which the volume belongs. When it belongs to a group, an asterisk (*) is displayed.

MASTER

When the -e long option is used, it displays the volume name of master volume. When it is not a proxy volume, an asterisk (*) is displayed.

PROXY

When the -e long option is used, it displays the proxy volume status in one of the two ways as given below. When it is not a proxy volume, an asterisk (*) is displayed.

Join

The volume is being joined to a master volume.

Part

The volume is being parted from a master volume.

SKIP

Displays the skip-resynchronization mode setting, which can be either of the following. If the volume is neither a mirror volume nor a single volume, an asterisk (*) is displayed.

on

Skip resynchronization.

off

Execute resynchronization.

Note

Note that the interface for setting or changing this option is not available.

JRM

Displays the just resynchronization mode setting, which can be either on or off. If the volume is neither a mirror volume nor a single volume, an asterisk (*) is displayed.

on

JRM is on.

off

JRM is off.

MODE

When the `-e` long option is used, it displays the access mode attribute value (default access mode) for the current node, which can be either of the following. If it is either an area for private slice or an unallocated area, an asterisk (*) is displayed.

rw

Read and write mode.

ro

Read only mode.

CMODE

When `-e` long option is used, it displays the present access mode of the activated volume from the current node. If the volume is not activated, an asterisk (*) is displayed.

rw

Read and write mode.

ro

Read only mode.

LOCK

When the `-e` long option is used, displays the lock mode of current node, which can be either of the following. If it is either a private area or an unallocated area, an asterisk (*) is displayed.

on

The volume is locked from activating thereafter.

off

The volume is not locked from activating thereafter.

1STBLK

Displays the block (sector) number of the first block. The block number is the logical block number, which is the offset in the group to which the volume belongs, and not the physical block number indicating the offset on the physical disk. However, when the volume belongs to a single disk, the block number will match the physical block number on the single disk. Also, when it belongs to a mirror group to which a disk is directly connected or a switch group, the block number will match the physical block number on the disk.

LASTBLK

Displays the block (sector) number of the last block. The block number is the logical block number, which is the offset in the group to which the volume belongs, and not the physical block number indicating the offset on the physical disk. However, when the volume belongs to a single disk, the block number will match the physical block number on the single disk. Also, when it belongs to a mirror group to which a disk is directly connected or a switch group, the block number will match the physical block number on the disk.

BLOCKS

Displays the size in blocks (sectors).

STATUS

Displays the volume status of the current node, which can be any of the following.

ACTIVE

Ready for work.

STOP

Stopped.

INVALID

Stopped, and cannot be activated due to problem with data.

FREE

Not yet allocated as a volume.

PRIVATE

An area reserved for GDS control, so cannot be allocated as a volume.

PSLICE

When the -e long option is used, displays the physical slice attribute value, which can be either of the following. If it is either a private area or an unallocated area, an asterisk (*) is displayed.

on

Physical slice attribute of the volume is set to "on." Among slices comprising volumes, slices on single disks, on disks connected to switch groups and on disks directly connected to mirror groups are registered with VTOC and have physical slices. If a lower level group is the only object directly connected to mirror group, the volume will not have a physical slice, regardless to this attribute being set to "on." Also, when the volume belongs to either a stripe group or a concatenation group, this attribute value will never

be set to "on."

off

Physical slice attribute of the volume is set to "off." The volume has no physical slices, and none of the slices in the volume is registered to the VTOC.

For a shadow volume it always displays off no matter whether the shadow slice is registered with VTOC.

SNUM

When the -e long option is specified, the slice number of the slice configuring the volume is displayed. If the physical slice attribute is off or no physical slice configures the volume, an asterisk (*) is displayed.

PJRM

When the -e long option is specified, either of the following values is displayed to indicate the just resynchronization mechanism mode on the proxy volume. If it is not a parted proxy volume, an asterisk (*) is displayed.

on

JRM is "on".

off

JRM is "off".

Slice information:

OBJ

Displays slice as an object classification.

NAME

When -e long option is used, display the name of slice. When the slice is not a mirror slice that is temporarily detached from the mirror volume using the sdxslice -M command, an asterisk (*) is displayed.

CLASS

Displays the class name to which the slice belongs.

GROUP

Displays the highest level group name to which the slice belongs.
If it is a single slice, an asterisk (*) is displayed.

DISK

Displays the name of the disk or the lower level group (i.e. the group to which this slice belongs, among the groups that are directly connected to the relevant highest level group) to which the slice belongs. If the highest level group is a switch group, the disk name of the active disk is displayed. If the highest level group is a stripe group or a concatenation group, an asterisk (*) is displayed.

VOLUME

Displays the volume name to which the slice belongs.

JRM

When the `-e` long option is used, displays the just resynchronization mode setting, which can be either on or off.

When the slice is not a mirror slice that is temporarily detached from the mirror volume using the `sdxslice -M` command, an asterisk (*) is displayed.

on

JRM is on.

off

JRM is off.

MODE

When `-e` long option is used, this displays the access mode, which can be either of the following. When the slice is not a mirror slice that is temporarily detached from the mirror volume using the `sdxslice -M` command, an asterisk (*) is displayed.

rw

Read and write mode.

ro

Read only mode.

STATUS

Displays the slice status on the current node, which can be any of the following.

ACTIVE

Ready for work.

STOP

Stopped.

INVALID

Due to a problem with data, has been temporarily detached from the volume.

COPY

Specifies a copy is underway, to maintain data uniformity.

TEMP

Temporarily detached from volume. Slice is operating in isolation.

TEMP-STOP

Temporarily detached from volume. Slice is stopping in isolation.

NOUSE

Stopped, with no operations possible.

COPY

When the `-e` long option is used, this displays one of the following copying process statuses. When the slice is not in COPY status, an asterisk (*) is displayed.

run

Copying is underway.

bg

Copying is in process in the background, but you can access valid data.

intr

Copying has been interrupted. Executing the `sdxcopy -I` command interrupts the copy.

wait

Since many copying processes are in progress, it is waiting to be scheduled.

CURBLKS

When the `-e` long option is used, this displays the number of blocks (sectors) that have already been copied. When CURBLKS and the later described COPYBLKS match, all copying has been completed. When the slice is not in COPY status, or when copying is in process through TimeFinder or SRDF by use of GDS Snapshot, an asterisk (*) is displayed.

COPYBLKS

When the `-e` long option is used, this displays the number of blocks (sectors) that needs to be copied. Usually this is the same size as the volume size it is registered with, but when just resynchronization is in process, the number of blocks that actually needs to be copied will be displayed. When the slice is not in COPY status, an asterisk (*) is displayed.

DLY

When the `-e` long option is used, this displays the copy delay time in milliseconds.
When not in COPY status, an asterisk (*) is displayed.

CPTYPE

When the `-e` long option is used, one of the following values is displayed as copy function which is used for copying process between the master and the proxy. When copying is not in process, or if the volume to which the slice belongs is not a target volume of copying process between a master volume and a proxy volume, an asterisk (*) is displayed.

soft

Copying is in process using the soft copy function provided by the GDS `sfdsd` driver.

EC

Copying is in process using the Equivalent Copy function.

OPC

Copying is in process using the One Point Copy function. When the master and the proxy have been joined, and the copy source volume is active, the soft copy function may be involved in part of the copying process.

REC

Copying is in process using the Remote Equivalent Copy function.

ROPC

Copying is in process using the Remote One Point Copy function. When the master and the proxy have been joined, and the copy source volume is active, the soft copy function may be involved in part of the copying process.

TF

Copying is in process using EMC TimeFinder.

SRDF

Copying is in process using EMC SRDF.

CPSOURCE

When the -e long option is used, the volume name of source proxy volume which is used to restore data is displayed. When the volume to which the slice belongs is not a target master volume of restoring process, an asterisk (*) is displayed.

FUNC

Displays one of the following values to indicate the session type of the disk unit's copy function when the -e long option is specified. If there is no session, an asterisk (*) is displayed.

EC

Source or target of the Equivalent Copy session.

OPC

Source or target of the One Point Copy session.

REC

Source or target of the Remote Equivalent Copy session.

ROPC

Source or target of the Remote One Point Copy session.

TF

Source or target of a BCV pair with EMC TimeFinder.

SRDF

Source or target of an SRDF pair with EMC SRDF.

CPSTAT

Displays one of the following values to indicate the session status of the disk unit's copy function when the -e long option is specified. If there is no session, an asterisk (*) is displayed.

equiv

Synchronized.

copy

In process of copying.

suspend

EC or REC session suspended.

split

BCV pair or SRDF pair split.

error

Suspended due to an error.

halt

Hardware suspended.

PARTNER

When the -e long option is specified, displays a destination slice name if the slice is the source in the session of the disk unit's copy function, or a source slice name if it is the destination. If there is no session, an asterisk (*) is displayed.

RETURNED VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a non-zero value is returned.

USAGE EXAMPLES

Displays all the objects within the current node.

```
sdxinfo
```

Displays information on all the disks registered with the class called "Class1."

```
sdxinfo -A -c Class1
```

Use this to check whether an object called "foo" is currently being used.

```
sdxinfo -o foo
```

CAUTION

Additional information may be displayed in accordance with new functions provided.

D.7 sdxattr - Set objects attributes

SYNOPSIS

```
sdxattr -C -c class -a attribute=value[,attribute=value,...]
```

```
sdxattr -D -c class -d disk -a attribute=value[,attribute=value]
```

```
sdxattr -G -c class -g group -a attribute=value[,attribute=value]
```

```
sdxattr -S -c class -s slice -a attribute=value[,attribute=value]
```

```
sdxattr -V -c class -v volume -a attribute=value[,attribute=value,...]
```

DESCRIPTION

Use sdxattr to change attribute values of objects (excluding shadow objects) on the current node.

You must be superuser to use this command.

PRIMARY OPTIONS

Primary options are used to specify the category of the object whose attributes are to be set.

-C

Class

Set the attributes of the class specified by *class*.

-D

Disk

Set the attributes of the disk specified by *disk*. The *class* indicates the class name with which *disk* is registered.

If *disk* is connected to a group, or if there is a volume within *disk*, it will result in an error, and the change in attribute will not be executed.

-G

Group

Set the attributes of the group specified by *group*. The *class* indicates the class name to which *group* belongs.

When *group* is connected to another group, this option will result in an error and you cannot change the attribute. Also, this option will result in an error if there are one or more activated volumes within the group. Stop all volumes before executing this command.

-S

Slice

Set the attributes of the slice specified by *slice*. The *class* indicates the class name to which the *slice* belongs.

-V

Volume

Set the attributes of the volume specified by *volume*. The *class* indicates the class name to which the *volume* belongs.

SUB OPTIONS

Sub options are as follows:

-a *attribute=value[,attribute=value,...]* (when using **-C**)

Sets the *attribute* attribute of the *class* to be *value*.

The *attribute* indicates the attribute name, and the *value* indicates the attribute value.

Always separate *attribute* and *value* with an equal (=) sign. Specifiers should be combined using a comma (,) as the delimiter.

You can use the following *attribute* and *value* pairs in any combination.

When attempting to set multiple attributes, if any of them result in an error, no attributes are set.

type=local or **type=shared**

Sets the class type attribute.

When there is an activated volume within *class*, you can change the type attribute from "local" to "shared," but not vice versa. In order to change it from "shared" to "local," you must execute the command after first stopping all volumes.

You cannot change it from "root" to "local" or "shared", or vice versa. Changing the type of a class fails if the class is a shared class that includes a disk of which physical scope includes 1 node, or that includes a switch group.

local

Sets the type attribute to "local."

Objects within "local" type class can only be used by current node.

For the scope attribute, current node's node identifier will be automatically set.

shared

Sets the type attribute to "shared."

By combining this with the scope attribute, objects in the class can be shared among multiple nodes including the current node.

scope=node:node:...

For a "shared" type class, changes the nodes which share the class.

When there is an activated volume within *class*, you can add new nodes, but you cannot remove a node that has already been included in the scope. In order to remove a node, you must execute the command after first stopping all volumes.

If the *node* settings are not all complete, this will result in an error.

The *node* indicates a node identifier that is defined by PRIMECLUSTER.

For scope, you can indicate up to sixteen nodes.

Changing the scope of a class fails if the class is a shared class that includes a disk of which physical scope is 1 node, or that includes a switch group.

hs=on or hs=off

It sets the operation of the hot spare.

You can make changes regardless to whether there is an activated volume within *class*.

on

Enables the hot spare.

off

Disables the hot spare. Spare disk automatic connection is restricted.

hsmode=exbox or hsmode=bybox

Changes the spare disk selection mode for automatic connection by hot spare.

This operation is available regardless whether or not there are active volumes within *class*.

exbox

Changes the spare disk selection method to the external mode. If an I/O error occurs in a disk of a disk array unit, this method selects a spare disk that belongs to a different disk case from that of the failed disk. If an I/O error occurs in a disk irrelevant to a disk array unit (such as an internal disk), it selects a spare disk that is connected to a different controller from that of the failed disk. When no applicable unconnected spare disk is found there, a spare disk that belongs to the same disk case or is connected to the same controller as that of the disk with the I/O error, is selected.

bybox

Changes the spare disk selection mode to the internal mode. If an I/O error occurs in a disk of a disk array unit, this method selects a spare disk that belongs to the same disk case as that of the failed disk. If an I/O error occurs in a disk irrelevant to a disk array unit (such as an internal disk), it selects a spare disk that is connected to the same controller as that of the failed disk. When no applicable unconnected spare disk is found there, spare disk automatic connection is restrained.

name=classname

Sets the class name to *classname*.

When there is an activated volume within *class*, it results in an error. First, you must stop all volumes.

When changing a class name through this operation, the paths of special files for volumes are also changed, so you must update the files in which the paths are described, such as */etc/vfstab*.

```
/dev/sfdsk/classname/rdisk/volume_name  
/dev/sfdsk/classname/disk/volume_name
```

-a attribute=value[,attribute=value] (when using -D)

Sets the *attribute* attribute of the *disk* to be *value*.

The *attribute* indicates the attribute name, and the *value* indicates the attribute value. Always separate *attribute* and *value* with an equal (=) sign. Specifiers should be combined using a comma (,) as the delimiter.

You can use the following *attribute* and *value* pairs in any combination.

When attempting to set multiple attributes, if any of them result in an error, no attributes are set.

type=keep, type=single, type=spare or type=undef

Sets the SDX disk type attribute. If *disk* is not connected to part of nodes included in the scope of *class*, or if a switch group exists in *class*, changing the type attribute of *disk* fails.

keep

Sets the type attribute to "keep."

Hereinafter, this disk will be regarded as a keep disk, and its format and data will remain when connected to a group.

You cannot change a single disk to keep disk.

single

Sets the type attribute to "single."

Single volume may be created within *disk* thereafter. You cannot change a keep disk to a single disk.

spare

Sets the type attribute to "spare."

The *disk* will be used as a spare disk thereafter.

When the available size of *disk* is smaller than the available size of the smallest mirror group within *class*, a warning message notifying that the hot spare function may not be available will be sent to standard error output.

undef

Sets the type attribute to "undef."

Hereinafter, this *disk* will be regarded as an undefined disk, which use is not yet determined.

name=*diskname*

Sets the name of a disk to *diskname*.

-a *attribute=value*[,*attribute=value*] (when using -G)

Sets the *attribute* attribute of the *group* to be *value*.

The *attribute* indicates the attribute name, and the *value* indicates the attribute value. Always separate *attribute* and *value* with an equal (=) sign. To set multiple attributes, specify sets of these specifiers in comma-delimited format.

Specify any of the following sets into *attribute* and *value*.

When multiple attributes are specified, the entire process is canceled in the event of an error in part of the processes.

name=*groupname*

Sets the name of the group to *groupname*.

actdisk=*disk*

Changes the active disk of the switch group specified by *group* to *disk*.

-a *attribute=value* (when using -S)

Sets the *attribute* attribute of the detached *slice* to be *value*. Both attribute values become invalid at the point when the slice is assembled with the volume using the -R option.

The *attribute* indicates the attribute name, and the *value* indicates the attribute value. Always separate *attribute* and *value* with an equal (=) sign. Specifiers should be combined using a comma (,) as the delimiter.

You can use the following *attribute* and *value* pairs in any combination.

When attempting to set multiple attributes, if any of them result in an error, no attributes are set.

jrm=off

Turns the slice's just resynchronization mechanism mode to "off".

It could be turned "off" regardless to the slice status.

To set the jrm to "on," attach the slice to the volume and then detach it again.

mode=rw or mode=ro

Changes the access mode of current node for slice.

It will result in an error when *slice* is activated. Execute after stopping it.

rw

Sets access mode for read and write.

ro

Sets access mode for read only. Opening a read-only volume in write mode will result in an error.

-a attribute=value[,attribute=value,...] (when using **-V**)

Sets the *attribute* attribute of the *volume* to be *value*.

The *attribute* indicates the attribute name, and the *value* indicates the attribute value. Always separate *attribute* and *value* with an equal (=) sign. Specifiers should be combined using comma (,) as the delimiter.

You can use the following *attribute* and *value* pairs in any combination.

When attempting to set multiple attributes, if any of them result in an error, no attributes are set.

jrm=on or jrm=off

Turn the JRM mode "on" or "off".

You can make changes regardless to whether *volume* is activated or not.

If *volume* belongs to a group other than a mirror group, this command results in an error.

on

JRM is "on".

off

JRM is "off".

lock=on or lock=off

Changes the lock mode of current node for *volume*.

You can make changes regardless to whether *volume* is activated or not.

If *class* is the root class, this command results in an error.

on

The volume is locked from activating thereafter.

off

The volume is not locked from activating thereafter.

mode=rw or mode=ro

Changes the access mode of current node for *volume*.

When *volume* is activated, it results in an error. First, you must stop the *volume*.

rw

Sets access mode for read and write.

ro

Sets access mode for read only.

Opening a read-only volume in write mode will result in an error.

name=volumentname

Sets the name of the volume to *volumentname*.

It will result in an error when *volume* is activated. Execute after stopping *volume*.

pjrm=off

The parted proxy volume's just resynchronization mechanism mode for rejoining or restoring is turned "off".

This can be changed regardless of whether *volume* is active or inactive.

This attribute value becomes invalid as *volume* is rejoined to the master volume with the Rejoin or RejoinRestore option of the *sdxproxy* command.

The value cannot be set to "on". To turn "on", the volume must be rejoined to the master volume once and then be parted again.

pslice=on or pslice=off

Turns the physical slice attribute value to be "on" or "off", respectively.

If *volume* is activated, to avoid an error, stop *volume* before indicating this option. This option will also result in an error if there is a detached slice within *volume*. In such case, attach the slice before indicating this option.

on

The physical slice attribute value of volume is set to be "on". Among the slices consisting *volume*, any slice on a single disk, and any slices on disks that are directly connected to a mirror group will be registered to the VTOC.

You cannot change this option to "on" when *volume* belongs to a group that cannot create a physical slice (stripe group, concatenation group, or a mirror group whose only directly-connecting group is a lower level group), or when there is a maximum number (six) of volumes with its physical slice attribute set to "on" within the same group or single disk.

off

The physical slice attribute value of volume is set to be "off".
If *class* is the root type, the value cannot be changed to "off."

-c class

The *class* indicates the class name to which the object belongs, or is to be changed, where object is the target of the change.

-d disk

The *disk* indicates the disk name that is the target of the change.

-g group

The *group* indicates the group name that is the target of the change.

-s slice

The *slice* indicates the slice name that is the target of the change.
Slice name should be specified in *disk.volume* format.

-v volume

The *volume* indicates the volume name that is the target of the change.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.
Otherwise, a non-zero value is returned.

D.8 sdxswap - Swap disk

SYNOPSIS

```
sdxswap -I -c class -d disk [-e delay=msec,nowaitsync]
```

```
sdxswap -O -c class -d disk
```

DESCRIPTION

Use sdxswap to make a disk (excluding a shadow disk) registered with GDS exchangeable and to restore the disk after swapping.
You must be superuser to use this command.

This command is primarily used for swapping faulty disks.

PRIMARY OPTIONS

You can use either of the following options.

-I

swapIn

Returns the disk (specified by *disk*) to a usable state, and restores its original status and configuration. You must execute the command after a crashed disk has been swapped.

The *disk* indicates the disk name that was made exchangeable with the -O option. The *class* indicates the class name with which the *disk* is registered.

The physical disk size of *disk* must be equal to or larger than the original physical disk size.

When the highest level group of *disk* is a mirror group, slice configuration or volume contents is copied as needed, and returns from the `sdxcswap` command once the copying is complete (when using `-e nowaitsync` option, before the copying process).

If a spare disk is substituted in place of *disk*, then once the redundancy of all related volumes has been restored, the spare disk is disconnected.

-O

swapOut

Detaches all slices within the disk (specified by *disk*), and sets the disk to exchangeable status. This command must be executed before swapping a faulty disk.

The *class* indicates the class name with which the *disk* is registered.

If *disk* is already nonusable, the status is rechecked and a warning message is sent to standard error output.

The following details explain unexchangeable and exchangeable conditions when *disk* is not connected to a group and when the highest level group of *disk* is a mirror group, a stripe group, a concatenation group or a switch group.

When *disk* is not connected to a group

When volumes exist in *disk*, the *disk* cannot be made exchangeable.

When the highest level group of *disk* is a mirror group

When volumes exist in the highest level mirror group of *disk* and detaching slices within the *disk* can change the volume configurations and statuses, the *disk* cannot be made exchangeable.

For example, if there are volumes in the highest level mirror group of *disk*, and if only the disk specified by *disk* is connected to that group, detaching slices within the *disk* will change the configurations and statuses of the volumes. Therefore, the *disk* cannot be made exchangeable.

When the highest level group of *disk* is a stripe group

When the highest level group of *disk* is a stripe group, the *disk* cannot be made exchangeable by detaching slices within the *disk*.

When the highest level group of *disk* is a concatenation group

When *disk* is a disk connected to the highest level concatenation group or the active disk connected to a lower level switch group, it cannot be made exchangeable no matter whether or not volumes exist.

When *disk* is the inactive disk connected to a lower level switch group, it can be made exchangeable no matter whether or not volumes exist.

When the highest level group of *disk* is a switch group

When *disk* is the inactive disk, the *disk* can be made exchangeable regardless whether or not there are volumes.

When *disk* is the active disk, the *disk* can be made exchangeable by detaching slices within the *disk* only if the switch group includes no switch volume or connected inactive disk.

When the switch group includes volumes, remove those volumes in order to make the active disk exchangeable. When it includes the connected inactive disk, perform an active disk switch with the `sdxattr -G` command and then make the former active disk exchangeable.

SUB OPTIONS

Sub options are as follows:

-c class

The *class* indicates the class name to which *disk* belongs, where *disk* is the target of the operation.

-d disk

The *disk* indicates the disk name that is the target of the operation.

-e delay=msec (when using -l)

When restoring the disk, data contained in the volume is copied as needed. This option delays the issuing of the input/output request to the disk at the time of copying, by milliseconds specified by *msec*.

This option allows you to adjust the influence on the application accessing the volume.

The value is set to 0 by default.

Values from 0 to 1000 may be specified for *msec*.

-e nowaitsync (when using -l)

Returns the command before the copying is complete.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.
Otherwise, a non-zero value is returned.

D.9 sdxfix - Restore a failed object

SYNOPSIS

```
sdxfix -C -c class
```

```
sdxfix -D -c class -d disk [-e online]
```

```
sdxfix -V -c class {-g group|-d disk} -v volume
```

DESCRIPTION

Use `sdxfix` to restore failed objects (excluding shadow objects). In such a case, you need to restore integrity from backup data or check by executing the "fsck" command. The "sdxfix" command can be executed with superuser access privileges only.

PRIMARY OPTIONS

You can use one of the following options.

-C

Class

Restores closed class to a normal status on the current node.

This command can restore *class* when the *class* includes a configuration database normally accessible and:

- less than 3 disks in ENABLE status and 1 or more disks normally accessible
- three to 5 disks in ENABLE status and 2 or more disks normally accessible
- six or more disks in ENABLE status and 3 or more disks normally accessible

After restoration is complete, objects within the *class* are restored to the previous status. However, if *class* is a local class, volumes that were in STOP status before the class closure will be in ACTIVE status. Additionally, if *class* is a shared class, volumes that were in ACTIVE status before the class closure will be in STOP status.

-D

Disk

Restores the state of disk *disk* that detected an I/O error.

Reads all *disk* area, and if there's no problem, then clears the error status.

If there's a volume in the highest-level group to which *disk* belongs, or *disk* has a single volume, the volume should be stopped or inactive (STOP or INVALID) on all nodes (except when -e online is specified).

When *disk* is connected to a switch group, restoration fails. To clear an I/O error in a disk connected to a switch group, use the `sdxswap -O` command to make the disk exchangeable and then use the `sdxswap -I` command to make the disk useable.

-V

Volume

Restores a slice with invalid data (INVALID) or a not-ready slice (NOUSE) specified by a set of *disk* and *volume* or by a set of *group* and *volume* to the STOP status to restore the *volume* with invalid data (INVALID) to the STOP status.

The *volume* should be stopped or inactive (STOP or INVALID) on all nodes. The slice state specified in the combination of *disk* and *volume*, or *group* and *volume* should be INVALID or NOUSE.

Reads the entire slice specified in the combination of *disk* and *volume*, or *group* and *volume*, and if there's no problem, changes the state to STOP, and then, changes the state of stopped slice to INVALID.

SUB OPTIONS

You can use the following sub-options.

-c *class*

Specify a name of the class to which the object belongs.

-d *disk* (when using -D)

Specify a name of the disk.

-d *disk* (when using **-V**)

When *volume* is a mirror volume, specify a name of the disk that is connected to the mirror group to which the *volume* belongs. This disk should have the INVALID mirror slice that needs to be restored to STOP.

Specify a single disk name when *volume* is a single volume.

When *volume* is a switch volume, specify a disk name of the active disk connected to a switch group that includes the *volume* into *disk*. Do not set an inactive disk name to *disk*.

-e online (when using **-D**)

Restores the object even when the highest level group to which *disk* belongs or the single disk specified by *disk* includes an active volume.

-g *group* (when using **-V**)

When *volume* is a mirror volume, specify a name of lower-level group that is connected to the mirror group to which the *volume* belongs. This *group* should have the INVALID mirror slice that needs to be restored to STOP.

Specify a name of the highest-level stripe group when *volume* is a stripe volume.

Specify a name of the highest-level concatenation group when *volume* belongs to the highest-level concatenation group.

-v *volume* (when using **-V**)

Specify a name of the volume.

RETURNED VALUE

When it is normally terminated, "0" is returned. Otherwise, a non-zero value is returned.

D.10 **sdxcopy** - Synchronization copying operation

SYNOPSIS

```
sdxcopy -B -c class -v volume, ... [-e delay=msec,nowaitsync]
```

```
sdxcopy -C -c class -v volume, ...
```

```
sdxcopy -I -c class -v volume, ...
```

```
sdxcopy -P -c class -v volume, ... -e delay=msec
```

DESCRIPTION

Use **sdxcopy** to access to synchronization copying for volume objects (excluding shadow volumes) specified by *volume*. You must be superuser to use this command.

PRIMARY OPTIONS

You can use one of the following options.

-B

Begin

Attaches slices that have been detached from mirror volumes specified by *volume*,... and executes synchronization copying. The command returns control after the synchronization copying is complete (before the copying process starts when using the *-e nowaitsync* option). *class* indicates the class name to which the *volume* belongs.

Slices on the *volume* that have a status of INVALID are attached, and then a synchronization copying is executed.

Slices with a status of TEMP* or NOUSE are not attached.

If there are any slices currently involved in a synchronization copy on the *volume*, this command will terminate with an error.

This command resumes copying from the point where it was interrupted with the *-I* option.

Synchronization copying will be executed while the *volume* is activated or stopped.

-C

Cancel

Cancels synchronization copying in process or interrupted on the mirror volume or mirror volumes specified by *volume*,.... After the cancel is completed, the system returns from the command. The *class* indicates the class name to which *volume* belongs.

-B option executes the copying process again.

-I

Interrupt

Interrupts the synchronization copying currently underway on a mirror volume or mirror volumes specified by *volume*,.... The command returns after the interrupt is completed. The *class* indicates the class name to which *volume* belongs.

-B option executes the copying process from the point where it was interrupted.

-P

Parameter

Changes parameter related to the synchronization copying in process or interrupted on the mirror volume or mirror volumes specified by *volume*,.... The *class* indicates the class name to which *volume* belongs.

Copying in process will resume after the parameter is changed.

The current status of synchronization copying that is either in process or interrupted can be checked by executing the *sdxinfo -S* command.

SUB OPTIONS

Sub options are as follows:

-c class

The *class* indicates the class name to which *volume* belongs.

-e delay=msec (when using -B,-P)

Delays the issuing of the input/output request to the disk at the time of copying by milliseconds specified by *msec*. This option allows you to adjust the influence on the application accessing the volume.

The value is set to 0 by default. If the copying process is either completed or canceled, the delay time will return to default (0).

Values from 0 to 1000 may be specified for *msec*.

-e nowaitsync (when using **-B**)

Returns the command before the copying is complete.

-v volume,...

The *volume* indicates the volume name that is the target of the operation.

To indicate multiple *volumes*, separate each volume name with a comma (,) as the delimiter.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.

Otherwise, a non-zero value is returned.

D.11 sdxroot - Root file system mirroring definition and cancellation

SYNOPSIS

```
sdxroot -M -c class -d disk[, disk, ...]
```

```
sdxroot -R -c class -d disk[, disk, ...]
```

DESCRIPTION

Use `sdxroot` to complete or cancel mirroring definition of system disks including root file systems. The "`sdxroot`" command can be executed with superuser access privileges only.

System disk means the physical disk on which the running Solaris is installed. Specifically, it means the entire disk that includes a slice currently running as any one of the following file systems (or a swap area).

`/`, `/usr`, `/var`, or `swap`

PRIMARY OPTIONS

You can use one of the following options.

-M

Make

Checks that one or more system disks specified by *disk* are ready for mirroring (registered with a class and connected to a group) and creates remaining mirroring definitions (such as updates of files `/etc/vfstab` and `/etc/system`). After returning from this command, reboot the system immediately. After the system is rebooted, system disk mirroring will commence.

Specify a disk that includes a slice currently running as `/(root)`, `/usr`, and `/var` or a swap area for *disk*. Among `/(root)`, `/usr`, and `/var`, and swap areas, although it is not required to specify disks with swap areas only, be sure to specify disks with `/(root)`, `/usr`, and `/var`. Additionally, the disk specified by *disk* must be ready for mirroring (registered with a class and connected to a group).

When synchronization copying is being performed in groups to which the system disks are connected, the `sdxroot` command results in an error. In this situation, cancel the synchronization copying using the `sdxcopy -C` command, or execute this command after the synchronization copying is completed.

-R

Remove

Checks that one or more system disks specified by *disk* are ready for mirroring cancellation (disks are disconnected from the system disk's groups and only one system disk is connected to each group) and creates remaining mirroring cancellation definitions (such as update of the `/etc/vfstab`, `/etc/system` file). After returning from this command, reboot the system immediately. After the system is rebooted, system disk mirroring will be cancelled.

To totally cancel system disk management with GDS, after the system is rebooted, it is necessary to delete system disk related volumes, groups, disks, and classes.

Specify a disk that includes a slice currently running as `/`(root), `/usr`, and `/var`, or a swap area for *disk*. The disk specified by *disk* must be ready for mirroring cancellation (disks are disconnected from the system disk's groups and only one system disk is connected to each group).

SUB OPTIONS

Sub options are as follows:

-c *class*

class indicates the class name to which *disk* belongs.

-d *disk[,disk,...]*

disk indicates a target disk. When connecting multiple disks, the disk names should be combined using commas (,).

RETURNED VALUE

Upon successful completion, a value of 0 is returned.

Otherwise, a non-zero value is returned.

USAGE EXAMPLES

Assuming the disk on which `/`(root), `/usr`, and `/var` are installed and the disk allocated as a swap area are different, examples of the procedures for system disk mirroring and system disk mirroring cancellation are given below.

Procedure for system disk mirroring

1. Stop the running application programs.

In order to ensure mirroring definition, all the running application programs must be stopped. For the mirroring definition to be in effect, the system must be rebooted after going through this procedure.

When higher safety is required, create system disk backups.

2. Register the system disks with the root class.

In this example, the installation disk of `/`(root), `/usr`, and `/var` is "c0t0d0," and the disk allocated as a swap area is "c0t0d1."

```
# sdxdisk -M -c System -a type=root -d c0t0d0=Root1:keep,  
c1t0d0=Root2:undef,c0t0d1=Swap1:keep,c1t0d1=Swap2:undef
```

3. Connect the system disks to a group.

```
# sdxdisk -C -c System -g Group1 -d Root1,Root2 -v 0=root:on,  
1=usr:on,3=var:on,4=home:on,5=opt:on  
# sdxdisk -C -c System -g Group2 -d Swap1,Swap2 -v 0=swap:on
```

Information

When System Disks Have Unopen Physical Slices

After returning from the `sdxdisk -C` command, volumes created for unopen physical slices are started, and synchronization copying is performed. In this event, cancel the synchronization copying using the `sdxcopy -C` command, or after the synchronization copying is completed move to step 4. Physical slices on which file systems are mounted and those accessed as raw devices are considered to be open. Physical slices may not be opened if they are not shown with the `mount(1M)` command and if process IDs are not shown with the `fuser(1M)` command.

4. Check that mirroring definition is completed.

```
# sdxroot -M -c System -d Root1,Swap1
```

5. Reboot the system.

```
# shutdown -y -g0 -i6
```

6. Check that mirroring is in effect.

Using the `mount` command and the `sdxinfo` command, make sure that the system disks have been mirrored properly.

Procedure for system disk mirroring cancellation

1. Stop the running application programs.

In order to ensure mirroring cancellation, all the running application programs must be stopped. For the mirroring cancellation to be in effect, the system must be rebooted after going through this procedure.

When higher safety is required, create system disk backups.

2. Disconnect disks other than those used as system disks after this cancellation from the groups.

```
# sdxdisk -D -c System -g Group1 -d Root2
# sdxdisk -D -c System -g Group2 -d Swap2
```

3. Check that mirroring cancellation is completed.

```
# sdxroot -R -c System -d Root1,Swap1
```

4. Reboot the system.

```
# shutdown -y -g0 -i6
```

5. Check that the mirroring has been cancelled.

Using the `mount` command and the `sdxinfo` command, make sure that the system disk mirroring has been cancelled properly.

6. Cancel system disk management.

```
# sdxvolume -F -c System -v root
# sdxvolume -F -c System -v usr
# sdxvolume -F -c System -v var
# sdxvolume -F -c System -v home
# sdxvolume -F -c System -v opt
# sdxvolume -F -c System -v swap
# sdxvolume -R -c System -v root
# sdxvolume -R -c System -v usr
# sdxvolume -R -c System -v var
```

```
# sdxvolume -R -c System -v home
# sdxvolume -R -c System -v opt
# sdxvolume -R -c System -v swap
# sdxgroup -R -c System -g Group1
# sdxgroup -R -c System -g Group2
# sdxdisk -R -c System -d Root1
# sdxdisk -R -c System -d Root2
# sdxdisk -R -c System -d Swap1
# sdxdisk -R -c System -d Swap2
```

D.12 sdxparam - Configuration parameter operations

SYNOPSIS

```
sdxparam -G [-p param,...]
```

```
sdxparam -S -p param=val [,param=val,...] [-e default]
```

DESCRIPTION

Use sdxparam to perform operations on GDS configuration parameter.
You must be superuser to use this command.

PRIMARY OPTIONS

You can use either of the following options.

-G

Get

Displays the current value of configuration parameter or parameters specified by *param*. When using cluster system, parameter value of the current node will be displayed.

If the -p option is omitted, all configuration parameters are displayed.

-S

Set

Sets the value specified by *val* to the configuration parameter or parameters specified by *param*. When using cluster system, this option sets the parameter value of the current node.

The new value becomes valid upon returning from the command, and rebooting the system will not change the value.

SUB OPTIONS

Sub options are as follows:

-e default (when using **-S**)

Resets all configuration parameter values to default.

When indicated at the same time as -p option, this option is ignored.

-p *param*,...(when using **-G**)

Displays the configuration parameter *param* value.

-p *param=val[,param=val,...]* (when using **-S**)

Sets *val* to configuration parameter *param*.

You can indicate the following combination to *param* and *val*.

copy_concurrency=num

Sets the maximum number of synchronization copying you can execute simultaneously to *num*.

The value is set to 8 by default.

Values from 1 to 1024 may be specified for *num*.

copy_delay=msec

This option delays the synchronization copying by milliseconds specified by *msec*, when the copying is initiated by an event other than hot spare.

The value is set to 0 by default.

Values from 0 to 1000 may be specified for *msec*.

spare_copy_delay=msec

This option delays the synchronization copying by milliseconds specified by *msec*, when the copying is initiated by hot spare.

The value is set to 50 by default.

Values from 0 to 1000 may be specified for *msec*.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.

Otherwise, a non-zero value is returned.

CAUTION

The default values and the range of values you can use for configuration parameters may change in the future.

D.13 **sdxcfig** - Object configuration operations

SYNOPSIS

```
sdxcfig Backup -c class[-o outfile] [-e update]
```

```
sdxcfig Convert -e remove[,update] -c class -d disk,...  
[-i infile] [-o outfile]
```

```
sdxcfig Convert -e remove[,update] -c class -g group,...  
[-i infile] [-o outfile]
```

```
sdxcfig Convert -e rename[,update] -c class=classname  
[-i infile] [-o outfile]
```



```
sdxconfig Convert -e replace[,update] -c class -d disk=devi ce[, disk=devi ce,...]  
[-i infile] [-o outfile]
```

```
sdxconfig Convert -e replace[,update] -c class  
-p devi ce=newdevi ce[, devi ce=newdevi ce,...]  
[-i infile] [-o outfile]
```

```
sdxconfig Remove -c class[-e keepid]
```

```
sdxconfig Restore -c class -i infile [-e chkps,skipsync]
```

DESCRIPTION

Use sdxconfig to perform object configuration operations for classes specified by *class* (excluding shadow classes). The sdxconfig command can be executed with superuser access privileges only.

This command must be executed in multi-user mode.

PRIMARY OPTIONS

You have a choice of the following options.

Backup

Outputs the object configuration of a local class or a shared class specified by *class* to a file specified by *outfile* (standard output by default) in configuration table format. Specify the class name targeted for configuration table creation for *class*.

If *class* includes switch groups, proxy objects, DISABLE disks, or TEMP slices, creation of the *class* configuration table fails.

Convert

Converts the configuration table of a class specified by *class* according to sub option specifications. Specify the class name contained in the configuration table for *class*.

Remove

Removes the object configuration of a local class specified by *class* from the system. All objects (volumes, groups and disks) within the *class* are removed. Specify the target local class name for *class*.

Even if the class object configuration is removed using this option, contents (data) of the removed volumes are not lost. By restoring the object configuration with the Restore option, the volume configuration and contents can be restored.

If *class* includes proxy objects, ACTIVE volumes, or TEMP or COPY slices, deletion of the *class* fails.

Restore

Restores the object configuration of a class specified by *class* according to the configuration table declared in a configuration file specified by *infile*. Specify the class name contained in the configuration table for *class*.

Even if the class object configuration is restored with this option, volume areas on the physical disks registered with the class are not initialized. After the object configuration is deleted with the Remove option, by restoring the object configuration using this option, the volume configuration and contents can be restored.

However, if the configuration table contains mirror volumes with the mirroring multiplicity of two or higher, after returning from the

sdxconfig command, synchronization copying of the mirror volumes is performed automatically (excepting when using `-e skipsync`). In this event, destination slices are overwritten with data of source slices automatically selected, and data previously saved on the destination slices will be lost.

The class specified by *class* will be restored as a local class on the current node. To restore the *class* as a shared class, after this command execution it is necessary to change the type attribute and scope attribute of the *class* using the `sdxattr -C` command.

If the class specified by *class* already exists, this command results in an error. Additionally, if the physical disk size contained in the configuration table and the actual physical disk size do not match, restoration of the *class* object configuration fails.

SUB OPTIONS

Sub options are as follows.

-c class

class indicates the target class name.

-c class=classname (when using **Convert -e rename)**

Changes the class name in the configuration table from *class* to *classname*.

-d disk,... (when using **Convert -e remove)**

Removes *disk*,... from the configuration table. Specify the disk name of an undefined disk, a spare disk, a single disk, or a disk directly connected to a mirror group to be removed for *disk*.

If *disk* is a single disk, volumes and slices within the *disk* are also removed. If *disk* is the only disk connected to a mirror group, volumes and slices in the mirror group and the mirror group itself are also removed.

This option can be used along with the `-g` option.

If *disk* is connected to a concatenation group or a stripe group in the configuration table, removing the *disk* fails.

-d disk=device[,disk=device,...] (when using **Convert -e replace)**

Changes the physical disk of a disk specified by *disk* to *device* in the configuration table. *device* can also indicate a physical disk not connected to the domain.

Specify a disk name for *disk* and a physical disk name for *device*. It is necessary to separate *disk* and *device* with an equal sign (=). To change multiple physical disks, specify sets of these specifiers in comma-delimited format.

The physical disk names can be specified in one of the following formats.

```
cCt TdD    (for normal hard disks)
mphpd/     (for MPHD disks)
mplb/      (for MPLB disks)
emcpowerN  (for emcpower disks)
```

Where *C* is the controller number, *T* is the target ID, *D* is the disk number, *I* is the MPHD instance number or the MPLB instance number and *N* is the emcpower device number. The emcpower device number can be a single digit number from 0 to 9 or double or more digits number, and cannot contain a lower case letter corresponding to the slice number.

This option cannot be used along with the `-p` option.

-e chkps (when using **Restore**)

Checks consistency of disk identification information (class and disk names) stored in the private slices of physical disks to be registered with *class* and the configuration table contained in the configuration file specified by *infile*.

Restoration of *class* does not take place if any of the following conditions is not satisfied.

- All physical disks contained in the configuration table have the private slices.
- The sizes of the private slices match between all physical disks contained in the configuration table.
- The class names stored in the private slices match between all physical disks contained in the configuration table.
- For all physical disks contained in the configuration table, the disk names stored in the private slices match the disk names assigned to the physical disks in the configuration table.

-e keepid (when using **Remove**)

Retains the private slices and disk identification information stored in the private slices of all disks registered with *class*.

By using this option, when *class* is restored using physical disks that were removed from the *class* or that were copied with the copy functions of disk units, configuration consistency can be checked with the `-e chkps` option of the `sdxconfig Restore` command.



If the object configuration of class is deleted with this option, physical disks that were deleted from the class cannot be registered with a class with the `sdxdisk -M` command. Before registering those deleted physical disks with a class with the `sdxdisk -M` command, it is necessary to delete the private slice (slice whose tag is 33) with the `format(1M)` command. The private slices of all physical disks deleted from class must be deleted. For deleting the private slices of all physical disks deleted from class, restore the object configuration with the `Restore` option once, and then execute the `Remove` option without this option.

-e remove (when using **Convert**)

Removes disks or groups from the configuration table.

-e rename (when using **Convert**)

Renames the class of the configuration table.

-e replace (when using **Convert**)

Changes physical disks in the configuration table.

-e skipsync (when using **Restore**)

Leaves synchronization copying of mirror volumes created within *class* undone, assuming that equivalency of all mirror volumes contained in the configuration file specified by *infile* is ensured on user's hand. Even if slices are nonequivalent, their statuses will not be in `INVALID`.

-e update (when using **Backup, Convert**)

Overwrites the file specified by *outfile* with the configuration table when the *outfile* is an existing file.

-g *group*,... (when using **Convert -e remove**)

Removes *group*,... from the configuration table. All objects (volumes, slices, disks, lower level groups) within the *group* are deleted. Specify the deleted group name for *group*.

This option can be used along with the -d option.

When *group* is connected to a higher level group other than a mirror group in the configuration table, removing *group* fails.

-i *infile* (when using **Convert, Restore**)

Converts the configuration table or restores the object configuration of a class specified by *class* according to the configuration file specified by *infile*. Specify the path to a configuration file for *infile*, using the absolute path name or the relative path name from the current directory.

When using Convert, it is not required to specify this option. By default, a configuration table from standard input is converted.

-o *outfile* (when using **Backup, Convert**)

Sends the created or converted configuration table to a configuration file specified by *outfile*. Specify the path to a configuration file for *outfile*, using the absolute path or the relative path from the current directory.

If the file specified by *outfile* already exists, this command results in an error (excepting when using -e update).

By default, the configuration table is output to standard output.

-p *device=newdevice[,device=newdevice,...]* (when using **Convert -e replace**)

Changes the physical disk specified by *device* to another physical disk specified by *newdevice* in the configuration table. *newdevice* can also indicate a physical disk not connected to the domain.

Specify a physical disk name described in the configuration table for *device* and a new physical disk name for *newdevice*. It is necessary to separate *device* and *newdevice* with the equal sign (=). To change multiple physical disks to new disks, specify sets of these specifiers in comma-delimited format.

The physical disk names for *device* and *newdevice* can be specified in one of the following formats.

```
cCt TdD    (for normal hard disks)
mphpd/     (for MPHD disks)
mplb/      (for MPLB disks)
emcpowerN (for emcpower disks)
```

Where *C* is the controller number, *T* is the target ID, *D* is the disk number, *I* is the MPHD instance number or the MPLB instance number and *N* is the emcpower device number. The emcpower device number can be a single digit number from 0 to 9 or double or more digit number, and cannot contain a lower case letter corresponding to the slice number.

This option cannot be used along with the -d option.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.

Otherwise, a non-zero value is returned.

D.14 sdxproxy - Proxy object operations

SYNOPSIS

```
sdxproxy Break -c class -p proxy [-e force,restore]
```

sdxproxy Cancel -c class -p proxy

```
sdxproxy Join -c class -m master -p proxy
[-a mvol=pvol:jrm[:pslice][,mvol=pvol:jrm[:pslice],...]]
[-e delay=msec,softcopy,syncmode,waitsync]
```

```
sdxproxy Part -c class -p proxy,... [-a attribute=value]
[-e instant,mode=val,unlock]
```

```
sdxproxy Rejoin -c class -p proxy,...
[-e delay=msec,softcopy,waitsync]
```

```
sdxproxy RejoinRestore -c class -p proxy,...
[-e delay=msec,instant,nowaitsync,softcopy]
```

sdxproxy Relate -c class -m master -p proxy

```
sdxproxy Restore -c class -p proxy,... [-e instant,nowaitsync]
```

```
sdxproxy Root -c class -m master,... [-e boot]
```

```
sdxproxy Root -c class -p proxy,... [-e boot]
```

```
sdxproxy Root -c class -m master,... -p proxy,... [-e boot]
```

sdxproxy Swap -c class -p proxy

```
sdxproxy Update -c class -p proxy,... [-e instant,nowaitsync]
```

DESCRIPTION

Use sdxproxy to perform operations on proxy objects.
The sdxproxy command can be executed with superuser access privileges only.



See

.....

- [A.1.8 Proxy Configuration Preconditions](#)
- [A.1.9 Number of Proxy Volumes](#)
- [A.1.10 Proxy Volume Size](#)
- [A.1.11 Proxy Group Size](#)
- [A.2.18 Using the Advanced Copy Function in a Proxy Configuration](#)
- [A.2.19 Instant Snapshot by OPC](#)
- [A.2.21 Using EMC TimeFinder or EMC SRDF in a Proxy Configuration](#)

.....

PRIMARY OPTIONS

You have a choice of the following options.

Break

Cancels master-proxy relationship between a pair of volumes or groups. You can cancel the relationship when the objects are either joined or parted.

proxy indicates a proxy volume or a proxy group to be cancelled. A proxy volume within a proxy group cannot be specified for cancellation.

Even after the relationship is cancelled, the proxy volume or proxy volumes within the proxy group can be used as normal volumes that retain original volume data and attributes.

You can also cancel the relationship when master volume and proxy volume are in use. However, if the joined master volume is in use, the data integrity of the proxy volume should be ensured in the file system layer or database layer that is managing data, after the relationship is cancelled. For instance, if you are using the master volume as a file system, use the `lockfs(1M)` command to write-lock the file system, before canceling the relationship.

The command will result in an error when:

- Copying is in process from the master volume to the proxy volume (except for when using `-e force`)
- Copying is in process from the proxy volume to the master volume (except for when using `-e restore`)

Cancel

Cancels (releases) sessions of the copy functions of disk units existing between parted proxies and masters.

Specify a parted proxy group or a parted proxy volume into *proxy*. When *proxy* is a proxy group, all parted proxies within the proxy group become targets. A parted proxy volume within a proxy group can also be specified. However, when a BCV pair or an SRDF pair exists between the master group and the proxy group, sessions cannot be cancelled specifying a parted proxy volume within the proxy group.

Sessions can be canceled even if master volumes and proxy volumes are in use. Even after canceling the sessions, the masters and the proxies are left parted. The volume statuses of the masters and the proxies also remain unchanged. However, if sessions are canceled when copying from masters to proxies and vice versa are in process, data becomes INVALID. If that happens, perform copying again, and the volume statuses will be restored when the copying is complete.

Join

A pair of volumes or a pair of groups are related and joined as master and proxy.

When joining a pair of volumes, synchronization copying of the master volume to the proxy volume is performed after returning from the command (when using `-e waitsync`, before returning from the command).

When joining a pair of volumes, the following conditions must be satisfied.

- The master volume size and the proxy volume size match.
- The master volumes and the proxy volumes belong to different mirror groups or single disks.

When joining a pair of groups, proxy volumes with the same offset and size as master volumes within the master group will be created in the proxy group, and synchronization copying from the master volumes to the proxy volumes is performed after returning from the command (when using `-e waitsync`, before returning from the command). The access mode of the proxy volumes created in the proxy group is set to "ro" (read-only). If a keep disk is connected to the master group or the proxy group and geometry such as the cylinder size does not match between the master group and the proxy group, geometry of the proxy group is changed conforming to that of the master group.

When joining a pair of groups, there are the following conditions and restrictions.

- For the root class, the smallest physical disk size directly connected to the proxy group must be larger than the last block number of a volume within the master group.
- For a local class or a shared class, the proxy group size must be larger than the last block number of a volume within the master group.
- If the master group has no volume and the proxy group already has volumes, joining them results in an error.
- The master group and the proxy group must be mirror groups.

When the copy functions of disk units are available, synchronization copying from masters to proxies is performed with those copy functions (except for when using `-e softcopy`).

A proxy volume that is joined cannot be accessed or activated. In order to access the proxy volume, part the volume from master using the Part option, or break the relationship with master using the Break option.

You can create multiple snapshots by joining another proxy to a master, which is already joined with a proxy. However, the total number of slices which belong to the master volume and slices which belong to the proxy volumes that are related to the master volume must not exceed 32.

Master volumes that are already related to a proxy cannot be joined to another master as a proxy. Also, a proxy that is already related to a master cannot be joined with another proxy.

The command will result in an error when:

- The proxy volume is activated.
- There is a slice that is temporarily detached or in copying process among any of the slices comprising the proxy volume or the master volume.
- Copying is in process between the master volume and the other proxy volume.
- The master volume is in INVALID status.

Part

Separates a proxy or proxies in joined status from the master. The master-proxy relationship will be maintained after parted. The parted proxy volume will be the snapshot containing the copy of the master volume data at the time of parting. By using the parted proxy volume, you can for instance, create a backup of the master volume at the time it was parted, or use it for other purposes.

proxy indicates a proxy volume or a proxy group in joined status. When proxy group is indicated, all proxy volumes in the group will be parted. You can also indicate a proxy volume in joined status in a proxy group.

After the parting is complete, the proxy volume will be activated as an independent volume and become accessible using the following special file.

```
/dev/sfdisk/class/rdisk/volume_name  
/dev/sfdisk/class/disk/volume_name
```

When the proxy volume belongs to a "shared" type class, it will be activated on all nodes defined in the class scope.

You can part master and proxy volumes even if the master volume is active, but the data integrity of the parted proxy volume must be ensured in the file system layer or database layer that is managing data. For instance, if you are using the master volume as a file system, use the `lockfs(1M)` command to write-lock the file system, before parting.

The command will result in an error when:

- Copying is in process from the master volume to the proxy volume (except for when using `-e instant`)

- Copying is in process from the proxy volume to the master volume

Rejoin

Rejoins one or more parted proxy with the master.

proxy indicates a parted proxy volume, or proxy group. The mirror or single type proxy can be specified. When a proxy group is indicated, all proxy volumes in the group will be rejoined. A parted proxy volume in a proxy group can also be specified.

Synchronization copying of master volumes to proxy volumes is performed after returning from the command (when using `-e waitsync`, before returning from the command). When the copy functions of disk units are available, synchronization copying is performed with those copy functions (except for when using `-e softcopy`).

When more than one proxy volume related to the same master volume is specified simultaneously, this command will result in an error.

The command will result in an error when:

- The proxy volume is active.
- There is a slice that is in copying process among any of the slices comprising the proxy volume or the master volume.
- Copying is in process between the master volume and the other proxy volume.
- The master volume is in INVALID status.

Note

When proxy volumes are running as system volumes, they cannot be stopped and thus rejoining fails. To rejoin such proxy volumes, firstly switch the boot environment using the `sdxproxy Root` command to free up the proxy volume

RejoinRestore

Rejoins a proxy or proxies in parted status with the master and restores the master volume data using the proxy volume data. Master volume data are restored by synchronization copying from the proxy volume to the master volume. When executing the command using this option, the master volume data are overwritten with the proxy volume data.

proxy indicates a proxy volume or a proxy group in parted status. When a proxy group is indicated, all proxy volumes in the group will be rejoined and data of the related master volumes will be restored. You can also indicate a proxy volume in parted status in a proxy group.

Synchronization copying of master to proxy volume is performed before returning from the command (when using `-e waitsync`, after returning from the command). When the copy functions of disk units are available, synchronization copying from masters to proxies is performed with those copy functions (except for when using `-e softcopy`).

When more than one proxy volume related to the same master volume is specified simultaneously, this command will result in an error.

The command will result in an error when:

- The master volume or the proxy volume is activated.
- There is a slice that is in copying process among any of the slices comprising the proxy volume or the master volume.
- Copying is in process between the master volume and the other proxy volume.
- The master volume is in INVALID status.

Relate

Relates and parts a pair of volumes or a pair of groups as a master and a proxy. This operation does not change data, statuses and attributes of the master and the proxy. To the related master and proxy, sessions by the copy functions of disk units are not set.

To relate a pair of volumes, the volumes must conform to the following conditions.

- The master volume and the proxy volume belong to different groups or single disks.
- The master volume size and the proxy volume size match.
- The master volume and proxy volume types are mirror or single.

To relate a pair of groups, the groups must conform to the following conditions.

- The master group and the proxy group are mirror groups.
- The layout (offsets and sizes) of volumes of the master group match with that of the proxy group.

For masters to whom proxies are already related, other additional proxies can be related. However, the number of slices comprising a master volume and all proxy volumes related to the master volume is limited to 32 in total.

Masters to whom proxies are already related cannot be related as proxies to other masters, or for proxies already related to masters, other proxies cannot be related.

This command will result in an error when:

- A slice being copied or temporarily detached exists in the master volume or the proxy volume.
- Copying is in process between the master volume and another proxy volume.

Restore

Copies data from a parted proxy to a master and restores contents of the master. With the OPC function, data of the proxy at the moment is copied (overwritten) to the master. The command returns control after the copying is complete (right after the copying starts when using the `-e instant` option and the `-e nowaitsync` option). If the OPC function is unavailable, the command fails.

Specify one or more proxy groups or parted proxy volumes for *proxy*. When *proxy* is a proxy group, all parted volumes within the proxy group become targets. A parted proxy volume within a proxy group can also be specified. Do not specify multiple proxy volumes related to the same master volume simultaneously.

The Restore operations can be performed even if proxy volumes are active, but it is necessary to secure consistency of data copied to master volumes in the file system layer or database layer that is managing data. For example, if the proxy volume is used as a file system, write-lock the file system with the `lockfs(1M)` command and then perform restoration.

This command will result in an error when:

- The master volume is active.
- A slice being copied exists in the master volume or the proxy volume.
- Copying is in process between the master volume and another proxy volume.
- A proxy volume joined to the master volume exists.
- The proxy volume is in INVALID status.

Root

Configures master volumes and proxy volumes specified by *master,...* and *proxy,...* for using them as file systems or swap areas in an alternative boot environment. When a master group or a proxy group is specified, all volumes that belong to the specified group will be configured.

Volumes to be used in an alternative boot environment must conform to the following conditions.

- The volumes are related directly or indirectly as the master and the proxy (alternative volumes) to volumes declared as file systems or swap areas in the `/etc/vfstab` file (current volumes).
- The volumes are parted.
- The access mode is "rw" (read and write).
- The volumes are in status other than INVALID (invalid data).
- The volumes are not running as file systems or swap areas.

It is not required to specify alternative volumes for all current volumes, but a volume to be used as the root file system in the alternative boot environment (alternative root volume) must always be specified.

Before returning from the command, the device names and the special file names contained in the `/etc/system` file and the `/etc/vfstab` file on the specified alternative root volume are changed to those for the specified alternative volume. Current volumes of which alternative boot volumes were not specified are included in `/etc/vfstab` on the alternative root volume without change. After parting the current root volume and the alternative root volume using the `Part` option, to edit `/etc/vfstab` on the current root volume, to edit `/etc/vfstab` on the alternative root volume, or to perform configuration change such as volume creation or deletion, firstly configure the alternative boot environment using this command. When executing the `sdxproxy` command with this option after these configurations are changed, after returning from the command, check whether contents of `/etc/vfstab` on the alternative root volume are correct. If an alternative volume that is a swap area used as a dump device is specified, configuration of the alternative volume for using it as a dump device takes place when starting the alternative boot environment.

For modifying the `elilo.conf` and `fstab` files on the alternative root volume, the alternative root volume is temporarily mounted on the `/.GDSPROXY` directory. This temporal mount point can be changed by specifying the mount point path in the environment variable `PROXY_ROOT`.

When alternative boot environment configuration is completed, the boot device names for the current boot environment and the alternative boot environment are output to standard output (excepting when using `-e boot`). Be sure to take a note of the output boot device names. By specifying the boot device name with the `boot` command in `OpenBoot`, the environment can be switched to the alternative boot environment. Similarly, by selecting the boot device name for the current boot environment, the environment can be switch back to the current boot environment. With successful boot environment switchover, the boot environment will be the default boot environment.

Swap

Swaps the master's slices with the proxy's slices.

proxy indicates a proxy volume or a proxy group in the joined status. A proxy volume within a proxy group cannot be specified for swapping.

You can swap the slices when master is in use.

The command will result in an error when:

- There is a slice that is in copying process among any of the slices comprising the proxy volume or the master volume.
- Copying is in process between the master volume and the other proxy volume.
- The proxy volume is in INVALID status.
- Between a master and a proxy, EC session, a BCV pair, or a SRDF pair exists.

Update

Copies data from a master to a parted proxy and updates contents of the proxy. With the `OPC` function, data of the master at the moment is copied (overwritten) to the proxy. The command returns control after the copying is complete (right after the copying starts when using the `-e instant` option and the `-e nowaitsync` option).

Updated proxy volumes become snapshots that have copies (replicas) of data of master volumes at the moment. By use of the updated proxy volumes, creating backups of master volumes at the moment and running other services become possible.

Specify one or more proxy groups or parted proxy volumes into *proxy*. When a proxy group is specified, all the parted proxy volumes within the proxy group become targets. A parted proxy volume within a proxy group can also be specified. Do not specify multiple proxy volumes related to the same master volume simultaneously.

The Update operations can be performed even if master volumes are active, but it is necessary to secure consistency of data copied to proxy volumes in the file system layer or database layer that is managing data. For example, if the master volume is used as a file system, write-lock the file system with the `lockfs(1M)` command and then perform update.

This command will result in an error when:

- The proxy volume is active.
- A slice being copied exists in the master volume or the proxy volume.
- Copying is in process between the master volume and another proxy volume.
- The master volume is in an INVALID status.

SUB OPTIONS

Sub options are as follows.

-a *attribute=value* (when using **Part)**

Sets *attribute* that is the attribute of the parted proxy volume to *value*. This attribute value becomes invalid when the proxy volume is rejoined to the master volume with the `Rejoin` or `RejoinRestore` option.

attribute indicates the attribute name, and *value* indicates the attribute value. It is necessary to separate *attribute* and *value* with an equal sign (=).

You can specify one of the following combinations to *attribute* and *value*.

pjrm=on or pjrm=off (default is on)

Sets the just resynchronization mechanism mode for proxies.

on

Turns "on" the just resynchronization mechanism mode for proxies.

off

Turns "off" the just resynchronization mechanism mode for proxies.

-a *mvol=pvol:jrm[:pslice]* [,*mvol=pvol:jrm[:pslice]* ,...] (when using **Join**)

Indicates the proxy volume attributes.

mvol must always be followed by an equal (=) sign, and *pvol*, *jrm* and *pslice* are delimited by a colon (:). When indicating attributes for more than one proxy volume, combine specifiers with a comma (,) as the delimiter.

When joining a pair of groups, *mvol* indicates the volume name of the master volume within the master group. *pvol* indicates the volume name of the proxy volume that will be created in the proxy group corresponding to the master volume specified by *mvol*, *jrm* indicates the just resynchronization mechanism mode for the volume (on or off), and *pslice* indicates the physical slice attribute (on or off). You must specify attributes for all proxy volumes created for respective master volumes within the master group. If *pslice* is omitted, the physical slice attribute of the proxy volume will be equivalent to that of the corresponding master volume.

When joining a pair of volumes, *mvol* indicates the volume name of the master volume, *pvol* indicates the volume name of the proxy volume, *jrm* indicates the just resynchronization mechanism mode for the volume (on or off), and *pslice* indicates the physical slice attribute (on or off). *mvol* and *pvol* should match with *master* and *proxy* respectively. When not using this option, the proxy volume attributes will be the same as before it was joined.

If *class* is the root type, "off" cannot be set to *pslice*.

-c class

class indicates the class name to which the master object or proxy object that will be the target of operation belongs

-e boot (when using Root)

Sets an alternative boot environment as the default boot environment. After returning from the `sdxproxy` command, reboot the system immediately, and the environment will be switched to the alternative boot environment.

Once the boot-device parameter is changed, the boot-device parameter values before and after changed are output to standard output. Be sure to take a note of the output values. By re-specifying the original boot-device parameter value with the `eeeprom(1M)` command or the `setenv` command in OpenBoot, it is possible to restore the default boot environment from the current boot environment. Additionally, by specifying the original boot-device parameter value with the `boot` command in OpenBoot, it is possible to set the original current boot environment as the default boot environment again.

If the root volume's slice status is changed under synchronization copy completion or due to an I/O error, or if a GDS daemon ends abnormally and it is restarted, the boot device for the current boot environment is set as the default boot device again. Therefore, when this option is used, after returning from the `sdxproxy` command it is necessary to reboot the system immediately.

-e delay=*msec* (when using Join, Rejoin, RejoinRestore)

Delays the issuing of the input/output request to the disk at the time of synchronization copying between master volume and proxy volume by milliseconds, specified by *msec*.

This option allows users to adjust the influence on the application accessing the master volume.

When copying is performed with a disk unit's copy function, this option setting is ignored.

When this option is omitted, the delay will be set to 0.

Values from 0 to 1000 may be specified for *msec*.

-e force (when using Break)

Forcibly breaks master-proxy relationship even when copying is in process between master and proxy.

When using this option, copying process will be cancelled and the status of the proxy volume will become INVALID.

-e instant (when using Part)

Parts proxy volumes and creates virtual snapshots of master volumes with the OPC function even if copying from the master to the proxy is in process. After returning from the command, the parted proxy volume will become accessible before the copying is complete, and will serve as a snapshot containing data of the master volume at the time of parted. When copying is in process from proxy to master, or if a hard copy function is not available, the command will result in an error.

-e instant (when using **RejoinRestore**)

Restoration will be completed instantly, and the command will be returned. After returning from the command, synchronization copying from proxy volume to master volume will automatically begin. Although the copying is still in process, restoration will appear to be complete. After returning from the command, the master volume can be activated and accessed, before the copying is complete. Master volume data will appear to have been overwritten by the proxy volume data at the time of executing the command.

-e instant (when using **Restore, Update**)

Instantly completes restore or update and returns from the command. After returning from the command, background copying with the OPC function is performed. Before the background copying is complete, you may start the copy destination volumes for access to valid data.

-e mode=*val* (when using **Part**)

Indicates the access mode of the proxy volume which will be activated.

val indicates either of the following options.

rw

Sets access mode for read and write.

ro

Sets access mode for read only.

Opening a read-only volume in write mode will result in an error.

Although proxy volume will be activated in the access mode specified by *val*, the proxy volume's access mode attribute will remain unchanged. Access mode specified by *val* is valid only while the proxy volume is activated ("Current Access Mode"), and will become invalid once the proxy volume is stopped. When the proxy volume is restarted, it will start in the mode set by access mode attribute ("Default Access Mode"), except for when the access mode is specified at the time of restart.

When this option is omitted, proxy volume will be activated in the access mode according to the access mode attribute set on each node.

-e nowaitsync (when using **RejoinRestore, Restore, Update**)

Returns control from the command right after copying starts. After returning from the command, wait until the copying is complete to start the copy destination volumes. To start copy destination volumes without waiting until copying is complete, use the **-e instant** option. When the **-e instant** option is specified simultaneously, this option is ignored.

-e restore (when using **Break**)

Cancels copying from a proxy to a master when such copying is in process and forces the command to break the relationship between the master and the proxy.

If copying is canceled and relationships are broken with this option, the master volumes after this operation becomes INVALID.

-e softcopy (when using **Join, Rejoin** or **RejoinRestore**)

The copy function of a disk unit will not be used for synchronization copying between master volume and proxy volume.

-e syncmode (when using **Join**)

When the REC function is used for synchronization copying from the master volume to the proxy volume, the transmission mode of REC is set to the synchronous mode. The default is the asynchronous Through mode. When the REC function is not used, this option setting is ignored.

-e unlock (when using **Part**)

The proxy volume will be activated regardless to whether it is locked.

Lock mode will not be changed unless you change it with the `sdxattr -V` command.

-e waitsync (when using **Join** or **Rejoin**)

When performing synchronization copying, returns from the command after the copying is complete.

-m master (when using **Join**, **Relate**)

Specifies the master volume or the master group that is joined or related.

master indicates the volume name of the master volume, or the group name of the master group.

-m master,... (when using **Root**)

Specifies one or more master volumes or master groups as the object of the operation. When multiple volumes or groups are specified, they must be in the same class.

master indicates the name of the master volume, or the name of the master group.

To specify multiple volumes and groups, separate each name with a comma (,).

-p proxy (when using **Break**, **Cancel**, **Join**, **Swap** or **Relate**)

proxy indicates a proxy volume or a proxy group that is the target of the operation.

proxy indicates the volume name of the proxy volume, or the group name of the proxy group.

-p proxy,... (when using **Part**, **Rejoin**, **RejoinRestore**, **Restore**, **Root** or **Update**)

proxy indicates one or more proxy volumes or proxy groups that will be the target of the operation. When indicating more than one volume or group, they must belong in the same class.

proxy indicates the volume name of the proxy volume or the group name of the proxy group that will be the target of the operation

When indicating more than one volume name or group name, combine them with a comma (,) as the delimiter.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.

Otherwise, a non-zero value is returned.

D.15 sdxshadowdisk - Shadow disk operations

SYNOPSIS

```
sdxshadowdisk -C -c class -g group -d disk, ...  
                [-a attribute=value [, attribute=value ]]
```

```
sdxshadowdisk -D -c class -g group -d disk
```

```
sdxshadowdisk -M -c class -d device=disk[:type] [, device=disk[:type] , ...]
```

```
sdxshadowdisk -R -c class -d disk
```

DESCRIPTION

Use `sdxshadowdisk` to perform operations on shadow disks specified by *disk*. This command can be executed with superuser access privileges only.

PRIMARY OPTIONS

You have a choice of the following options.

-C

Connect

Connects one or more shadow disks (single type or undefined type) specified by *disk*,... to a group specified by *group*. *class* indicates the name of the shadow class with which *disk* is registered.

If there is no shadow group with the name specified by *group*, it is created automatically.

You cannot connect a shadow disk of the single type to an existing shadow group. Also, multiple shadow disks of the single type cannot be connected to the same shadow group simultaneously.

The type attribute of a shadow disk connected to a shadow group will be changed to match the type attribute of that group (mirror, stripe or concatenation). Shadow disks and lower level groups that are connected to the same shadow group will be mirrored, striped or concatenated, depending on their type attributes.

Details about connecting shadow disks to mirror type, stripe type, and concatenation type shadow groups are described below.

When connecting to a shadow group of the mirror type

Shadow disks and lower level shadow groups connected to the same shadow group of the mirror type will be mirrored one another. When only one shadow disk or one lower level shadow group is connected to a shadow group of the mirror type, the shadow volume created within that shadow group will not be mirrored. When configuring a mirroring environment with "n"-way multiplexing, "n" numbers of shadow disks or lower level shadow groups must be connected. A maximum of eight-way multiplex mirroring is supported.

If a shadow disk is connected to a shadow group of the mirror type including a shadow volume, synchronization copying of the shadow volume is not performed. To ensure synchronization for a shadow volume of the mirror type, the mirror volume must be properly synchronized with GDS or SynfinityDisk that manages the mirror volumes corresponding to shadow volumes.

By connecting a shadow disk of the single type including a shadow volume to a group of the mirror type, the shadow volume can also be changed from the single type to the mirror type.

The available size of a shadow group of the mirror type (available capacity as shadow volumes) will be the same size as that of the smallest shadow disk or lower level shadow group connected. If connecting *disk* results in a decrease in the available size of

group, a warning message will be sent to standard error output.

When connecting to a shadow group of the stripe type

Shadow disks specified by *disk...* will be connected to *group* in the order they are listed. Disks connected to a stripe group in another domain should be connected in the same order. Alternatively, destination disks copied with the copy functions of disk units from disks connected to a stripe group should be connected in the same order. For the disk connecting order, check the DISKS field displayed with the `sdxinfo -G` command. Respective shadow disks and lower level shadow groups connected to the same shadow group of the stripe type will configure stripe columns, and will be striped in the order they were connected. When only one shadow disk or one lower level shadow group is connected to a shadow group of the stripe type, a shadow volume cannot be created within that shadow group. When striping "n" number of columns, "n" number of shadow disks or lower level shadow groups must be connected. Striping of two or more columns up to 64 columns is supported.

When a shadow group of the stripe type specified by *group* already exists, stripe columns will be added after the existing stripe columns in *group*, in the order they are specified by *disk,...* However, a shadow disk with the available size smaller than the stripe width cannot be connected to the existing shadow group of the stripe type. In addition, you cannot increase stripe columns by connecting shadow disks to a stripe group with a shadow volume, or to a stripe group connected to a higher level shadow group.

The available size of a shadow group of the stripe type (available capacity as shadow volumes) equals the available size of the smallest shadow disk (or lower level shadow group) multiplied by the number of stripe columns, and rounded down to the common multiple of the stripe width times stripe columns and the cylinder size. If connecting *disk* decreases the available size of *group*, a warning message will be sent to standard error output.

You cannot connect a shadow disk of the single type to a shadow group of the stripe type.

When connecting to a shadow group of the concatenation type

Shadow disks connected to the same shadow group of the concatenation type will be concatenated in the order they are specified by *disk...* Disks connected to a concatenation group in another domain should be connected in the same order. Alternatively, destination disks copied with the copy functions of disk units from disks connected to a concatenation group should be connected in the same order. For the disk connecting order, check the DISKS field displayed with the `sdxinfo -G` command. A maximum of 64 disks can be concatenated.

The available size of a shadow group of the concatenation type (available capacity as shadow volumes) equals the total of the available size of connected shadow disks.

The available size of an existing shadow group of the concatenation type can be increased by connecting shadow disks. When a shadow group of the concatenation type specified by *group* already exists, shadow disks will be concatenated in the order they were specified by *disk,...* after the disk that was last concatenated in *group*. However, you cannot add a shadow disk to a lower level shadow group of the concatenation type if the highest level shadow group of the stripe type already has a shadow volume. Also if the order of connecting shadow groups from the higher level is the mirror type, the stripe type and the concatenation type, a shadow disk cannot be connected to the lowest level shadow group of the concatenation type.

You cannot connect a shadow disk of the single type to a shadow group of the concatenation type.

-D

Disconnect

Disconnects a shadow disk specified by *disk* from a shadow group specified by *group*. *class* indicates the name of the shadow class with which the *disk* is registered, and *group* indicates the name of the shadow group to which *disk* is connected.

The disconnected shadow disk will have the original type attribute again (single or undefined).

If only *disk* is connected to *group*, *group* will automatically be removed upon disconnecting *disk*. However, when *disk* is the only object connected to *group* and *group* is connected to a higher level shadow group, disconnection will result in an error. In such a case, disconnect *group* from the higher level shadow group using the `sdxshadowgroup -D` command, and then disconnect *disk*.

You cannot disconnect *disk* if the disconnection will result in a change on the status of any of the existing shadow volumes within *group*.

Conditions that do not allow you to disconnect a shadow disk from a shadow group of the mirror type, stripe type or concatenation type are as below.

When disconnecting from a shadow group of the mirror type

For example, you cannot disconnect *disk* from a shadow group of the mirror type specified by *group* if a shadow volume exists within the group, and *disk* is the only object connected to *group*.

When disconnecting from a shadow group of the stripe type

You cannot disconnect a shadow disk from a shadow group of the stripe type including an existing shadow volume, or from a shadow group of the stripe type connected to a higher level shadow group.

When disconnecting from a shadow group of the concatenation type

The only disk you can disconnect from a shadow group of the concatenation type is the shadow disk that was concatenated last.

A shadow disk containing shadow volume data cannot be disconnected from a shadow group of the concatenation type.

You cannot disconnect a shadow disk from a lower level shadow group of the concatenation type if the highest level shadow group has an existing shadow volume. Also, if the order of connecting shadow groups from the higher level is the mirror type, stripe type and the concatenation type, a shadow disk cannot be disconnected from the lowest level shadow group of the concatenation type.

-M

Make

Registers one or more physical disks specified by *device* with a shadow class. *class* indicates the name of the destination shadow class. Once physical disks are registered, they can then be managed by GDS. Accordingly, the user can perform operations on the disk by use of the disk name specified by *disk*. However, *device* will be no longer managed by GDS if the current node is rebooted or if the GDS daemon on the current node is re-launched because the configuration information of a shadow class is only retained on the memory of the current node but not stored on the private slice.

If no shadow class with the same name as *class* exists, then it is automatically created. The type attribute of the shadow class is "local," and objects in the shadow class are only available on the current node.

A shadow class can include physical disks that are not registered with other classes in the current domain and on which the private slices of GDS or SynfinityDisk exist. In other words, a shadow class can include physical disks that are registered with classes in other domains and physical disks to which the private slices of SDX disks are copied with the copy functions of disk units. Physical disks can be registered with the same shadow class if they are registered with classes that have the same names in other domains or if they are destinations to which the private slices of SDX disks registered with classes with the same names are copied with the copy functions of disk units. In addition, disks with the private slices of different sizes cannot be registered with the same shadow class.

While contents on physical disks (excluding keep disks) registered by the *sdxdisk* command are initialized, contents on physical disks registered by the *sdxshadowdisk* command are not changed.

-R

Remove

Removes a shadow disk specified by *disk* from a shadow class specified by *class*. *class* indicates the name of the shadow class with which *disk* is registered.

The removed shadow disk is no longer managed by GDS.

When the last shadow disk is removed from *class*, the shadow class definition is automatically removed as well.

A shadow disk cannot be removed when a shadow volume exists within *disk*, or when *disk* is connected to a shadow group.

SUB OPTIONS

Sub options are as follows:

-a attribute=value[,attribute=value] (when using **-C**)

When using the **-C** option and defining a new group name with the **-g** option, a new shadow group is automatically created. This option sets *attribute* to value for the created *group* attribute.

The *attribute* indicates the attribute name, and the *value* indicates the attribute value. The equal sign (=) is always necessary between attribute and value. When specifying multiple attributes, each specifier set must be separated by a comma (,).

If no shadow group is created, specifying *value* different from the existing *group* attribute value will result in an error. You cannot change the attribute value of an existing *group*.

You can specify one of the following combination to *attribute* and *value*.

If multiple attributes are specified and any error occurs, the entire process is canceled.

type=mirror, type=stripe or type=concat (default is **mirror**)

Sets the type attribute of *group*.

mirror

Sets type attribute to "mirror".

stripe

Sets type attribute to "stripe".

concat

Sets type attribute to "concatenation".

width=blks (default is 32)

Sets the stripe width of *group*. *blks* indicates the stripe width in blocks (decimal numbers). One block is 512 bytes. For *blks*, you can specify an integer that is the power of two, is equal to or larger than one and is equal to or smaller than the available size of the smallest shadow disk among disks specified by *disk*,... If *group* is not the stripe type, this option will cause in an error.

-c class

class indicates the name of the shadow class to which the target shadow disk is registered or is to be registered.

-d device=disk[:type] [,device=disk[:type],...] (when using **-M**)

device indicates the name of the physical disk, *disk*, the name of the disk, and *type*, the type attribute of the shadow disk. An equal sign (=) always follows *device*, and if *type* is specified it must be separated from *disk* by a colon (:). To register multiple *devices*, separate each specifier set as above with a comma (,).

The physical disk name can be specified in either the following formats:

`cCdTd` (for normal hard disks)
`mphd/` (for MPHD disks)
`mplb/` (for MPLB disks)
`emcpowerN` (for emcpower disks)

Where *C* is the controller number, *T* is the target ID, *D* is the disk number, *I* is the MPHD instance number or the MPLB instance number, and *N* is the emcpower device number.

The emcpower device number is a single digit number or double or more digit numbers by numbers from 0 to 9, and cannot contain a lower case letter corresponding to the slice number.

If *device* is registered with a class in another domain, the same disk name as that in the domain must be specified to *disk*. If *device* is a destination to which the private slice of an SDX disk is copied with a disk unit's copy function, the same disk name as the SDX disk name must be specified to *disk*.

One of the following types can be specified to *type*. The default value for the registered shadow disk is the undefined type.

single

Single type.

undef

Undefined type.

If "single" is specified to *type*, *device* is registered as a shadow disk of the single type. For the shadow disk of the single type, a shadow volume of the single type can be created with the `sdxshadowvolume` command even if it is not connected to a shadow group.

-d *disk* (when using -D, -R)

disk indicates the name of the shadow disk that is the object of the operation.

-d *disk*,... (when using -C)

disk indicates the name of the shadow disk that is the object of the operation. To connect multiple shadow disks, separate each disk name with a comma (,).

-g *group* (when using -C,-D)

group indicates the name of the shadow group to which the shadow disk as the object of the operation is connected or is to be connected.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.

Otherwise, a non-zero value is returned.

D.16 sdxshadowgroup - Shadow group operations

SYNOPSIS

```
sdxshadowgroup -C -c class -h hgroup -l lgroup,...  
[-a attribute=value[,attribute=value]]
```

```
sdxshadowgroup -D -c class -h hgroup -l lgroup
```

```
sdxshadowgroup -R -c class -g group
```

DESCRIPTION

Use `sdxshadowgroup` to perform operations on shadow groups.
This command can be executed with superuser access privileges only.

PRIMARY OPTIONS

You have a choice of the following options.

-C

Connect

Connects one or more shadow groups (stripe type or concatenation type) specified by *lgroup*,... to a shadow group (mirror type or stripe type) specified by *hgroup*. *class* indicates the name of the shadow class to which *lgroup* belongs.

When no shadow group with the same name as *hgroup* exists, it is created automatically.

A shadow group specified by *hgroup* is referred to as a higher level shadow group, and a shadow group specified by *lgroup* is referred to as a lower level shadow group.

Lower level shadow groups and shadow disks connected to the same higher level shadow group are mirrored or striped according to the type attribute of the higher level shadow group. Connecting a shadow group to a higher level shadow group does not change the type attribute of the lower level shadow group.

You cannot connect shadow groups when:

- *lgroup* is the mirror group.
- *hgroup* is the concatenation group.
- Type attributes of *lgroup* and *hgroup* are the same.

In addition, a shadow group already including a shadow volume cannot be connected to another shadow group.

Details about connecting shadow groups to mirror type and stripe type shadow groups are described below.

When connecting to a shadow group of the mirror type

One or more shadow groups (stripe type or concatenation type) specified by *lgroup*,... can be connected to a shadow group of the mirror type specified by *hgroup*.

Shadow disks and lower level shadow groups connected to the same shadow group of the mirror type will be mirrored one another. When only one shadow disk or one lower level shadow group is connected, the shadow volume created within that shadow group of the mirror type will not be mirrored. When configuring a mirroring environment with "n"-way multiplexing, "n" numbers of shadow disks or lower level shadow groups must be connected. A maximum of eight-way multiplex mirroring is supported.

If a lower level shadow group is connected to a shadow group of the mirror type with a shadow volume, synchronization copying for the shadow volume is not performed. To ensure synchronization for a shadow volume of the mirror type, the mirror volume must be properly synchronized with GDS or SynfinityDisk that manages the mirror volume corresponding to the shadow volume.

The available size of a shadow group of the mirror type (available capacity as shadow volumes) will be the same as that of the smallest shadow disk or lower level shadow group connected. If connecting *lgroup* decreases the available size of *hgroup*, a warning message will be sent to standard error output.

When connecting to a shadow group of the stripe type

One or more shadow groups (concatenation type) specified by *lgroup*,... can be connected to a shadow group of the stripe type indicated by *hgroup*. Shadow groups indicated by *lgroup*,... will be connected to *hgroup* in the order they are listed. Lower level groups connected to a stripe group in another domain should be connected in the same order. Alternatively, destination disks copied with the copy functions of disk units from lower level groups connected to a stripe group should be connected in the same order. For the order of connecting lower level groups, check the DISKS field displayed with the `sdxinfo -G` command.

Respective shadow disks and lower level shadow groups connected to the same shadow group of the stripe type will configure stripe columns, and will be striped in the order they were connected. When only one shadow disk or one lower level shadow group is connected, a shadow volume cannot be created within that shadow group. When striping "n" number of columns, "n" number of shadow disks or lower level shadow groups must be connected. Striping of two or more columns up to 64 columns is supported.

When a shadow group of the stripe type indicated by *hgroup* already exists, stripe columns will be added after the existing stripe columns in *hgroup*, in the order they are indicated by *lgroup*,... However, a shadow group with the available size smaller than the stripe width cannot be connected to an existing shadow group of the stripe type. In addition, you cannot increase stripe columns by connecting shadow groups to a stripe group with a shadow volume, or to a stripe group connected to a higher level shadow group.

The available size of a shadow group of the stripe type (available capacity as shadow volumes) equals the available size of the smallest shadow disk (or the lower level shadow group) connected multiplied by the number of stripe columns, and rounded down to the common multiple of the stripe width times stripe columns and the cylinder size. If connecting *lgroup* decreases the available size of *hgroup*, a warning message will be sent to standard error output.

-D

Disconnect

Disconnects a shadow group specified by *lgroup* from a shadow group specified by *hgroup*. *class* indicates the name of the shadow class to which *lgroup* belongs, and *hgroup* indicates the name of the higher level shadow group to which *lgroup* is connected.

When *lgroup* is the only object connected to *hgroup*, *hgroup* will automatically be removed upon disconnecting *lgroup*. However, when *lgroup* is the only object connected to *hgroup*, and *hgroup* is connected to a higher level shadow group, disconnection will result in an error. In such a case, disconnect *hgroup* from its higher level shadow group, and then disconnect *lgroup*.

You cannot disconnect *lgroup* if the disconnection may result in a change in the status of any existing shadow volume within *hgroup*.

Conditions that do not allow you to disconnect a shadow group from a shadow group of the mirror type or a shadow group of the stripe type are as below.

When disconnecting from a higher level shadow group of the mirror type

For example, you cannot disconnect *lgroup* from a shadow group of the mirror type specified by *hgroup* if a shadow volume exists within that group, and *lgroup* is the only object connected to *hgroup*.

When disconnecting from a shadow group of the stripe type

You cannot disconnect a lower level shadow group from a shadow group of the stripe type with an existing shadow volume, or from a shadow group of the stripe type connected to a higher level shadow group.

-R

Remove

Removes the shadow group definition specified by *group*. *class* indicates the name of the shadow class to which *group* belongs.

Shadow disks and lower level shadow groups connected to *group* will be disconnected. The disconnected shadow disk will have the original type attribute (single or undefined).

The definition cannot be removed when a shadow volume exists within *group*, or when *group* is connected to a higher level shadow group.

SUB OPTIONS

Sub options are as follows:

-a *attribute=value[,attribute=value]* (when using **-C**)

When using the **-C** option and defining a new group name with the **-h** option, a new shadow group, *hgroup*, is automatically created. This option sets *attribute* to *value* for the created *hgroup* attribute.

attribute indicates the attribute name, and *value* indicates the attribute value. An equal sign (=) is always necessary between *attribute* and *value*. When specifying multiple attributes, each specifier must be separated by a comma (,).

If no shadow group is created, specifying an attribute value different from the attribute value of the existing *hgroup* will result in an error. You cannot change the attribute value of an existing *hgroup*.

You can specify one of the following combinations to *attribute* and *value*.

If multiple attributes are specified and any error occurs, the entire process is canceled.

type=mirror or **type=stripe** (default is **mirror**)

Sets the type attribute of *hgroup*.

mirror

Sets the type attribute to "mirror".

stripe

Sets the type attribute to "stripe".

width=blks (default is 32)

Sets the stripe width of *hgroup*. *blks* indicates the stripe width in blocks (decimal numbers). One block is 512 bytes. For *blks*, you can specify an integer that is the power of two, equal to or larger than one and equal to or smaller than the available size of the smallest shadow group among groups specified by *lgroup*,... If *hgroup* is not the stripe type, this option will cause an error.

-c *class*

class indicates the name of the shadow class to which the shadow group as the object of the operation belongs.

-g *group* (when using **-R**)

group indicates the name of the shadow group that is the object of the operation.

-h *hgroup* (when using **-C,-D**)

hgroup indicates the name of the higher level shadow group to which the lower level shadow group as the object of the operation is connected or is to be connected.

-l *lgroup* (when using **-D**)

lgroup indicates the name of the lower level shadow group as the object of the operation.

-l *lgroup*,... (when using **-C**)

lgroup indicates the name of the lower level shadow group as the object of the operation. To connect multiple shadow groups, separate each group name with a comma (,).

RETURNED VALUE

Upon successful completion, a value of 0 is returned.
Otherwise, a non-zero value is returned.

D.17 sdxshadowvolume - Shadow volume operations

SYNOPSIS

```
sdxshadowvolume -F -c class [-v volume,...]
```

```
sdxshadowvolume -M -c class {-g group | -d disk} -v volume -s size
```

```
sdxshadowvolume -N -c class [-v volume,...] [-e mode=val]
```

```
sdxshadowvolume -R -c class -v volume
```

DESCRIPTION

Use `sdxshadowvolume` to perform operations on shadow volumes specified by *volume*.
This command can be executed with superuser access privileges only.

PRIMARY OPTIONS

You have a choice of the following options.

-F

offline

Stops one or more shadow volumes specified by *volume*,... By default all shadow volumes within *class* are turned off. Volumes turned off cannot be accessed.

This results in an error if the volume is in use.

-M

Make

Creates a shadow volume specified by *volume* to the highest level shadow group specified by *group* or to a shadow disk of the single type specified by *disk*. *size* indicates the number of blocks on *volume*, *class* indicates the name of the shadow class associated with *group* or *disk*.

The access mode of the created shadow volume is ro (read-only). Synchronization copying is not performed on shadow volumes of the mirror type. The attribute of just resynchronization mode is off. In addition, the physical slice attribute is always off even though the slice configuring the shadow volume is registered with VTOC. A maximum of 1024 shadow volumes can be created to the same *group* or *disk*.

When creation is completed, the shadow volume is activated and can be accessed using the following special files:

```
/dev/sfdsk/class/rdisk/volume
/dev/sfdsk/class/dsk/volume
```

Create shadow volumes conforming to the following rules in order to use the shadow volumes to access data in the corresponding logical volumes.

- Equal to the corresponding logical volumes in size. For volume sizes, check the BLOCKS field displayed with the `sdxinfo -V` command.
- Whose first block numbers must be consistent with the first block numbers of the corresponding logical volumes. Therefore, create shadow volumes within the same shadow group or shadow disk in ascending order in conformity to the order of the first block numbers of the corresponding logical volumes. For the first block numbers of volumes, check the 1STBLK field displayed with the `sdxinfo -V` command.

The features of created shadow volumes when *group* is the mirror type and the stripe type are as below.

When group is a mirror type

In a shadow group of the mirror type, shadow volumes of the mirror type with mirror-multiplexing equal to the number of shadow disks or lower level shadow groups connected (maximum of eight) are created. When only one shadow disk or one lower level shadow group is connected, the created shadow volume will not be mirrored.

Synchronization copying is not performed on shadow volumes even if the shadow volumes of the mirror type are created. To ensure synchronization for a shadow volume of the mirror type, the mirror volume must be properly synchronized with GDS or SynfinityDisk that manages the mirror volume corresponding to the shadow volume.

When group is a stripe group

In a shadow group of the stripe type, shadow volumes of the stripe type with columns equal to the number of shadow disks or lower level shadow groups connected (maximum of 64) are created. When only one shadow disk or one lower level shadow group is connected, a shadow volume cannot be created.

-N

oNline

Activates one or more shadow volumes specified by *volume,...* By default all shadow volumes within *class* are activated. The activated shadow volumes become accessible.

Synchronization copying is not performed on shadow volumes if the shadow volumes of the mirror type are activated. To ensure synchronization for a shadow volume of the mirror type, the mirror volume must be properly synchronized with GDS or SynfinityDisk that manages the mirror volume corresponding to the shadow volume.

-R

Remove

Removes a shadow volume specified by *volume* and releases the disk area that the shadow group or the shadow disk of the single type was occupying.

This results in an error if the shadow volume is active.

No data stored on *volume* will be lost due to this removal of *volume*.

SUB OPTIONS

Sub options are as follows:

-c class

class indicates the name of the shadow class to which the shadow volume as the object of the operation belongs, or in which the shadow volume is to be created.

-d disk (when using -M)

disk indicates the name of a shadow disk of the single type in which the shadow volume of the single type will be created.

-e mode=*val* (when using -N)

Specifies the access mode for one or more shadow volumes that will be activated.

You can specify either of the following options to *val*.

rw

Sets the access mode to read and write.

ro

Sets the access mode to read only. Opening a read-only volume in the write mode will result in an error.

Although shadow volumes are activated in the access mode specified by *val*, the access mode attributes for the shadow volumes will remain unchanged. The access mode specified by *val* ("current access mode") is valid only while the shadow volume is active and will become invalid once the shadow volume is stopped. When the shadow volume is restarted, it will start in the mode according to the access mode attribute ("default access mode") unless the access mode is specified by this option.

In order to start a shadow volume that is already activated on the current node in a different access mode, you must first stop the shadow volume.

-g group (when using -M)

group indicates the name of a shadow group in which the shadow volume will be created.

-s size (when using -M)

Specifies the size of the *volume* being created in blocks (decimal numbers). One block is 512 bytes.

When *group* is the stripe type, the size of *volume* created will be the *size* rounded up to a common multiple of the stripe width multiplied by stripe columns and the cylinder size. In other cases, the size of *volume* created will be the *size* rounded up to the integer multiple of the cylinder size.

-v volume (when using -M,-R)

volume indicates the name of the shadow volume as the object of the operation.

-v volume,... (when using -F,-N)

volume,... indicates the names of one or more shadow volumes as the objects of the operation. To specify multiple shadow volumes, separate each volume name with a comma (.).

RETURNED VALUE

Upon successful completion, a value of 0 is returned.
Otherwise, a non-zero value is returned.

D.18 Volume Creation Using Command

In this section, operation outline of volume creation is explained.
Please use as a reference when configuring the environment.
For details, see the Command Reference.

Information

For the order of mirroring system disks, see "USAGE EXAMPLES" in "D.11 sdxroot - Root file system mirroring definition and cancellation."

Note

For PRIMECLUSTER Systems

In order to define the configuration of GDS objects such as classes and volumes, PRIMECLUSTER resources must be registered in advance. For details on registering resources, see "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide."

(1) Creating a mirror volume

The procedure which carries out the mirroring of physical disks c1t0d0 and c2t0d0, and creates a volume as an example is shown.

1) Registering disks to class

Register the physical disks with a class. When the specified class does not exist, it will be created automatically.

Example) Registering physical disks c1t0d0 and c2t0d0 to local class "Local", and name these "Disk1" and "Disk2".

```
# sdxdisk -M -a type=local -c Class1 -d c1t0d0=Disk1,c2t0d0=Disk2
```

Information

When registering disks to shared class, you must specify the "scope" attribute with the -a option. See an example below.

```
# sdxdisk -M -c Class1 -a type=shared,scope=node1:node2 -d  
c1t0d0=Disk1,c2t0d0=Disk2
```

2) Connecting the disks to a mirror group

Connect the disks to a mirror group. When the specified mirror group does not exist, it will be created automatically.

Example: Connecting "Disk1" and "Disk2" to mirror group "Group1".

```
# sdxdisk -C -c Class1 -g Group1 -d Disk1,Disk2
```

3) Creating a mirror volume

Create a volume within the mirror group.

Example: Creating a volume of 1,000 blocks within mirror group "Group1", and assigning a volume name "Volume1".

```
# sdxvolume -M -c Class1 -g Group1 -v Volume1 -s 1000
```

After returning from the command, synchronization copying will automatically begin.

4) Confirming the completion of procedure

Confirm that the synchronization copying is complete.

Example: Confirming synchronization copying of volume "Volume1" is complete.

```
# sdxinfo -S -o Volume1
OBJ    CLASS  GROUP  DISK    VOLUME  STATUS
-----
slice  Class1  Group1  Disk1   Volume1 ACTIVE
slice  Class1  Group1  Disk2   Volume1 COPY
```

If the all the displayed slices' STATUS fields are "ACTIVE," synchronization copying is complete.

If the synchronization copying is still in progress, "COPY" will be displayed in the STATUS field.

Using the -e long option, you can check the progress of the synchronization copying.

For details, see "[D.6 sdxinfo - Display object configuration and status information.](#)"

(2) Creating a single volume

The procedure which creates a single volume as an example using physical disk c1t0d0 is shown.

1) Registering a disk to class

Register the physical disk with a class. When the specified class does not exist, it will be created automatically.

Example: Registering physical disk c1t0d0 to shared class "Shared", and assigning the name "Disk1".

```
# sdxdisk -M -c Class1 -a type=shared,scope=node1:node2 -d
c1t0d0=Disk1:single
```

2) Creating a single volume

Create a volume within the single disk.

Example: Creating a volume of 1,000 blocks within single disk "Disk 1", and assigning a volume name "Volume1".

```
# sdxvolume -M -c Class1 -d Disk1 -v Volume1 -s 1000
```

(3) Creating a large-capacity volume (using concatenation)

The procedure which concatenates physical disks c1t0d0 and c2t0d0, and creates a volume as an example is shown.

1) Registering disks to class

Register the physical disks with a class. When the specified class does not exist, it will be created automatically.

Example: Registering physical disks c1t0d0 and c2t0d0 to shared class "Shared", and assigning the names "Disk1" and "Disk2" respectively.

```
# sdxdisk -M -c Class1 -a type=shared,scope=node1:node2 -d
c1t0d0=Disk1,c2t0d0=Disk2
```

2) Connecting the disks to a concatenation group

Connect the disks to a concatenation group. When the specified concatenation group does not exist, it will be created automatically.

Example: Connecting "Disk1" and "Disk2" to concatenation group "Group1".

```
# sdxdisk -C -c Class1 -g Group1 -d Disk1,Disk2 -a type=concat
```

3) Creating a large-capacity volume

Create a volume within the concatenation group.

Example: Creating a volume of 1,000,000,000 blocks within concatenation group "Group1", and assigning a volume name "Volume1".

```
# sdxvolume -M -c Class1 -g Group1 -v Volume1 -s 1000000000 -a pslice=off
```

(4) Creating a stripe volume

The procedure which carries out striping of physical disk c1t0d0 and c2t0d0, and creates volume as an example is shown.

1) Registering disks to class

Register the physical disks with a class. When the specified class does not exist, it will be created automatically.

Example: Registering physical disks c1t0d0 and c2t0d0 to shared class "Shared", and assigning the names "Disk1" and "Disk2" respectively.

```
# sdxdisk -M -c Class1 -a type=shared,scope=node1:node2 -d  
c1t0d0=Disk1,c2t0d0=Disk2
```

2) Connecting the disks to a stripe group

Connect the disks to a stripe group. When the stripe group does not exist, it will be created automatically.

Example: Connecting "Disk1" and "Disk2" to stripe group "Group1".

```
# sdxdisk -C -c Class1 -g Group1 -d Disk1,Disk2 -a type=stripe,width=32
```

3) Creating a stripe volume

Create a volume within the stripe disk.

Example: Creating a volume of 1,000 blocks within stripe group "Group1", and assigning a volume name "Volume1".

```
# sdxvolume -M -c Class1 -g Group1 -v Volume1 -s 1000 -a pslice=off
```

(5) Creating a mirror volume (Combining striping and mirroring)

Example: Configuring stripe groups with physical disks c1t0d0 and c2t0d0, and c3t0d0 and c4t0d0 respectively, and creating a volume by mirroring those two stripe groups.

1) Registering disks to class

Register the physical disks with a class. When the specified class does not exist, it will be created automatically.

Example: Registering physical disks c1t0d0, c2t0d0, c3t0d0 and c4t0d0 to shared class "Shared", and assigning the names "Disk1", "Disk2", "Disk3" and "Disk4" respectively.

```
# sdxdisk -M -c Class1 -a type=shared,scope=node1:node2 \  
-d c1t0d0=Disk1,c2t0d0=Disk2,c3t0d0=Disk3,c4t0d0=Disk4
```

2) Connecting the disks to a stripe group

Connect the disks to a stripe group. When the stripe group does not exist, it will be created automatically.

Example: Connecting "Disk1" and "Disk2" to stripe group "Group1".

```
# sdxdisk -C -c Class1 -g Group1 -d Disk1,Disk2 -a type=stripe,width=32
```

Connecting Disk3 and Disk4 to stripe group Group2.

```
# sdxdisk -C -c Class1 -g Group2 -d Disk3,Disk4 -a type=stripe,width=32
```

3) Connecting the stripe groups to a mirror group

Connect the stripe groups to a mirror group. When the specified mirror group does not exist, it will be created automatically.

Example: Connecting the stripe group "Group1" and "Group2" to mirror group "Group3".

```
# sdxgroup -C -c Class1 -h Group3 -l Group1,Group2 -a type=mirror
```

The "-a type=mirror" option is omissible.

4) Creating a mirror volume

Create a volume within the highest level mirror group.

Example: Creating a volume of 1,000 blocks within mirror group "Group3", and assigning a volume name "Volume1"

```
# sdxvolume -M -c Class1 -g Group3 -v Volume1 -s 1000 -a pslice=off
```

After returning from the command, synchronization copying will automatically begin.

5) Confirming the completion of procedure

Confirm that the synchronization copying is complete.

Example: Confirming synchronization copying of volume "Volume1" is complete.

```
# sdxinfo -S -o Volume1  
OBJ    CLASS  GROUP  DISK    VOLUME  STATUS  
-----  
slice  Class1  Group3  Group1  Volume1  ACTIVE  
slice  Class1  Group3  Group2  Volume1  COPY
```

If the all the displayed slices' STATUS fields are "ACTIVE," synchronization copying is complete.

If the synchronization copying is still in progress, "COPY" will be displayed in the STATUS field.

Using the -e long option, you can check the progress of the synchronization copying.

For details, see "[D.6 sdxinfo - Display object configuration and status information.](#)"

D.19 Snapshot Creation Using Command

In this section, operation outline of snapshot creation using command is explained.
Please use it as a reference when configuring the environment.
For details, see the Command Reference.

1) Joining the proxy volume with master volume

Join the proxy volume with the master volume, and copy the master volume data to the proxy volume. Before joining them, stop the proxy volume.

Information

When Class1 is a shared class, stop Volume2 on all nodes in the class scope by using the `sdxvolume -F` command with the `-e allnodes` option.

Example: Master volume Volume1 and proxy volume Volume2 will be joined.

```
# sdxvolume -F -c Class1 -v Volume2
# sdxproxy Join -c Class1 -m Volume1 -p Volume2
```

After returning from the command, synchronization copying will automatically be performed.

Information

`-m` option and `-p` option can indicate groups as well as volumes. When indicating a group, all volumes within the group will be copied. When specifying groups, it is necessary to use the `-a` option.

2) Confirming the completion of copying

Confirm that the synchronization copying is complete.

Example: Confirming synchronization copying of proxy volume (Volume2) is complete.

```
# sdxinfo -S
OBJ    CLASS  GROUP  DISK    VOLUME  STATUS
-----
slice  Class1  Group1  Disk1   Volume1  ACTIVE
slice  Class1  Group1  Disk2   Volume1  ACTIVE
slice  Class1  Group2  Disk3   Volume2  STOP
slice  Class1  Group2  Disk4   Volume2  STOP
```

If all the displayed slices' STATUS fields are "STOP," synchronization copying is complete.

If the synchronization copying is still in progress, "COPY" will be displayed in the STATUS field.

3) Creating a snapshot

In order to create a snapshot, part the proxy volume from the master volume after confirming that synchronization copying is complete.

```
# sdxproxy Part -c Class1 -p Volume2
```

Note

In order to create a snapshot properly, you should either stop the application, or secure integrity at the file system layer or database layer that is managing the data. For instance, if you are using the file system, use the `lockfs(1M)` command to write-lock the file system before separating the proxy volume from the master volume.

4) Creating a backup

In order to make a data backup from the snapshot you created, follow the procedures of whatever backup tool you are using.

If you are not using the snapshot function after creating the backup, cancel the master-proxy relationship.

5) Rejoining the proxy volume with master volume

In order to synchronize the proxy volume with the master volume again, remove the created snapshot, and rejoin the proxy volume with the master volume. Before rejoining them, stop the proxy volume.

Information

If Class1 is a shared class, use the `-e allnodes` option with the `sdxvolume -F` command to stop Volume2 on all the nodes within the class scope.

```
# sdxvolume -F -c Class1 -v Volume2
# sdxproxy Rejoin -c Class1 -p Volume2
```

Synchronization copying will be performed as in procedure 1).

When creating more than one backup using the same proxy volume, repeat procedures 2) to 4).

6) Canceling the master-proxy relationship

Cancel the master-proxy relationship, and finish using the snapshot function.

```
# sdxproxy Break -c Class1 -p Volume2
```

D.20 Volume Expansion Using Commands

This section discusses the procedures for expanding the sizes of `/`, `/usr`, and `/var` file systems during services using the Snapshot function of GDS Snapshot.

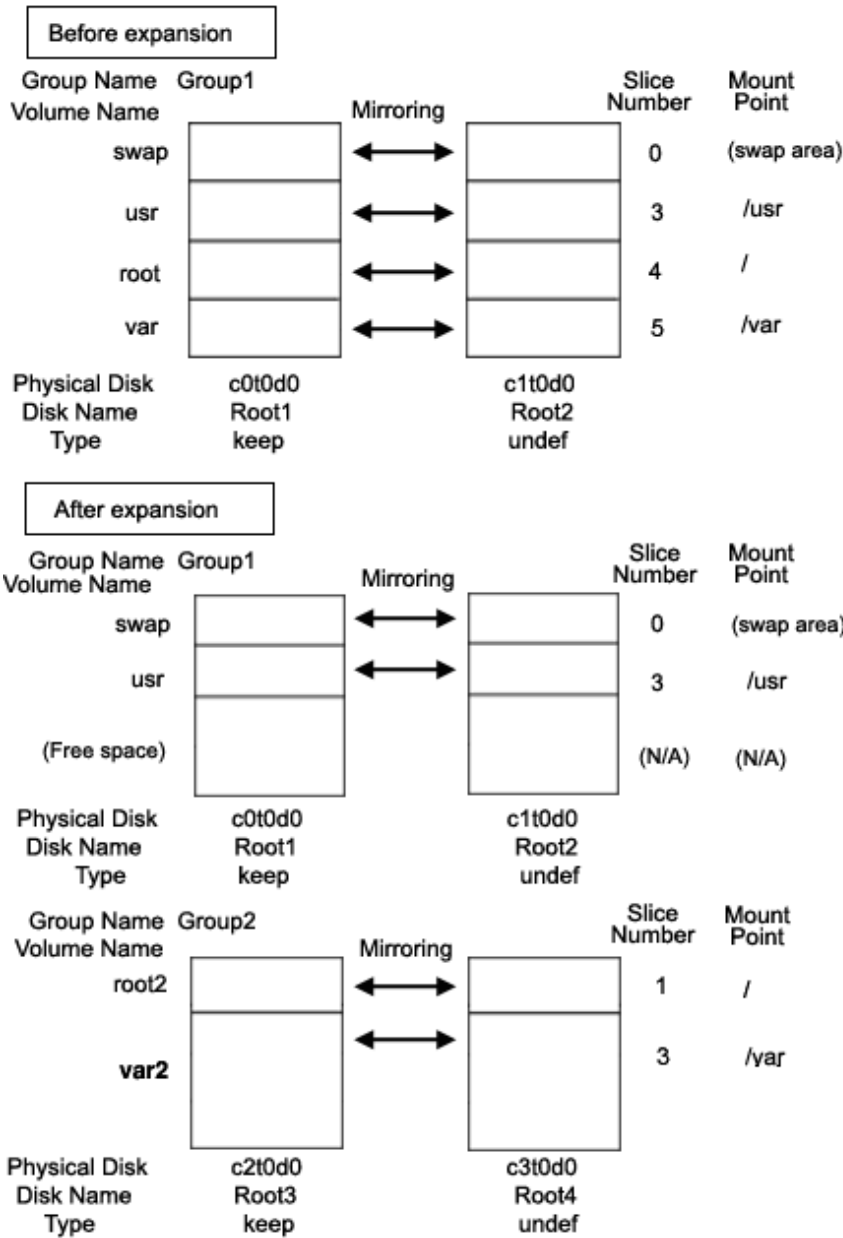
Note

In this procedure, to expand the `ufs` file system size without stopping the operation, use the Solaris operating system's `growfs(1M)` command.

[Procedure]

The following describes the procedure for expanding the `/var` file system area in the configuration as shown below.

As shown in the following figure, a `/` file system must be migrated to another volume regardless of whether or not it is expanded in size.



 **Note**

The cylinder sizes of keep disks Root1 and Root3 must be equal.

 **Information**

System disk mirroring is not a requirement. The configuration without c0t1d0 and c3t0d0 that are shown in the figure above is also supported. However, in systems that require high availability, it is recommended to use the mirroring configuration as above.

1) Mirroring system disks

1-1) Exit all active application programs.

To ensure safe mirroring, exit all running application programs. When higher security is required, you should make backups of system disks.

1-2) Register system disks with the root class.

```
# sdxdisk -M -c System -a type=root -d c0t0d0=Root1:keep,c1t0d0=Root2:undef
```

1-3) Connect system disks to groups respectively.

```
# sdxdisk -C -c System -g Group1 -d Root1,Root2 -v 0=swap:on,3=usr:on,4=root:on,5=var:on
```

1-4) Confirm that the mirror definition is complete.

```
# sdxroot -M -c System -d Root1
```

1-5) Reboot the system.


```
# shutdown -y -g0 -i6
```

1-6) Confirm that mirroring is complete.

Use the mount(IM) command or the sdxinfo command to verify that the system disks have been mirrored properly.

2) Creating proxy volumes

Create proxy volumes that expand / file systems and the sizes. At this point, only one keep disk should be connected to the group in which proxy groups are created. The following describes the procedure for creating proxy volumes in the configuration as shown below.



 Point

Sufficient free space is required after the last block of proxy volumes in a /var file system to be expanded in size.

2-1) View the current volume sizes.

```
# sdxinfo -v -c System
OBJ  NAME  CLASS  GROUP  SKIP  JRM  1STBLK  LASTBLK  BLOCKS  STATUS
-----
volume swap  System  Group1  off  on      0  1049759  1049760  1049760  ACTIVE
volume *     System  Group1  *    *    1049760  1071359  21600  PRIVATE
volume *     System  Group1  *    *    1071360  2946239  1874880  FREE
volume usr   System  Group1  off  on    2946240  2965679  19440  ACTIVE
volume root  System  Group1  off  on    2965680  3140639  174960  ACTIVE
volume var   System  Group1  off  on    3140640  3376079  235440  ACTIVE
```

The volume sizes can be viewed in the BLOCKS field provided by the `sdxinfo -V` command. In this example, the root size is 174960 blocks and the var size is 235440 blocks.

2-2) Register the disks with the root class.

```
# sdxdisk -M -c System -d c2t0d0=Root3:keep,c3t0d0=Root4:undef
```

Note

When registering multiple keep disks with a class together, as many or more undefined disks must also be registered.

2-3) Connect only one keep disk to the group.

```
# sdxdisk -C -c System -g Group2 -d Root3
```

2-4) Create volumes.

The volume sizes should correspond to those shown in step 2-1).

```
# sdxvolume -M -c System -g Group2 -v root2 -s 174960
# sdxvolume -M -c System -g Group2 -v var2 -s 235440
```


2-5) View the sizes of the created volumes.

Verify that the sizes of the volumes created in step 2-4) correspond to those shown in step 2-1).

```
# sdxinfo -V -c System
OBJ    NAME    CLASS  GROUP  SKIP  JRM  1STBLK  LASTBLK  BLOCKS  STATUS
-----
volume swap    System Group1  off  on      0  1049759  1049760  1049760  ACTIVE
volume *       System Group1  *    *    1049760  1071359   21600  PRIVATE
volume *       System Group1  *    *    1071360  2946239  1874880  FREE
volume usr     System Group1  off  on    2946240  2965679   19440  ACTIVE
volume root    System Group1  off  on    2965680  3140639  174960  ACTIVE
volume var     System Group1  off  on    3140640  3376079  235440  ACTIVE
volume *       System Group2  *    *      0    21599    21600  PRIVATE
volume root2   System Group2  off  on     21600   196559   174960  ACTIVE
volume var2    System Group2  off  on    196560   431999   235440  ACTIVE
volume *       System Group2  *    *    432000  3376079  2944080  FREE
```

3) Joining the proxy volumes

Data in / and /var can be copied in proxy volumes by joining the proxy volumes to the volumes in / and /var (master volumes). The following describes the procedure for joining proxy volumes in the configuration as shown below.



3-1) Join the proxy volumes.

```
# sdxvolume -F -c System -v root2,var2
# sdxproxy Join -c System -m root -p root2
# sdxproxy Join -c System -m var -p var2
```

3-2) Confirm that synchronization copying is complete.

```
# sdxinfo -S -c System
OBJ   CLASS  GROUP  DISK   VOLUME  STATUS
-----
slice System  Group2 Root1   swap    ACTIVE
slice System  Group2 Root2   swap    ACTIVE
slice System  Group1 Root1   usr     ACTIVE
slice System  Group1 Root2   usr     ACTIVE
slice System  Group1 Root1   root    ACTIVE
slice System  Group1 Root2   root    ACTIVE
slice System  Group1 Root1   var     ACTIVE
slice System  Group1 Root2   var     ACTIVE
slice System  Group2 Root3   root2   STOP
slice System  Group2 Root3   var2    COPY
```

When synchronization copying is in process, COPY is displayed in the STATUS field for slices of proxy volumes. When the STATUS of every proxy volume is in STOP, the synchronization copying is completed.

4) Parting the proxy volumes

Once synchronization copying is completed, the master volume and the proxy volume become synchronized. Snapshots of a master volume can be collected in a proxy volume by parting those synchronized master volume and proxy volume.

4-1) Secure consistency of the file systems.

To secure consistency of snapshot file systems, the file systems must be prevented from being updated. However, file systems such as /, /usr, and /var are necessary for the system to operate, they cannot be unmounted during the system operation. You can follow the procedure as below in order to reduce write occurrences into system disks and write occurrences into system disks that are not updated yet.

- a. Activate the system in the single user mode. (This can be skipped.)
- b. Exit all active application programs writing in the system disks. (This can be skipped.)
- c. Execute the sync(1M) command to write file system data not written in yet but updated on the memory into the disks.

Note

The system may hang up if a / file system is locked not to be written in with the lockfs(1M) command. In addition, GDS and GDS Snapshot commands hang up if a /var file system is locked not to be written in. Therefore, / and /var should not be applied the write lock.

Even if all the steps a., b., and c. are enforced, it cannot completely restrain the file systems from being updated. As a result, inconsistency similar to that after a system panic might occur in snapshot file systems.

When a., b., and c. are all enforced, a snapshot file system will be similar to a file system after a panic occurs in the single user mode.

When only c. is enforced skipping a. and b., s snapshot file system will be similar to a file system after a panic occurs during the system operation.

In any of these occasions, a file system may have inconsistency, and the file system need to be checked for consistency and repaired as described in step 5-1).

4-2) Part the proxy volumes.

```
# sdxproxy Part -c System -p root2,var2
```

4-3) When the system was activated in the single user mode in a. of step 4-1), reboot it in the multi-user mode.

4-4) When application programs were exited in b. of step 4-1), launch the application programs.

5) Configuring an alternative boot environment

Configure the environment to boot from the proxy volumes.

5-1) Check and repair the file systems on the proxy volumes.

There may be inconsistency in file systems on proxy volumes, and so verify and repair them using the fsck(1M) command.

```
# fsck -y /dev/sfdsk/System/rdisk/root2
# fsck -y /dev/sfdsk/System/rdisk/var2
```

5-2) Configure the alternative boot environment.

```
# sdxproxy Root -c System -p root2,var2
```

Once the alternative boot environment is configured, the following message is output.

```
SDX:sdxproxy: INFO: completed definitions of alternative boot environment:
current-boot-device=/pci@1f,4000/scsi@3/disk@0,0:a /pci@1f,4000/scsi@3/
disk@1,0:a
alternative-boot-device=/pci@1f,4000/scsi@3/disk@2,0:a
```

Be sure to keep a copy of the output boot device names in the current boot environment (current-boot-device value) and in the alternative boot environment (alternative-boot-device value).

6) Expanding the snapshot size

6-1) Cancel the master-proxy relationship.

```
# sdxproxy Break -c System -p root2
# sdxproxy Break -c System -p var2
```

6-2) Expand the snapshot volume size.

The following example shows how to expand the snapshot volume size in /var to 706320 blocks.

```
# sdxvolume -s -c System -v var2 -s 706320
```

6-3) Expand the snapshot file system size.

The following example shows how to expand the snapshot file system size to 706320 blocks.

```
# growfs -s 706320 /dev/sfdsk/System/rdsk/var2
```



See

.....
For details about the growfs(1M) command, see the Solaris operating system manual.
.....

7) Mirroring the snapshot

A snapshot volume can be mirrored in the manner that a disk is added to the group to which the snapshot volume belongs.

```
# sdxdisk -C -c System -g Group2 -d Root4
```

8) Switching to the alternative boot environment

By switching to the alternative boot environment, /var file systems can be switched to the volumes expanded in size.

8-1) Boot the system through the alternative boot environment.

Execute the boot command in the OpenBoot environment. For the boot command, you can specify one of the boot devices in the alternative boot environment output in the message as shown in step 5-2).

```
# shutdown -y -g0 -i0  
ok boot /pci@1f,4000/scsi@3/disk@2,0:a
```



Information

.....
The boot command may fail as shown below if the number of the boot device names in the alternative boot environment output in step 5-2) exceeds the number of boot devices in the alternative boot environment.
.....

```
ok boot /pci@1f,4000/scsi@3/disk@2,0:a  
Rebooting with command: boot /pci@1f,4000/scsi@3/disk@2,0:a  
Boot device: /pci@1f,4000/scsi@3/disk@2,0:a File and args:  
Evaluating: boot /pci@1f,4000/scsi@3/disk@2,0:a  
  
Can't open boot device  
  
ok
```

If this happens, you should specify another boot device name in the alternative boot environment output in step 5-2) and try the boot command again.
.....

8-2) Confirm that it was booted normally.

Using the mount(1M) command or the sdxinfo command, make sure that it was booted normally in the alternative boot environment and that GDS objects do not contain errors. Additionally, according to need, you should also confirm that file system contents in the alternative boot environment are proper and that applications can normally run.

Information

In the case of an abnormal boot, restore the original boot environment. To restore the original boot environment, specify one of the boot devices in the current boot environment output in the message as shown in step 5-2) using the boot command in the OpenBoot environment.

```
# shutdown -y -g0 -i0
ok boot /pci@1f,4000/scsi@3/disk@0,0:a
```

9) Removing unnecessary volumes

After a normal boot in the alternative boot environment is verified, remove volumes in the / file system in the previous boot environment and in the /var file system before expansion.

```
# sdxvolume -F -c System -v root,var
# sdxvolume -R -c System -v root
# sdxvolume -R -c System -v var
```

Appendix E GDS Messages

This section explains the messages for GDS drivers, daemons and commands.

E.1 Web-Based Admin View Messages (0001-0099)

For message numbers 0001 to 0099, see "Web-Based Admin View Operation Guide" for details.

E.2 Driver Messages

Driver messages are output to a log file or console via the syslog interface.

Message Numbers

The message numbers used to identify messages described in E.2.1 to E.2.4 do not appear in messages actually output by GDS drivers.

Variable Names

Italicized words in the messages are variable names, and the actual output will vary depending on the situation. The meaning and the format of the variable names used in the message explanations are described below.

Variable names	Descriptions
<i>v_devno</i>	Volume device number (hexadecimal number)
<i>v_maj</i>	Volume major number (decimal number)
<i>v_min</i>	Volume minor number (decimal number)
<i>p_devno</i>	Physical slice device number (hexadecimal number)
<i>p_maj</i>	Physical slice major number (decimal number)
<i>p_min</i>	Physical slice minor number (decimal number)
<i>pdevtree</i>	Physical disk device tree path name. e.g.) "/pci@4,2000/pci@1/sunw,isp200@4/sd@2,0"
<i>pdevinst</i>	Physical disk instance name. e.g.) sd77
<i>blknodk</i>	I/O requested disk block number (decimal number)
<i>blknosl</i>	I/O requested slice block number (decimal number)
<i>length</i>	I/O transfer requested block size (in bytes, decimal number)
<i>resid</i>	Residing block size (in bytes, decimal number)
<i>errno</i>	System call error number (decimal number)
<i>oflag</i>	Open flag (hexadecimal number)
<i>second</i>	Elapsed time (in seconds, decimal number)
<i>lbolt</i>	Elapsed time after system boot (in ticks, hexadecimal number)
<i>details</i>	Details

Explanation

Messages output by the driver are shown below in the order of severity. There are three levels of severity.

Level of importance	Descriptions	facility.level
PANIC	This message is displayed when an event that will stop the system is detected.	kern.notice
WARNING	This message is output when an abnormal event is detected. It will not always result in an error.	kern.warning
NOTICE	This message is output to record driver operation and the like.	kern.notice

facility.level is the priority of a message passed from the GDS driver to syslogd(1M). The output destinations of messages are defined in the /etc/syslog.conf configuration file of syslogd(1M) and can be changed by modifying the definitions in /etc/syslog.conf. For details, see the description about syslog.conf(4) in this manual.

E.2.1 Panic Messages (20000)

20000

```
sfdsk: no valid root disk found, disk ID information is invalid
```

Explanation

Disk ID information was invalid for all root disks (disks where root file system will be mounted).

Resolution

Check for improper root disk swap or connection change of I/O cables.

When the cause cannot be specified, collect investigation material including the memory dump and contact your local customer support.

E.2.2 Warning Messages (22000-22099)

22000

```
WARNING: sfdsk: read error on slice:

volume info : devno(maj,min)=v_devno(v_maj,v_min)

device info : devno(maj,min)=p_devno(p_maj,p_min)

                devname=pdevtree (pdevinst)

error info : blk in disk=blknodk, blk in slice=blknos1

                length=length, resid=resid, errno=errno
```

Explanation

Read request sent to slice terminated abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

22001

WARNING: sfdsk: write error on slice:

```
volume info : devno(maj,min)=v_devno(v_maj,v_min)
device info  : devno(maj,min)=p_devno(p_maj,p_min)
              devname=pdevtree (pdevinst)
error info   : blk in disk=blknodk, blk in slice=blknos1
              length=length, resid=resid, errno=errno
```

Explanation

Write request sent to slice terminated abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk driver.

22002

WARNING: sfdsk: read error on disk:

```
device info : devno(maj,min)=p_devno(p_maj,p_min)
              devname=pdevtree (pdevinst)
error info   : blk in disk=blknodk, blk in slice=blknos1
              length=length, resid=resid, errno=errno
```

Explanation

Read request sent to disk terminated abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

22003

WARNING: sfdsk: write error on disk:

```
device info : devno(maj,min)=p_devno(p_maj,p_min)
              devname=pdevtree (pdevinst)
error info   : blk in disk=blknodk, blk in slice=blknos1
              length=length, resid=resid, errno=errno
```

Explanation

Write request sent to disk terminated abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

22004

WARNING: sfdsk: read and writeback error on slice:

```
volume info : devno(maj,min)=v_devno(v_maj,v_min)

device info : devno(maj,min)=p_devno(p_maj,p_min)

                devname=pdevtree (pdevinst)

error info : blk in disk=blknodk, blk in slice=blknos1

                length=length, resid=resid, errno=errno
```

Explanation

Read request and writeback request sent to slice terminated abnormally.
Writeback is a process to read data from other slices in the event of a read error.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

22005

WARNING: sfdsk: open error on slice:

```
volume info : devno(maj,min)=v_devno(v_maj,v_min)

device info : devno(maj,min)=p_devno(p_maj,p_min)

                devname=pdevtree (pdevinst)

error info : blk in disk=blknodk, blk in slice=blknos1

                oflag=oflag, errno=errno
```

Explanation

Open request sent to slice returned abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

22006

WARNING: sfdsk: open error on disk:

```
device info : devno(maj,min)=p_devno(p_maj,p_min)

                devname=pdevtree (pdevinst)

error info : oflag=oflag, errno=errno
```

Explanation

Open request sent to disk returned abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

22007

WARNING: sfdsk: close error on disk:

```
device info : devno(maj,min)=p_devno(p_maj,p_min)
              devname=pdevtree (pdevinst)
error info  : oflag=oflag, errno=errno
```

Explanation

Close request sent to disk returned abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

22008

WARNING: sfdsk: NVURM read error on disk:

```
volume info : devno(maj,min)=v_devno(v_maj,v_min)
device info  : devno(maj,min)=p_devno(p_maj,p_min)
              devname=pdevtree (pdevinst)
error info   : blk in disk=blknodk, blk in slice=blknos1
              length=length, resid=resid, errno=errno
```

Explanation

NVURM read request sent to disk terminated abnormally.

NVURM is volume update area map information which is stored on the disk for just resynchronization.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

22009

WARNING: sfdsk: NVURM write error on disk:

```
volume info : devno(maj,min)=v_devno(v_maj,v_min)
device info  : devno(maj,min)=p_devno(p_maj,p_min)
              devname=pdevtree (pdevinst)
error info   : blk in disk=blknodk, blk in slice=blknos1
              length=length, resid=resid, errno=errno
```

Explanation

NVURM write request sent to disk terminated abnormally.

NVURM is volume update area map information which is stored on the disk for just resynchronization.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

22010

```
WARNING: sfdsk: volume status log write error on disk:
```

```
device info : devno(maj,min)=p_devno(p_maj,p_min)
```

```
devname=pdevtree (pdevinst)
```

```
error info : blk in disk=blknodk, blk in slice=blknos1
```

```
length=length, resid=resid, errno=errno
```

Explanation

Write request of volume status log sent to disk terminated abnormally.

Volume status log records if the volume closed normally in the event of a system crash.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

22011

```
WARNING: sfdsk: failed to abort I/O requests on disk:
```

```
device info : devno(maj,min)=p_devno(p_maj,p_min)
```

```
devname=pdevtree (pdevinst)
```

```
error info : errno=errno
```

Explanation

The request to cancel I/O requests on the physical device indicated in device info (an ioctl request for the mphp or mplb driver) ended abnormally.

Resolution

Collect investigation material and contact your local customer support.

E.2.3 Information Messages (24000-24099)

24000

```
sfdsk: driver started up
```

Explanation

The driver has been installed into the system.

24001

```
sfdisk: received shutdown request
```

Explanation

A shutdown request from the sdxservd daemon has been received.

24002

```
sfdisk: volume status log updated successfully,details
```

Explanation

Write request sent to the volume status log terminated normally.

Volume status log records if the volume closed normally when a system failure occurred.

24003

```
NOTICE: sfdisk: I/O error on slice:

    volume info: devno(maj,min)=v_devno(v_maj,v_min)

    device info: devno(maj,min)=p_devno(p_maj,p_min)

                devname=device

    error info:  blk in disk=blknodk, blk in slice=blknos1

                length=length, resid=resid, errno=errno

                flags=b_flags
```

Explanation

I/O request sent to slice terminated abnormally.

24004

```
NOTICE: sfdisk: read error and writeback success on slice:

    volume info : devno(maj,min)=v_devno(v_maj,v_min)

    device info : devno(maj,min)=p_devno(p_maj,p_min)

                devname=pdevtree (pdevinst)

    error info :  blk in disk=blknodk, blk in slice=blknos1

                length=length
```

Explanation

Read request sent to slice terminated abnormally, but has been recovered by writeback process.
Writeback is a process to read data from other slices in the event of a read error.

24005

```
NOTICE: sfdsk: trying to open slice:

    volume info: devno(maj,min)=v_devno(v_maj,v_min)

    device info: devno(maj,min)=p_devno(p_maj,p_min)

                devname=pdevtree(pdevinst)
```

Explanation

A request for opening the slice is issued.

24006

```
NOTICE: sfdsk: copy timeout. no response from sdxservd daemon:

    volume info: devno(maj,min)=v_devno(v_maj,v_min)
```

Explanation

Synchronization copying was not performed since there was no response from the sdxservd daemon.

24007

```
NOTICE: sfdsk: processing has taken long time on disk:

    device info: devno(maj,min)=p_devno(p_maj, p_min)

                devname=pdevtree (pdevinst)

    request info: lapsed seconds=second, start lbolt=lbolt

                details
```

Explanation

While *second* seconds have passed since the I/O request indicated in *details* was issued for the physical device indicated in device info, the I/O request is not complete yet.

24008

```
NOTICE: sfdsk: processing has taken long time on volume:

    volume info: devno(maj,min)=v_devno(v_maj,v_min)

    request info: lapsed seconds=second, start lbolt=lbolt

                details
```

Explanation

While *second* seconds have passed since the I/O request indicated in *details* was issued for the volume indicated in volume info, the I/O request is not complete yet.

24009

NOTICE: sfdsk: abort I/O requests on disk:

device info: devno(maj,min)=p_devno(p_maj,p_min)

devname=pdevtree (pdevinst)

Explanation

The process to cancel I/O requests on the physical disk indicated in device info has started.

24010

NOTICE: sfdsk: succeeded to abort I/O requests on disk:

device info: devno(maj,min)=p_devno(p_maj,p_min)

devname=pdevtree (pdevinst)

Explanation

The request to cancel I/O requests on the physical device indicated in device info (an ioctl request for the mphp or mplb driver) ended normally.

E.2.4 Internal Error Messages (26000)

26000

severity: sfdsk: internal error, details

Explanation

Internal error occurred. *details* indicates the cause of the error.
severity indicates the severity of the message.

Resolution

Collect investigation material and contact your local customer support.

E.3 Daemon Messages

Daemon messages are output to the GDS log file.

/var/opt/FJSVsdx/msglog/sdxservd.log

Output format to log file is as follows.

Mon Day HH:MM:SS SDX:daemon: severity: message

Mon gives the month the message was output, *Day* the date, *HH* the hour, *MM* the minute, *SS* the second, *daemon* the daemon program name, *severity* the severity of the message, and *message* the message text.

Depending on the settings of the syslog, daemon message will also be output to syslog log file and console.

Message Numbers

The message numbers used to identify messages described in E.3.1 to E.3.5 do not appear in messages actually output by GDS daemons.

Variable Names

Italicized words in the messages are variable names, and the actual output will vary depending on the situation. The meaning and the format of the variable names used in the message explanations are described below.

Variable names	Descriptions
<i>class</i>	Class name
<i>disk</i>	Disk name
<i>group</i>	Group name
<i>lgroup</i>	Lower level group name
<i>hgroup</i>	Higher level group name
<i>volume</i>	Volume name
<i>disk.volume</i>	Slice name
<i>object.volume</i>	Slice name
<i>proxy</i>	Proxy object name
<i>object</i>	Object name
<i>type</i>	Type attribute value
<i>status</i>	Object status
<i>device</i>	Physical disk name (<i>cCtTdD</i> , <i>mphdI</i> , <i>mplbI</i> , <i>emcpowerN</i>) where <i>C</i> is the controller number, <i>T</i> is the target ID, <i>D</i> is the disk number, <i>I</i> is the instance number and <i>N</i> is the emcpower device number.
<i>pslice</i>	Physical slice name (<i>cCtTdDs</i> {0-7}, <i>mphdIs</i> {0-7}, <i>mplbIs</i> {0-7}, <i>emcpowerN</i> {a-h}) where <i>C</i> is the controller number, <i>T</i> is the target ID, <i>D</i> is the disk number, <i>I</i> is the instance number, <i>N</i> is the emcpower device number, {0-7} is the slice number with 1 digit from 0 to 7, and {a-h} is the slice number with one alphabet from a to h. May be indicated by device tree file name. e.g.) "sd@2,0:a,raw"
<i>psdevtree</i>	Device tree path name to a physical slice e.g.) /devices/pci@1f,4000/scsi@3/sd@0,0:a
<i>v_devno</i>	Volume device number (hexadecimal number)
<i>v_maj</i>	Volume major number (decimal number)
<i>v_min</i>	Volume minor number (decimal number)
<i>p_devno</i>	Physical slice device number (hexadecimal number)
<i>p_maj</i>	Physical slice major number (decimal number)
<i>p_min</i>	Physical slice minor number (decimal number)

Variable names	Descriptions
<i>blknodk</i>	I/O requested disk block number (decimal number)
<i>blknosl</i>	I/O requested slice block number (decimal number)
<i>length</i>	I/O transfer requested block size (in bytes, decimal number)
<i>resid</i>	Residing block size (in bytes, decimal number)
<i>errno</i>	System call error number (decimal number)
<i>sdxerrno</i>	GDS defined internal error number (decimal number)
<i>node</i>	Node identifier, or node name
<i>oflag</i>	Open flag (hexadecimal number)
<i>ioctlcmd</i>	The ioctl command name
<i>timeofday</i>	Character string indicating current time and date
<i>details</i>	Details
<i>sdxfunc</i>	GDS function name
<i>exitstat</i>	Value indicating command end status (decimal number)
<i>cmdline</i>	Character indicating command line

Explanation

Messages output by the daemon are shown below in the order of its severity. There are four levels of severity.

Level of importance	Descriptions	facility.level
HALT	This message is output when an event that will halt all services provided by GDS is detected.	user.crit
ERROR	This message is output when an event that will halt some services is detected.	user.err
WARNING	This message is output when an abnormal event is detected. It will not always result in a service halt.	user.warning
INFO	This message is output to record the daemon operation and the like. Usually, it can be ignored.	user.info

facility.level is the priority of a message passed from the GDS driver to syslogd(1M). The output destinations of messages are defined in the /etc/syslog.conf configuration file of syslogd(1M) and can be changed by modifying the definitions in /etc/syslog.conf. For details, see the description about syslog.conf(4) in this manual.

E.3.1 Halt Messages (40000-40099)

40000

```
HALT: failed to create a new thread, errno=errno
```

Explanation

Function pthread_create() terminated abnormally. Process cannot be continued. The daemon process will be exited.

Resolution

When error number information *errno* is insufficient for identifying the cause and recovery, collect investigation material and contact your local customer support.

40001

HALT: cannot open driver administrative file, errno=errno

Explanation

GDS driver(sfdsk) administrative file cannot be opened. Process cannot be continued. The daemon process will be exited. This message is output when files under /dev/sfdsk directory cannot be accessed.

Resolution

Collect investigation material, and contact your local customer support.

40002

HALT: startup failed, no enough address space

Explanation

Startup failure due to unsuccessful memory allocation. Process cannot be continued. The daemon process will be exited.

Resolution

Confirm you have sufficient memory or swap area.

40003

HALT: failed to respawn daemon daemon, osfunc=osfunc, errno=errno

Explanation

The *daemon* terminated abnormally and failed to restart. The failure was caused by abnormal termination of OS *osfunc* function. The error number is *errno*. This message is output via syslog.

Resolution

When error number information is insufficient for identifying the cause, collect investigation material, and contact your local customer support.

40004

HALT: cannot start node-down recovery for remote node node, no enough space, osfunc=osfunc, errno=errno

Explanation

Could not recover the crashed remote node *node* due to unsuccessful memory allocation. Process cannot be continued. The daemon process will be exited.

Resolution

The OS *osfunc* function terminated abnormally, and error number is *errno*. Confirm you have sufficient memory or swap area, and recover.

E.3.2 Error Messages (42000-42099)

42000

```
ERROR: read error on status slice object.volume, class=class:

volume info:devno(maj,min)=v_devno(v_maj,v_min)

device info:devno(maj,min)=p_devno(p_maj,p_min)

        devname=device

error info: blk in disk=blknodk, blk in slice=blknosl,

        length=length, resid=resid, errno=errno
```

Explanation

Read request sent to slice *object.volume* in *status* status terminated abnormally. Read request sent to volume configured by this slice, or to slice accessible in isolation was returned as error.

You must attempt recovery as soon as possible since the application may cease normal operation.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and promptly recover the disk.

42001

```
ERROR: write error on status slice object.volume, class=class:

volume info:devno(maj,min)=v_devno(v_maj,v_min)

device info:devno(maj,min)=p_devno(p_maj,p_min)

        devname=device

error info: blk in disk=blknodk, blk in slice=blknosl,

        length=length, resid=resid, errno=errno
```

Explanation

Write request sent to slice *object.volume* in *status* status terminated abnormally. Write request sent to volume configured by this slice, or to slice accessible in isolation was returned as error.

You must attempt recovery as soon as possible since the application may cease normal operation.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and promptly recover the disk.

42002

```
ERROR: open error on status slice object.volume, class=class:

volume info:devno(maj,min)=v_devno(v_maj,v_min)

device info:devno(maj,min)=p_devno(p_maj,p_min)

        devname=device
```

```
error info: oflag=oflag, errno=errno
```

Explanation

Open request sent to slice *object.volume* in *status* status terminated abnormally. Open request sent to volume configured by this slice, or to slice accessible in isolation was returned as error.

You must attempt recovery as soon as possible since the application may cease to operate normally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and promptly recover the disk.

42003

```
ERROR: read error and writeback error on status slice object.volume, class=class:
```

```
volume info:devno(maj,min)=v_devno(v_maj,v_min)
```

```
device info:devno(maj,min)=p_devno(p_maj,p_min)
```

```
devname=device
```

```
error info: blk in disk=blknodk, blk in slice=blknosl,
```

```
length=length, resid=resid, errno=errno
```

Explanation

Read request and writeback request sent to slice *object.volume* in *status* status terminated abnormally. Read request sent to volume configured by this slice, or to slice accessible in isolation was returned as error.

You must attempt recovery as soon as possible since the application may not operate normally.

Writeback is a process to read data from other slices in the event of a read error.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and promptly recover the disk.

42004

```
ERROR: disk: device: ready for swapping due to I/O error, class=class
```

Explanation

Because an I/O error occurred on the disk *disk* in the class *class*, the physical disk *device* was excluded from GDS management and made exchangeable.

Resolution

Exchange the physical disk *device*, and then include it into GDS management executing `diskpowon(1M)`, the hot-swap disk attachment command that is a machine administrative function provided by Enhanced Support Facility (ESF).

For details, see "Machine Administration Guide."

However, if the error message "ERROR: device: setup for hot-swap failed" is output following this message on the console, recovery based on the resolution method related in that message is necessary.

42005

```
ERROR: device: setup for hot-swap failed
```

Explanation

The physical disk *device* was excluded from GDS management and made exchangeable, but preparations for hot swap with a machine administrative function of Enhanced Support Facility (ESF) failed.

Resolution

Exchange the physical disk *device*, and then include it into GDS management executing the `sdxswap -I` command or performing physical disk recovery in GDS Management View.

42006

```
ERROR: volume: closed down volume, class=class
```

Explanation

volume has closed down.

Resolution

Promptly attempt recovery by identifying the cause of failure by referring to GDS log message that was output immediately before the error.

For information on recovery, see "[F.1.3 Volume Status Abnormality](#)."

42007

```
ERROR: class: cannot startup class, no valid configuration database, sdxerrno=errno
```

Explanation

Could not startup *class*, since valid *class* configuration database could not be found.

This message is output when all disks (or the majority of disks) registered with *class* are unavailable.

Resolution

See "[F.1.4 Class Status Abnormality](#)."

42008

```
ERROR: class: cannot startup class, too few valid configuration database replicas, sdxerrno=errno
```

Explanation

Could not startup *class*, due to insufficient number of valid databases.

This message is output when many disks registered with *class* are unavailable.

Resolution

See "[F.1.4 Class Status Abnormality](#)."

42009

```
ERROR: class: closing down class, no valid configuration database
```

Explanation

class was closed since no valid *class* configuration database was found.

This message is output when all disks (or the majority of disks) registered with *class* are unavailable.

Resolution

See "[F.1.4 Class Status Abnormality](#)."

42010

ERROR: *class*: closing down class, too few valid configuration database replicas

Explanation

class was closed due to insufficient number of valid database.

This message is output when many disks registered with *class* are unavailable.

Resolution

See "[F.1.4 Class Status Abnormality](#)."

42011

ERROR: failed to send request message on node *node*, details

Explanation

Sending request message from *node* was unsuccessful.

Resolution

Collect investigation material and contact your local customer support.

42012

ERROR: timeout on receiving reply message from node *node*, details

Explanation

Timeout occurred while receiving a reply message from a remote node *node*.

Resolution

Collect investigation material and contact your local customer support.

42013

ERROR: rejected request message on remote node *node*, details

Explanation

Processing a request message on a remote node *node* was unsuccessful.

Resolution

Investigate the *node* message log and take necessary actions. If recovery is impossible, collect investigation material and contact your local customer support.

42014

ERROR: class: failed to start type volumes, status volume volume exists, node=node

Explanation

Starting volumes within the class *class* failed on the node *node* since the *status* volume *volume* exists. *type* is the *class* type attribute.

Resolution

volume is in abnormal status. First, you must recover normal status.
For information on recovery, see "[F.1.3 Volume Status Abnormality](#)."

42015

ERROR: class: failed to start and standby type volumes, status volume volume exists, node=node

Explanation

Starting and putting on standby volumes within the class *class* failed on the node *node* since the *status* volume *volume* exists. *type* is the *class* type attribute.

Resolution

volume is in abnormal status. First, you must recover normal status.
For information on recovery, see "[F.1.3 Volume Status Abnormality](#)."

42016

ERROR: class: failed to stop and standby type volumes, status volume volume exists, node=node

Explanation

Stopping and putting on standby volumes within the class *class* failed on the node *node* since the *status* volume *volume* exists. *type* is the *class* type attribute.

Resolution

volume is in abnormal status. First, you must recover normal status.
For information on recovery, see "[F.1.3 Volume Status Abnormality](#)."

42017

ERROR: class: failed to stop type volumes, status volume volume exists, node=node

Explanation

Stopping volumes within the class *class* failed on the node *node* since the *status* volume *volume* exists. *type* is the *class* type attribute.

Resolution

volume is in abnormal status. First, you must recover normal status.
For information on recovery, see "[F.1.3 Volume Status Abnormality](#)."

42018

ERROR: class: failed to start type volumes, class closed down, node=node

Explanation

Starting volumes failed since the class *class* has been closed down.
type is the *class* type attribute.

Resolution

Recover closed *class*. There may be multiple disk failures. Identify the cause based on object status, GDS log message, and syslog message.

For information on recovery, see "[F.1.4 Class Status Abnormality](#)."

42019

ERROR: class: failed to start and standby type volumes, class closed down, node=node

Explanation

Starting and putting on standby volumes failed. since the class *class* has been closed down.
type is the *class* type attribute.

Resolution

Recover closed *class*. There may be multiple disk failures. Identify the cause based on object status, GDS log message, and syslog message.

For information on recovery, see "[F.1.4 Class Status Abnormality](#)."

42020

ERROR: class: failed to stop and standby type volumes, class closed down, node=node

Explanation

Stopping and putting on standby volumes failed. Since the class *class* has been closed down.
type is the *class* type attribute.

Resolution

Recover closed *class*. There may be multiple disk failures. Identify the cause based on object status, GDS log message, and syslog message.

For information on recovery, see "[F.1.4 Class Status Abnormality](#)."

42021

ERROR: class: failed to stop type volumes, class closed down, node=node

Explanation

Stopping volumes failed since the class *class* has been closed down.
type is the *class* type attribute.

Resolution

Recover closed *class*. There may be multiple disk failures. Identify the cause based on object status, GDS log message, and syslog message.

For information on recovery, see "[F.1.4 Class Status Abnormality](#)."

42022

```
ERROR: class: closing down class, cluster-wide lock failure, sdxerrno=sdxerrno
```

Explanation

Abnormal exclusive control between cluster system nodes occurred. Since the process cannot be continued, *class* will be closed.

Resolution

Collect investigation material and contact your local customer support.

42023

```
ERROR: class: cannot startup class, cluster-wide lock failure, sdxerrno=errno
```

Explanation

Abnormal exclusive control between cluster system nodes occurred. Since the process cannot be continued, *class* could not be started.

Resolution

Collect investigation material and contact your local customer support.

42024

```
ERROR: class: closing down class, cluster communication failure, sdxerrno=sdxerrno
```

Explanation

Transmission failure between cluster system nodes occurred. Since the process cannot be continued, *class* will be closed.

Resolution

Collect investigation material and contact your local customer support.

42025

```
ERROR: class: cannot operate in cluster environment, created when cluster control facility not ready
```

Explanation

The class *class* cannot be used in a cluster environment since it was created when the cluster control facility was inactive. This message is output when one of the following operations is performed.

- a. After *class* was created automatic resource registration was performed on a node where automatic resource registration was not complete.
- b. In a cluster environment where automatic resource registration is complete, *class* was created in single user mode.
- c. A single node with *class* was changed over to a cluster system.

Resolution

See "(1) The error message "ERROR: class: cannot operate in cluster environment, ..." is output, and the operation cannot be conducted on the class class." in "F.1.9 Cluster System Related Error."

42026

```
ERROR: proxy: failed to copy with OPC, source=disk.volume, target=disk.volume, class=class
```

Explanation

While performing copying between proxy volume *proxy* and master volume with the OPC function, an I/O error occurred, and the copying process failed.

"source" specifies the original slice name, "target" the mirror slice name, and "class" to which the original slice and mirror slice belong.

Resolution

Identify the cause by referring to the log messages for GDS, disk driver, and Fujitsu Storage Systems ETERNUS that were output right before the occurrence of the error, and restore.

42027

```
ERROR: proxy: failed to copy with EC, source=disk.volume, target=disk.volume, class=class
```

Explanation

While performing copying between proxy volume *proxy* and master volume with the EC function, an I/O error occurred, and the copying process failed.

"source" specifies the original slice name, "target" the mirror slice name, and "class" to which the original slice and mirror slice belong.

Resolution

Identify the cause by referring to the log messages for GDS, disk driver, and Fujitsu Storage Systems ETERNUS that were output right before the occurrence of the error, and restore.

42028

```
ERROR: proxy: failed to copy with TimeFinder, source=disk, target=disk, class=class
```

Explanation

Copying failed due to an I/O error caused while conducting copying between the proxy group *proxy* and the master group with TimeFinder.

"source" is the name of the copy source disk, "target" is the name of the copy destination disk, and "class" is the name of the class to which the copy source and destination disks belong.

Resolution

Identify the cause by referring to the log messages for GDS, disk driver and Symmetrix device that were output right before the occurrence of the error, and restore.

42029

```
ERROR: proxy: failed to copy with SRDF, source=disk, target=disk, class=class
```

Explanation

Copying failed due to an I/O error caused while conducting copying between the proxy group *proxy* and the master group with SRDF. "source" is the name of the copy source disk, "target" is the name of the copy destination disk, and "class" is the name of the class to which the copy source and destination disks belong.

Resolution

Identify the cause by referring to the log messages for GDS, disk driver and Symmetrix device that were output right before the occurrence of the error, and restore.

42030

ERROR: proxy: failed to start OPC, source=disk.volume, target=disk.volume, class=class

Explanation

An error occurred between proxy volume *proxy* and the master volume when OPC was initiated. "source" indicates the copy source slice name, "target" indicates the copy destination slice name, and "class" indicates a class to which the copy source and destination slices belong.

Resolution

Identify the cause based on log messages of GDS, disk drivers, Fujitsu Storage Systems ETERNUS and such, and promptly recover the disk.

42031

ERROR: proxy: failed to stop OPC, source=disk.volume, target=disk.volume, class=class

Explanation

When stopping OPC between proxy volume *proxy* and master volume, an error occurred. "source" specifies the original slice name, "target" the mirror slice name, and "class" to which the original slice and mirror slice belong.

Resolution

Identify the cause by referring to the log messages for GDS, disk driver, Fujitsu Storage Systems ETERNUS that were output right before the occurrence of the error, and restore.

42032

ERROR: proxy: failed to start EC session, source=disk.volume, target=disk.volume, class=class

Explanation

An error occurred between proxy volume *proxy* and the master volume when EC was initiated. "source" indicates the copy source slice name, "target" indicates the copy destination slice name, and "class" indicates a class to which the copy source and destination slices belong.

Resolution

Identify the cause based on log messages of GDS, disk drivers, Fujitsu Storage Systems ETERNUS and such, and promptly recover the disk.

42033

ERROR: proxy: failed to stop EC session, source=disk.volume, target=disk.volume, class=class

Explanation

When stopping EC session between proxy volume *proxy* and master volume, an error occurred.

"source" specifies the original slice name, "target" the mirror slice name, and "class" to which the original slice and mirror slice belong.

Resolution

Identify the cause by referring to the log messages for GDS, disk driver, and Fujitsu Storage Systems ETERNUS that was output right before the occurrence of the error, and restore.

42034

```
ERROR: proxy: failed to suspend EC session, source=disk.volume, target=disk.volume, class=class
```

Explanation

When temporarily suspending EC session between proxy volume *proxy* and master volume, an error occurred.

"source" specifies the original slice name, "target" the mirror slice name, and "class" to which the original slice and mirror slice belong.

Resolution

Identify the cause by referring to the log messages for GDS, disk driver, and Fujitsu Storage Systems ETERNUS that were output right before the occurrence of the error, and restore.

42035

```
ERROR: proxy: failed to resume EC session, source=disk.volume, target=disk.volume, class=class
```

Explanation

When resuming EC session between proxy volume *proxy* and master volume, an error occurred.

"source" specifies the original slice name, "target" the mirror slice name, and "class" to which the original slice and mirror slice belong.

Resolution

Identify the cause by referring to the log messages for GDS, disk driver, and Fujitsu Storage Systems ETERNUS that were output right before the occurrence of the error, and restore.

42036

```
ERROR: proxy: failed to establish BCV pair, STD=disk, BCV=disk, class=class
```

Explanation

Error occurred while establishing a BCV pair between the proxy group *proxy* and the master group.

"STD" is the disk name of the standard device, "BCV" is the disk name of the BCV device, and "class" is the name of the class to which the standard and BCV devices belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or Symmetrix device, and promptly recover the disk.

42037

```
ERROR: proxy: failed to cancel BCV pair, STD=disk, BCV=disk, class=class
```

Explanation

Error occurred while canceling a BCV pair between the proxy group *proxy* and the master group.

"STD" is the disk name of the standard device, "BCV" is the disk name of the BCV device, and "class" is the name of the class to which the standard and BCV devices belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or Symmetrix device, and promptly recover the disk.

42038

ERROR: proxy: failed to split BCV pair, STD=disk, BCV=disk, class=class

Explanation

Error occurred while splitting a BCV pair between the proxy group *proxy* and the master group.

"STD" is the disk name of the standard device, "BCV" is the disk name of the BCV device, and "class" is the name of the class to which the standard and BCV devices belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or Symmetrix device, and promptly recover the disk.

42039

ERROR: proxy: failed to re-establish BCV pair, STD=disk, BCV=disk, class=class

Explanation

Error occurred while re-establishing a BCV pair between the proxy group *proxy* and the master group.

"STD" is the disk name of the standard device, "BCV" is the disk name of the BCV device, and "class" is the name of the class to which the standard and BCV devices belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or Symmetrix device, and promptly recover the disk.

42040

ERROR: proxy: failed to establish SRDF pair, source=disk, target=disk, class=class

Explanation

Error occurred while establishing an SRDF pair between the proxy group *proxy* and the master group.

"source" is the name of the source disk, "target" is the name of the target disk, and "class" is the name of the class to which the source and target disks belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or Symmetrix device, and promptly recover the disk.

42041

ERROR: proxy: failed to cancel SRDF pair, source=disk, target=disk, class=class

Explanation

Error occurred while canceling an SRDF pair between the proxy group *proxy* and the master group.

"source" is the name of the source disk, "target" is the name of the target disk, and "class" is the name of the class to which the source and target disks belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or Symmetrix device, and promptly recover the disk.

42042

ERROR: proxy: failed to split SRDF pair, source=disk, target=disk, class=class

Explanation

Error occurred while splitting an SRDF pair between the proxy group *proxy* and the master group.

"source" is the name of the source disk, "target" is the name of the target disk, and "class" is the name of the class to which the source and target disks belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or Symmetrix device, and promptly recover the disk.

42043

ERROR: proxy: failed to re-establish SRDF pair, source=disk, target=disk, class=class

Explanation

Error occurred while re-establishing an SRDF pair between the proxy group *proxy* and the master group.

"source" is the name of the source disk, "target" is the name of the target disk, and "class" is the name of the class to which the source and target disks belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or Symmetrix device, and promptly recover the disk.

42045

ERROR: volume: failed to change slice status, class=class

Explanation

Changing the status of slices composing the volume *volume* of the class *class* failed. If ACM's inter-server replication function has normally started the EC or REC copying to *volume*, the equivalency of *volume* is not ensured.

Resolution

Collect investigation material and contact your local customer support.

42046

ERROR: volume: failed to change read policy, class=class

Explanation

The volume *volume* of the class *class* is registered as the replication destination volume of ACM's inter-server replication function. The read policy of *volume* is set as follows according to the inter-server replication processing state.

- a. When *volume* is up and it is the EC/REC copy destination in equivalent status:

A slice that is read when a read request is issued from a higher level to *volume* is always a certain slice among ACTIVE slices composing *volume*.

- b. When *volume* is in a status other than a.:

A slice that is read when a read request is issued from a higher level to *volume* is selected from ACTIVE slices composing *volume* one after the other.

If this message is output when changing the read policy to the status of a., reading *volume* may read improper data.

If this message is output when changing the read policy to the status of b., the load of reading certain slices composing *volume* may be higher.

Resolution

Collect investigation material and contact your local customer support.

42047

ERROR: no enough address space, details

Explanation

The OS malloc function returned an error.

Resolution

Check whether memory and a swap area are sufficient.

42048

ERROR: volume: invalid volume type, class=class

Explanation

The volume *volume* of the class *class* is not of a volume type supported by ACM's inter-server replication function.

Resolution

Specify a volume of a supported type.

If *volume* is already registered in ACM replication volume information, the *volume* configuration may have been changed after the registration. In that event, with the relevant ACM commands, cancel the EC or REC copying to *volume* and remove *volume* from the ACM replication volume information.

To re-register *volume* in the ACM replication volume information, back up *volume* data as needed and then change the *volume* configuration.

E.3.3 Warning Messages (44000-44099)

44000

WARNING: read error on *status slice object.volume*, class=*class*:

volume info:devno(maj,min)=v_devno(v_maj,v_min)

device info:devno(maj,min)=p_devno(p_maj,p_min)

devname=*device*

error info: blk in disk=*blknodk*, blk in slice=*blknosl*,

length=*length*, resid=*resid*, errno=*errno*

Explanation

Read request sent to slice *object.volume* in *status* status terminated abnormally. Slice with abnormality will be detached.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44001

WARNING: write error on *status slice object.volume*, class=*class*:

volume info:devno(maj,min)=v_devno(v_maj,v_min)

device info:devno(maj,min)=p_devno(p_maj,p_min)

devname=*device*

error info: blk in disk=*blknodk*, blk in slice=*blknosl*,

length=*length*, resid=*resid*, errno=*errno*

Explanation

Write request sent to slice *object.volume* in *status* status terminated abnormally. Slice with abnormality will be detached.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44002

WARNING: open error on *status slice object.volume*, class=*class*:

volume info:devno(maj,min)=v_devno(v_maj,v_min)

device info:devno(maj,min)=p_devno(p_maj,p_min)

devname=*device*

error info: oflag=*oflag*, errno=*errno*

Explanation

Open request sent to slice *object.volume* in *status* status terminated abnormally. Slice with abnormality will be detached.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44003

WARNING: read error and writeback error on *status slice object.volume*, *class=class*:

```
volume info:devno(maj,min)=v_devno(v_maj,v_min)

device info:devno(maj,min)=p_devno(p_maj,p_min)

        devname=device

error info: blk in disk=blknodk, blk in slice=blknos1,

        length=length, resid=resid, errno=errno
```

Explanation

Read request and writeback request sent to slice *object.volume* in *status* status terminated abnormally. Slice with abnormality will be detached.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44004

WARNING: NVURM write error on disk *disk*, *class=class*:

```
volume info:devno(maj,min)=v_devno(v_maj,v_min)

        volume=volume, class=class

device info:devno(maj,min)=p_devno(p_maj,p_min)

        devname=device

error info: blk in disk=blknodk, blk in slice=blknos1,

        length=length, resid=resid, errno=errno
```

Explanation

NVURM write request sent to disk *disk* terminated abnormally. Although just resynchronization process on *volume* will be temporarily disabled, it will automatically attempt recovery.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44005

WARNING: *object.volume*: detached *status slice* by an I/O error, *class=class*

Explanation

Since an I/O error occurred on slice *object.volume* in *status* status, the slice was detached from the volume.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44006

```
WARNING: open error on private slice pslice, oflag=oflag, errno=errno
```

Explanation

Open request sent to disk private slice *pslice* terminated abnormally. It will automatically search for a normal alternate disk and attempt recovery.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44007

```
WARNING: read error on private slice p_devno(p_maj,p_min), offset=blknosl, length=length,  
resid=resid, errno=errno
```

Explanation

Read request sent to disk private slice *p_devno(p_maj,p_min)* terminated abnormally. It will automatically search for a normal alternate disk and attempt recovery.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44008

```
WARNING: write error on private slice p_devno(p_maj,p_min), offset=blknosl, length=length,  
resid=resid, errno=errno
```

Explanation

Write request sent to disk private slice *p_devno(p_maj,p_min)* terminated abnormally. It will automatically search for a normal alternate disk and attempt recovery.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44009

```
WARNING: close error on private slice p_devno(p_maj,p_min), errno=errno
```

Explanation

Close request sent to disk private slice *p_devno(p_maj,p_min)* terminated abnormally. It will automatically search for a normal alternate disk and attempt recovery.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44010

WARNING: *sdxfunc: pslice: open error, errno=errno*

Explanation

The open request sent to the physical slice *pslice* terminated abnormally.

The following messages may be output when a node is booted, investigation material is collected (fjsnap or sdxsnap.sh is executed), or the physical disk information update menu in GDS Management View is executed.

a) WARNING: pd_get_info: *pslice: open error, errno=6*

b) WARNING: pd_set_orig_all: *pslice: open error, errno=6*

Resolution

A disk failure may have occurred. Identify the cause by referring to disk driver log messages, and recover the disk.

However, in the following situations, GDS is behaving normally and the messages a) and b) may be ignored.

- In the messages a) and b), *pslice* is a physical slice of a disk unit previously removed. In this situation, delete the device special file for *pslice*, and these messages will no longer be output.
 - In the messages a) and b), *pslice* is a physical disk slice of a physical disk not registered with GDS.
-

44011

WARNING: *sdxfunc: pslice: read error, errno=errno*

Explanation

Read request sent to physical slice *pslice* terminated abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44012

WARNING: *sdxfunc: pslice: write error, errno=errno*

Explanation

Write request sent to physical slice *pslice* terminated abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44013

WARNING: *sdxfunc: pslice: ioctl error, request=ioctlcmd, errno=errno*

Explanation

The ioctl request sent to physical slice *pslice* terminated abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44014

WARNING: sdxfunc: pslice: close error, errno=errno

Explanation

Close request sent to physical slice *pslice* terminated abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44015

WARNING: volume: failed to enable JRM, no available NVURM space, class=class

Explanation

Resuming just resynchronization process on *volume* was unsuccessful due to insufficient number of normal disks.

Resolution

Check the disk status within *class*. Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44016

WARNING: volume: failed to retrieve NVURM from disk disk, class=class

Explanation

NVURM read request of *volume* from *disk* was unsuccessful. Just resynchronization will switch to copying of the entire disk.

NVURM is volume update area map information stored in the disk for just resynchronization.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44017

WARNING: disk: device: disabled disk, class=class

Explanation

disk is disabled since its data is invalid.

device is the physical disk name corresponding to *disk*.

Resolution

Disk configuration may have been wrongly changed or disk data could be damaged. Check for improper connection change of I/O cables and disk swap. Also check for disk driver log message regarding the disk in question, and attempt recovery.

44018

WARNING: volume: volume synchronization failed, class=class

Explanation

Synchronization copying of *volume* was unsuccessful.

Resolution

Attempt recovery by identifying the cause of failure by referring to GDS log message and disk driver message that were output immediately before the error.

44019

WARNING: volume: volume just resynchronization failed, class=class

Explanation

Just resynchronization of *volume* was unsuccessful.

Resolution

Attempt recovery by identifying the cause of failure by referring to GDS log message and disk driver message that were output immediately before the error.

44020

WARNING: class: unknown class file found

Explanation

Class file *class* which does not exist in class database file was found while booting the system.

44021

WARNING: invalid configuration database ID information, sdxerrno=sdxerrno, class=class:

psdevtree

Explanation

Since the configuration database ID information was invalid, private slice *psdevtree* was not used as the configuration database for class *class*.

Resolution

Disk configuration may have been wrongly changed or disk data could be damaged. Check for improper connection change of I/O cables and disk swap. Also check for disk driver log message regarding the disk in question, and attempt recovery.

44022

WARNING: class: too few valid configuration database replicas

Explanation

There are not sufficient valid configuration databases for *class*.
This message is displayed when the majority of disks registered with *class* are unavailable.
Leaving it as is may cause serious problems.

Resolution

For details, see "[F.1.4 Class Status Abnormality](#)."

44023

WARNING: cannot open message logging file, errno=errno

`/var/opt/FJSVsdx/msglog/daemon.log`

Explanation

Opening GDS log file was unsuccessful.
This message will be output via syslog. Although the message will not be output on GDS logfile, it does not affect other processes.

Resolution

Collect investigation material, and contact your local customer support regarding recovery.

44024

WARNING: cannot write message logging file, errno=errno

`/var/opt/FJSVsdx/msglog/sdxservd.log`

Explanation

Writing to GDS log file was unsuccessful.
This message will be output via syslog. Although the message will not be output on GDS logfile, it does not affect other processes.

Resolution

Collect investigation material and contact your local customer support.

44025

WARNING: failed to reply message to node *node*, details

Explanation

Replying to remote node *node* was unsuccessful.

Resolution

Collect investigation material and contact your local customer support.

44026

WARNING: class: failed to change class resource status on remote node node, status=new-status, sdxerrno=sdxerrno

Explanation

Changing class resource status to *new-status* on a remote node *node* was unsuccessful.

Resolution

Collect investigation material and contact your local customer support.

44027

WARNING: sdxfunc: no enough address space, osfunc=osfunc, errno=errno

Explanation

The OS *osfunc* function returned an error. The error number is *errno*.

Resolution

Confirm you have sufficient memory or swap area.

44028

WARNING: respawned daemon daemon successfully

Explanation

Although *daemon* terminated abnormally, it completed normally by restarting. This message will be output via syslog.

Resolution

After *daemon* ended abnormally, the recovery function restarted the daemon normally. Therefore, the abnormal end has no influence on the operation and restoration work is not necessary.

If you want to identify the cause of the *daemon* abnormal end, collect investigation material and contact your local customer support.

44029

WARNING: device: failed to restore VTOC on disk, sdxerrno=sdxerrno

Explanation

Recovering physical disk *device* format information was unsuccessful. Disk failure may have occurred.

Resolution

Use `format(1M)` command to recover format information.

44030

WARNING: cannot set use-nvramrc? parameter to true, NVRAMRC not empty

Explanation

Since a command is already registered with NVRAMRC, use-nvramrc? parameter could not be set to "true."

Resolution

For details of advisory notices, see "[A.2.9 System Disk Mirroring](#)."

44031

WARNING: volume: cannot copy to one or more slices in volume

Explanation

Copying process could not be performed on some slices within *volume*.

Resolution

Execute synchronization copying as needed.

44032

WARNING: device: write error, errno=errno

Explanation

A write error occurred in physical disk *device* and recovering the VTOC of *device* failed. A disk failure may have occurred.

Resolution

To recover the VTOC of *device*, use the format(1M) command. If *device* is write-locked by the disk unit's copy function, no action is required.

44033

WARNING: device: read error, errno=errno

Explanation

A read error occurred in physical disk *device* and recovering the VTOC of *device* failed. A disk failure may have occurred.

Resolution

To recover the VTOC of *device*, use the format(1M) command.

44034

WARNING: error may occur on device:

special

Explanation

An error may have occurred on the mirrored boot device *special*.

Resolution

Disk failure may have occurred. Identify the cause by referring to the disk driver log message and such, and recover the disk.

44035

WARNING: system booted from diagnostic boot device

Explanation

The diag-switch? parameter stored in NVRAM (non-volatile memory) on the main unit is true, and the system was booted from the diagnostic boot device.

Resolution

Change the diag-switch? variable to false where necessary and reboot the system.

44036

WARNING: proxy: too many EC/OPC sessions

Explanation

The number of EC or OPC sessions within the physical disk (LU) or the disk array body has reached the upper limit of supported concurrent sessions. For this reason, a new EC or OPC session cannot be started. Copying is started using the soft copy function.

Resolution

To make copying by EC or OPC available, cancel the relationship between the proxy volume *proxy* and the master volume, wait until the running session ends and try this command again. Alternatively, according to need, cancel the running session using the `sdxproxy Cancel` command, the `sdxproxy Break` command, or `[Operation]:[Proxy Operation]:[Break]` and try this command again.

44037

WARNING: proxy: failed to start OPC, source=disk.volume, target=disk.volume, class=class

Explanation

An error occurred when OPC was started between the proxy volume *proxy* and the master volume.

"source" is the copy source slice name, "target" is the copy destination slice name, and "class" is the name of the class to which the copy source slice and the copy destination slice belong.

Copying is started between *proxy* and the master using the soft copy function.

Resolution

The cause of this error may derive from a faulty disk array unit. Identify the cause by referring to the log messages for GDS, disk driver, and Fujitsu Storage Systems ETERNUS that were output right before the occurrence of the error, and restore.

44038

WARNING: proxy: failed to start EC session, source=disk.volume, target=disk.volume, class=class

Explanation

An error occurred when an EC session was started between the proxy volume *proxy* and the master volume.

"source" is the copy source slice name, "target" is the copy destination slice name, and "class" is the name of the class to which the copy source slice and the copy destination slice belong.

Copying is started between *proxy* and the master using the soft copy function.

Resolution

The cause of this error may derive from a faulty disk array unit. Identify the cause by referring to the log messages for GDS, disk driver, and Fujitsu Storage Systems ETERNUS that were output right before the occurrence of the error, and restore.

E.3.4 Information Messages (46000-46199)

46000

```
INFO: read error and writeback success on status slice object.volume, class=class:
```

```
volume info:devno(maj,min)=v_devno(v_maj,v_min)
```

```
device info:devno(maj,min)=p_devno(p_maj,p_min)
```

```
devname=device,
```

```
error info: blk in disk=blknodk, blk in slice=blknosl,
```

```
length=length
```

Explanation

Read request sent to slice *object.volume* in status *status* was unsuccessful, but was restored by writeback process. Writeback is a process to read data from other slices in the event of a read error.

46001

```
INFO: NVRAM configuration parameter has been updated:
```

```
parameter='param'
```

```
old      ='old_value'
```

```
new      ='new_value'
```

Explanation

Parameter *param* value stored in NVRAM (non-volatile memory) on the body unit was updated from *old_value* to *new_value*.

46002

```
INFO:volume: temporarily disabled JRM, class=class
```

Explanation

Although just resynchronization process on *volume* will be temporarily disabled, it will automatically attempt recovery.

46003

```
INFO: disk: failed to connect spare disk for disk disk, group=group, class=class, sdxerrno=sdxerrno
```

Explanation

Connecting spare disk *disk* to *group* instead of *disk* was unsuccessful.

46004

INFO: *disk*: failed to connect spare disk for group *lgroup*, the highest level group=*hgroup*, class=*class*, sdxerrno=*sdxerrno*

Explanation

The attempt to connect spare disk *disk* to the highest level group *hgroup* in place of group *lgroup* failed.

46005

INFO: *disk*: connected spare disk for disk *disk*, group=*group*, class=*class*

Explanation

Connected spare disk *disk* to *group* instead of *disk*.

46006

INFO: *disk*: connected spare disk for group *lgroup*, the highest level group=*hgroup*, class=*class*

Explanation

Spare disk *disk* was connected to the highest level group *hgroup* in place of group *lgroup*.

46007

INFO: *group*: free blocks are reduced in group, class=*class*

Explanation

Free blocks of *group* were reduced.

46008

INFO: *volume*: reallocated NVURM space and enabled JRM successfully, class=*class*

Explanation

Just resynchronization process of *volume* was resumed.

46009

INFO: *volume*: retrieved NVURM from disk *disk* successfully, class=*class*

Explanation

NVURM read request of *volume* sent from *disk* was successful. Just resynchronization process will be resumed. NVURM is volume update area map information stored in the disk for just resynchronization.

46010

INFO: *volume*: no need to retrieve NVURM, *sdxerrno*=*sdxerrno*, *class*=*class*

Explanation

NVURM was not retrieved. Entire copying of volume will start.

46011

INFO: *disk*: *pslice*: failed to open physical special files exclusively, *errno*=*errno*

Explanation

Opening *disk* physical slice *pslice* exclusively was unsuccessful.

46012

INFO: *disk*: *device*: disk ID information is invalid, *sdxerrno*=*sdxerrno*

Explanation

Disk ID information of *disk* is invalid. *disk* will be automatically disabled.

46013

INFO: *disk*: enabled disk, *class*=*class*

Explanation

Disabled *disk* was enabled.

46014

INFO: *volume*: volume synchronization started, *class*=*class*

Explanation

Started synchronization copying process of *volume*.

46015

INFO: *volume*: volume just resynchronization started, *class*=*class*

Explanation

Started just resynchronization process of *volume*.

46016

INFO: *volume*: volume synchronization canceled, class=*class*

Explanation

Canceled synchronization copying of *volume*.

46017

INFO: *volume*: volume just resynchronization canceled, class=*class*

Explanation

Canceled just resynchronization process of *volume*.

46018

INFO: *volume*: volume synchronization completed successfully, class=*class*

Explanation

Synchronization copying process of *volume* completed successfully.

46019

INFO: *volume*: volume just resynchronization completed successfully, class=*class*

Explanation

Just resynchronization process of *volume* completed successfully.

46020

INFO: *object*: failed to update configuration database, class=*class*

Explanation

Updating configuration database was unsuccessful since all (or majority of) disks registered with *class* were unavailable. Usually, *class* will be closed following this message.

46021

INFO: *sdxservd* daemon started up

Explanation

The *sdxservd* daemon started up. GDS will be started.

46022

INFO: started local volumes, *timeofday*

Explanation

All volumes under local class have been started.

46023

INFO: started root volumes

Explanation

All volumes under root class have been started.

46024

INFO: stopped all services by shutdown request, *timeofday*

Explanation

All GDS services were stopped in response to a shutdown request.

46025

INFO: cannot open class database file, *errno=errno*

Explanation

Could not open class database file while booting the system. Will automatically attempt recovery.

46026

INFO: class database file corrupted

Explanation

Corrupted class database file was detected while booting the system. Will automatically attempt recovery.

46027

INFO: *class*: cannot open class file, *errno=errno*

Explanation

Could not open class file *class* while booting the system. Will automatically attempt recovery.

46028

INFO: *class*: class file corrupted

Explanation

Corrupted class file *class* was detected while booting the system. Will automatically attempt recovery.

46029

INFO: class database file updated successfully

Explanation

Class database file was updated.

46030

INFO: class: class file updated successfully

Explanation

Class file *class* was updated.

46031

INFO: cannot write class database file, errno=errno

Explanation

Could not write to class database file. Will automatically attempt recovery.

46032

INFO: class: cannot write class file, errno=errno

Explanation

Could not write to class file *class*. Will automatically attempt recovery.

46033

INFO: cannot check configuration database ID information, sdxerrno=sdxerrno, class=class:
psdextree

Explanation

Opening or reading private slice *psdextree* was unsuccessful, and could not check configuration database ID information of *class*.

46034

INFO: cannot check configuration database, sdxerrno=sdxerrno, class=class:
psdextree

Explanation

Opening or reading private slice *psdextree* was unsuccessful, and could not check configuration database of *class*.

46035

INFO: configuration database corrupted, *sdxerrno=sdxerrno*, *class=class*:

psdextree ...

Explanation

Since there was a check-sum error in the configuration database, private slice *psdextree...* was not used as the configuration database for class *class*.

46036

INFO: configuration database defeated, *sdxerrno=sdxerrno*, *class=class*:

psdextree ...

Explanation

The configuration database for class *class* stored on private slice *psdextree...* was determined to be invalid in a validity check.

46037

INFO: *class*: valid configuration database replicas exist on:

psdextree ...

Explanation

The valid configuration database for class *class* was determined.
psdextree... is the private slice storing the valid configuration database.

46038

INFO: *class*: starting up *class*

Explanation

class will be started.

46039

INFO: cannot update configuration database replica, *sdxerrno=sdxerrno*, *class=class*:

psdextree

Explanation

Updating a replica of the configuration database for class *class* stored on private slice *psdextree* failed.

46040

INFO: *class*: relocated configuration database replicas on:
psdextree ...

Explanation

A replica of the configuration database for class *class* was relocated onto private slice *psdevtree...*

46041

INFO: *disk*: disconnected spare disk from group *group*, class=*class*

Explanation

Disconnected spare disk *disk* from *group*.

46042

INFO: *group*: free blocks are increased in group, class=*class*

Explanation

Number of free blocks in *group* increased.

46043

INFO: failed to create a new thread, errno=*errno*

Explanation

Function `pthread_create()` terminated abnormally.

46044

INFO: cannot open configuration parameter file *filename*, errno=*errno*

Explanation

Opening configuration parameter *filename* was unsuccessful.

46045

INFO: cannot read configuration parameter file, errno=*errno*

Explanation

Reading configuration parameter files was unsuccessful.

46046

INFO: received unexpected data from sfdsk driver and ignored

Explanation

Unexpected data was received from sfdsk driver and was ignored.

46047

INFO: received unexpected event from sfdsk driver and ignored, *details*

Explanation

Unexpected event was received from sfdsk driver and was ignored.
details displays the details about the event.

46048

INFO: *class*: class closed down, *node=node*

Explanation

class on *node* was closed.

46049

INFO: command executed:

cmdline

Explanation

Command *cmdline* was executed.

46050

INFO: command exited, *exit-status=exitstat*:

cmdline

Explanation

Processing *cmdline* is complete.

46051

INFO: trying to execute command:

cmdline

Explanation

The *cmdline* command is about to be executed.

46052

INFO: failed to execute command:

cmdline

Explanation

The *cmdline* command failed.

46053

INFO: *class*: changed class resource status on remote node *node*, *old-status* => *new-status*

Explanation

Class resource status on remote node *node* was changed from *old-status* to *new-status*.

46054

INFO: *class*: changed class resource status on current node *node*, *old-status* => *new-status*

Explanation

Class resource status on current node *node* was changed from *old-status* to *new-status*.

46055

INFO: *class*: started *type* volumes, node=*node*

Explanation

Starting volumes that belong to the class *class* was completed on the node *node*.
type is the *class* type attribute.

46056

INFO: *class*: started and stood by *type* volumes, node=*node*

Explanation

Starting and putting on standby volumes that belong to the class *class* was completed on the node *node*.
type is the *class* type attribute.

46057

INFO: *class*: stopped and stood by *type* volumes, node=*node*

Explanation

Stopping and putting on standby volumes that belong to the class *class* was completed on the node *node*.
type is the *class* type attribute.

46058

INFO: *class*: stopped *type* volumes, node=*node*

Explanation

Stopping volumes that belong to the class *class* was completed on the node *node*.
type is the *class* type attribute.

46059

INFO: cannot connect spare disk, cluster-wide lock failure, class=*class*, sdxerrno=*sdxerrno*

Explanation

Due to the occurrence of an abnormal exclusive control between cluster system nodes, the spare disk could not be connected.

46060

INFO: cannot connect spare disk, too few valid configuration database replicas, class=*class*,
disk=*disk*

Explanation

Could not connect spare disk *disk* due to insufficient number of valid configuration databases.

46061

INFO: cannot connect spare disk, hot spare disabled, class=*class*, disk=*disk*

Explanation

Spare disk *disk* could not be connected since the hot spare is disabled.

46062

INFO: *class*: started class-down recovery for remote node *node*

Explanation

Closed class *class* on remote node *node* will be recovered.

46063

INFO: *class*: class-down recovery failed, already class-down on current node *node*

Explanation

Attempted recovery of closed class *class* on remote node *node*. Recovery was unsuccessful since the class was also closed on current node *node*.

46064

INFO: *class*: class-down recovery failed, *sdxerrno*=*sdxerrno*

Explanation

Recovering closed class *class* was unsuccessful.

46065

INFO: *class*: class-down recovery completed successfully

Explanation

Recovering closed class *class* was successful.

46066

INFO: *class*: started node-down recovery for remote node *node*

Explanation

Started node-down recovery on remote node *node*.

46067

INFO: *class*: started shutdown recovery for remote node *node*

Explanation

Started shutdown recovery on remote node *node*.

46068

INFO: *class*: node-down recovery failed, already class-down on current node *node*

Explanation

Recovering node-down was unsuccessful since *class* was in closed status on current node *node*.

46069

INFO: *class*: shutdown recovery failed, already class-down on current node *node*

Explanation

Recovering shutdown was unsuccessful since *class* was in closed status on current node *node*.

46070

INFO: *class*: node-down recovery failed, *sdxerrno*=*sdxerrno*

Explanation

Recovering *class* from node-down was unsuccessful.

46071

INFO: *class*: shutdown recovery failed, *sdxerrno*=*sdxerrno*

Explanation

Recovering *class* from shutdown was unsuccessful.

46072

INFO: *class*: node-down recovery completed successfully

Explanation

Recovering *class* from node-down completed successfully.

46073

INFO: *class*: shutdown recovery completed successfully

Explanation

Recovering *class* from shutdown completed successfully.

46074

INFO: *object.volume*: failed to update slice error information, *class* closed down, *class*=*class*

Explanation

Updating error information of slice *object.volume* was unsuccessful due to *class* in closed status.

46075

INFO: *volume*: failed to disable JRM, *class* closed down, *class*=*class*

Explanation

Disabling just resynchronization process on *volume* was unsuccessful due to *class* in closed status.

46076

INFO: *object.volume*: failed to detach slice, class closed down, class=*class*

Explanation

Detaching slice *object.volume* was unsuccessful due to *class* in closed status.

46077

INFO: *volume*: failed to restart volume, class closed down, class=*class*

Explanation

Restarting *volume* was unsuccessful due to *class* in closed status.

46078

INFO: open error on *status slice object.volume*, class=*class*

Explanation

Open request sent to slice *object.volume* in *status* status terminated abnormally.

46079

INFO: *class*: trying to identify class master, *details*

Explanation

Trying to identify the class master for shared class *class*.

46080

INFO: *class*: identified class master, node=*node*

Explanation

The master for shared class *class* has been identified as node *node*.

46081

INFO: *class*: searching class master

Explanation

Searching the master for shared class *class*.

46082

INFO: *class*: class master found, node=*node*

Explanation

The master for shared class *class* was found to be node *node*.

46083

INFO: *class*: class master not found

Explanation

The master for shared class *class* cannot be found.

46084

INFO: *class*: got class master privilege

Explanation

Trying to obtain master privilege for shared class *class*.

46085

INFO: *class*: broadcasted class master information to remote nodes

Explanation

The master information for shared class *class* has been broadcast to the remote nodes.

46086

INFO: *class*: received confirmations of class master information from remote node *node*

Explanation

Received master confirmation for shared class *class* from remote node *node*.

46087

INFO: waiting for outstanding event operations, *details*

Explanation

Waiting for event operations in process.

46088

INFO: completed outstanding event operations

Explanation

Event operations in process have been completed.

46089

INFO: *class*: trying to release class master privilege, *details*

Explanation

Trying to release master privilege for shared class *class*.

46090

INFO: *class*: released class master privilege

Explanation

Master privilege for shared class *class* has been released.

46091

INFO: *proxy*: started to copy with OPC, *source=disk.volume*, *target=disk.volume*, *class=class*

Explanation

Copying between proxy volume *proxy* and master volume with the OPC function started.

"source" specifies the original slice name, "target" the mirror slice name, and "class" the name of the class to which the original slice and mirror slice belong.

46092

INFO: *proxy*: completed copying with OPC, *source=disk.volume*, *target=disk.volume*, *class=class*

Explanation

Copying between proxy volume *proxy* and master volume with the OPC function is completed.

"source" specifies the original slice name, "target" the mirror slice name, and "class" the name of the class to which the original slice and mirror slice belong.

46093

INFO: *proxy*: canceled copying with OPC, *source=disk.volume*, *target=disk.volume*, *class=class*

Explanation

Copying between proxy volume *proxy* and master volume with the OPC function was cancelled.

"source" specifies the original slice name, "target" the mirror slice name, and "class" the name of the class to which the original slice and mirror slice belong.

46094

INFO: *proxy*: EC session started, *source=disk.volume*, *target=disk.volume*, *class=class*

Explanation

EC session between proxy volume *proxy* and master volume started.

"source" specifies the original slice name, "target" the mirror slice name, and "class" the name of the class to which the original slice and mirror slice belong.

46095

INFO: proxy: completed copying with EC, source=disk.volume, target=disk.volume, class=class

Explanation

Copying between proxy volume *proxy* and master volume with the EC function is completed.

"source" specifies the original slice name, "target" the mirror slice name, and "class" the name of the class to which the original slice and mirror slice belong.

46096

INFO: proxy: canceled copying with EC, source=disk.volume, target=disk.volume, class=class

Explanation

Copying between proxy volume *proxy* and master volume with the EC function was cancelled.

"source" specifies the original slice name, "target" the mirror slice name, and "class" the name of the class to which the original slice and mirror slice belong.

46097

INFO: proxy: EC session stopped, source=disk.volume, target=disk.volume, class=class

Explanation

EC session between proxy volume *proxy* and master volume with the EC function was stopped.

"source" specifies the original slice name, "target" the mirror slice name, and "class" the name of the class to which the original slice and mirror slice belong.

46098

INFO: proxy: EC session suspended, source=disk.volume, target=disk.volume, class=class

Explanation

EC session between proxy volume *proxy* and master volume with the EC function has been temporarily suspended.

"source" specifies the original slice name, "target" the mirror slice name, and "class" the name of the class to which the original slice and mirror slice belong.

46099

INFO: proxy: EC session resumed, source=disk.volume, target=disk.volume, class=class

Explanation

EC session between proxy volume *proxy* and master volume with the EC function was resumed.

"source" specifies the original slice name, "target" the mirror slice name, and "class" the name of the class to which the original slice and mirror slice belong.

46100

INFO: *proxy*: established BCV pair, STD=*disk*, BCV=*disk*, class=*class*

Explanation

BCV pair between the proxy group *proxy* and the master group was established.

"STD" is the disk name of the standard device, "BCV" is the disk name of the BCV device, and "class" is the name of the class to which the standard and BCV devices belong.

46101

INFO: *proxy*: completed copying with TimeFinder, source=*disk*, target=*disk*, class=*class*

Explanation

Copying between the proxy group *proxy* and the master group with TimeFinder is complete.

"source" is the name of the copy source disk, "target" is the name of the copy destination disk, and "class" is the name of the class to which the copy source and destination disks belong.

46102

INFO: *proxy*: canceled copying with TimeFinder, source=*disk*, target=*disk*, class=*class*

Explanation

Copying between the proxy group *proxy* and the master group with TimeFinder was canceled.

"source" is the name of the copy source disk, "target" is the name of the copy destination disk, and "class" is the name of the class to which the copy source and destination disks belong.

46103

INFO: *proxy*: canceled BCV pair, STD=*disk*, BCV=*disk*, class=*class*

Explanation

BCV pair between the proxy group *proxy* and the master group was canceled.

"STD" is the disk name of the standard device, "BCV" is the disk name of the BCV device, and "class" is the name of the class to which the standard and BCV devices belong.

46104

INFO: *proxy*: split BCV pair, STD=*disk*, BCV=*disk*, class=*class*

Explanation

BCV pair between the proxy group *proxy* and the master group was split.

"STD" is the disk name of the standard device, "BCV" is the disk name of the BCV device, and "class" is the name of the class to which the standard and BCV devices belong.

46105

INFO: proxy: re-established BCV pair, STD=disk, BCV=disk, class=class

Explanation

BCV pair between the proxy group *proxy* and the master group was re-established.

"STD" is the disk name of the standard device, "BCV" is the disk name of the BCV device, and "class" is the name of the class to which the standard and BCV devices belong.

46106

INFO: proxy: established SRDF pair, source=disk, target=disk, class=class

Explanation

SRDF pair between the proxy group *proxy* and the master group was established.

"source" is the name of the source disk, "target" is the name of the target disk, and "class" is the name of the class to which the source and target disks belong.

46107

INFO: proxy: completed copying with SRDF, source=disk, target=disk, class=class

Explanation

Copying between the proxy group *proxy* and the master group with SRDF is complete.

"source" is the name of the copy source disk, "target" is the name of the copy destination disk, and "class" is the name of the class to which the copy source and destination disks belong.

46108

INFO: proxy: canceled copying with SRDF, source=disk, target=disk, class=class

Explanation

Copying between the proxy group *proxy* and the master group with SRDF was canceled.

"source" is the name of the copy source disk, "target" is the name of the copy destination disk, and "class" is the name of the class to which the copy source and destination disks belong.

46109

INFO: proxy: canceled SRDF pair, source=disk, target=disk, class=class

Explanation

SRDF pair between the proxy group *proxy* and the master group was canceled.

"source" is the name of the source disk, "target" is the name of the target disk, and "class" is the name of the class to which the source and target disks belong.

46110

```
INFO: proxy: split SRDF pair, source=disk, target=disk, class=class
```

Explanation

SRDF pair between the proxy group *proxy* and the master group was split.

"source" is the name of the source disk, "target" is the name of the target disk, and "class" is the name of the class to which the source and target disks belong.

46111

```
INFO: proxy: re-established SRDF pair, source=disk, target=disk, class=class
```

Explanation

SRDF pair between the proxy group *proxy* and the master group was re-established.

"source" is the name of the source disk, "target" is the name of the target disk, and "class" is the name of the class to which the source and target disks belong.

46112

```
INFO: file not found
```

Explanation

The *file* file was not found.

E.3.5 Internal Error Messages (48000)

48000

```
daemon: severity: module: internal error, details
```

Explanation

An internal error occurred. *details* gives the cause of the error, *daemon* gives the name of the daemon program, *severity* gives the severity of the message, and *module* the module name (usually the internal function name) where the event was detected.

The following messages may be output when an I/O error occurs.

- a) sdxservd: ERROR: *module*: internal error, sdxfunc=dbrw_read_dc(), rval=-1, sdxerrno=4120
- b) sdxservd: ERROR: *module*: internal error, sdxfunc=dbrw_read_dc(), rval=-1, sdxerrno=4121

The message c) below may be output when a node is booted. This means the return value of the OS read(2) system call was 0 and the error number was 2. If the return value is 0, the event is not an error and the error number has no meaning, but it is recorded as a WARNING because the return value is not normal (a positive value).

- c) sdxservd: WARNING: *module*: internal error, osfunc=read, rval=0, errno=2

Resolution

In the event of a) or b), some other I/O error message will be output. Refer to the explanation of resolution for that message, and take necessary actions.

In the event of c), if no error message was output around the same period of time, GDS was behaving normally and the message c) may be ignored. If some other error message was output around the same period of time, refer to the explanation and resolution for that message, and take necessary actions.

In other events, collect investigation material and contact your local customer support.

E.4 Command Messages

Command messages will be sent to standard output or standard error output. Output format is as follows.

```
SDX:command: severity: message
```

The *command* gives the command name, *severity* the severity of the message, and *message* the message text body.

Command message will also be output to the following log file.

```
/var/opt/FJVSdx/msglog/sdxservd.log
```

Output format is as follows.

```
Mon Day HH:MM:SS SDX:daemon: severity: message
```

Mon gives the month the message was output, *Day* the date, *HH* the hour, *MM* the minute, *SS* the second, *daemon* the daemon program name, *severity* the severity of the message, and *message* the message text body.

Message Numbers

The message numbers used to identify messages described in E.4.1 to E.4.5 do not appear in messages actually output by GDS commands.

Variable Names

Italicized words in the messages are variable names, and the actual output will vary depending on the situation. The meaning and the format of the variable names used in the message explanation are described below.

Variable names	Descriptions
<i>class</i>	Class name
<i>disk</i>	Disk name
<i>group</i>	Group name
<i>lgroup</i>	Lower level group name
<i>hgroup</i>	Higher level group name
<i>volume</i>	Volume name
<i>disk.volume</i>	Slice name
<i>object.volume</i>	Slice name
<i>master</i>	Master object name
<i>proxy</i>	Proxy object name
<i>object</i>	Object name (or physical disk name)
<i>status</i>	Object status

Variable names	Descriptions
<i>device</i>	Physical disk name (cCtTdD, mphpI, mplbI, emcpowerN) where C is the controller number, T is the target ID, D is the disk number, I is the instance number, and N is the emcpower device number.
<i>pslice</i>	Physical slice name (cCtTdDs{0-7}, mphpIs{0-7}, mplbIs{0-7}, emcpowerN{a-h}) where C is the controller number, T is the target ID, D is the disk number, I is the instance number, N is the emcpower device number, {0-7} is the slice number with 1 digit from 0 to 7, and {a-h} is the slice number with one alphabet from a to h.
<i>node</i>	Node identifier or node name.
<i>attribute</i>	Attribute name
<i>value</i>	Attribute value
<i>param</i>	Parameter name
<i>val</i>	Parameter value
<i>size</i>	Number of blocks (in sectors), (decimal number)
<i>option</i>	Command options.
<i>usage</i>	Syntax when using command.
<i>letter</i>	Characters
<i>details</i>	Details
<i>errno</i>	System call error number (decimal)
<i>sdxerrno</i>	Internal error number defined by GDS (decimal)
<i>string</i>	Other character strings

Explanation

Messages output by command are shown below in the order of severity. There are four levels of severity.

Level of importance	Descriptions
ERROR	This message is output to indicate the cause when a command terminates abnormally.
WARNING	This message is output as a warning when an abnormal event is detected. It will not always result in an abnormal termination of the command.
INFO	This message is output to record the command operation.
TO FIX	This message is output to show you how to fix the command. Usually it outputs command syntax.

E.4.1 Error Messages (60000-60499)

60000

```
ERROR: connection timeout
```

Explanation

Connection failed due to no response from sdxservd daemon.

Resolution

Confirm that sdxservd daemon process was started normally.

60001

ERROR: not privileged

Explanation

Executed user is not a superuser.

Resolution

You need to have superuser authority for execution.

60002

ERROR: option: illegal option

Explanation

option is illegal.

Resolution

See the Fix Messages described later, or the "[Appendix D Command Reference](#)."

60003

ERROR: syntax error

Explanation

Executed command has a syntax error.

Resolution

See the Fix Messages described later, or the "[Appendix D Command Reference](#)."

60004

ERROR: string: name too long

Explanation

Object name, file name or node identifier specified with *string* is too long.

Resolution

Specify the proper name.

60005

ERROR: object: name contains invalid character "letter"

Explanation

Object name *object* contains an invalid character *letter*.

Resolution

You can only use alphanumeric characters, "-" (hyphen), and "_" (underscore) when specifying an object name.

60006

ERROR: object: name starting with "-" or "-" is invalid

Explanation

Object name *object* starting with "-" (hyphen) or "_" (underscore) is invalid.

Resolution

Specify an object name that starts with an alphanumeric character.

60007

ERROR: device: illegal physical disk name

Explanation

Physical disk name *device* is illegal.

Resolution

See "[Appendix D Command Reference](#)."

60008

ERROR: object: object names must be unique within a class

Explanation

Object names within a class must be unique.

Resolution

Specify a unique object name.

60009

ERROR: cannot be executed in non-global zone

Explanation

This command cannot be executed in a non-global zone.

Resolution

Execute the command in the global zone.

60010

ERROR: *module: environment error, details*

Explanation

The command cannot be executed due to an error in the environment. The error causes only a failure in command execution but has no effect on the system. Volumes can be accessed continuously.

module is the name of the module on which this event was detected (usually an internal variable name).

details indicates the error in detail.

Resolution

Collect investigation material and contact your local customer support.

60011

ERROR: *class is a shadow class*

Explanation

The class *class* is a shadow class. The attempted operation is not supported for shadow classes.

Resolution

See "[Appendix D Command Reference](#)," and specify a correct command name and class name.

60012

ERROR: *attribute: invalid attribute name*

Explanation

The same *attribute* has already been specified.

Resolution

See "[Appendix D Command Reference](#)."

60013

ERROR: *attribute: attribute name duplicated*

Explanation

The same attribute *attribute* has already been specified.

Resolution

You can only specify an attribute once.

60014

ERROR: *attribute=value: invalid attribute value*

Explanation

Attribute value *value* is invalid.

Resolution

See "[Appendix D Command Reference](#)."

If "pslice=off: invalid attribute value" is output, the possible cause is that the FJSVsdxl package of PRIMECLUSTER GDS may not be installed normally. If this is the cause, reinstall FJSVsdxl.

60015

ERROR: node: node name duplicated

Explanation

The same node identifier has already been specified.

Resolution

You can only specify a node identifier once.

60016

ERROR: too many nodes in scope

Explanation

There are too many nodes specified in scope.

Resolution

Up to 16 nodes can be specified in a scope.

60017

ERROR: class: cannot operate shared objects, cluster control facility not ready

Explanation

Cannot operate shared objects since cluster control is not operating.

Resolution

Start cluster control and try again.

60018

ERROR: node: remote node cannot be specified for local class

Explanation

A remote node was specified for a local class.

Resolution

Specify the node identifier of the current node for a local class.

60019

ERROR: node: remote node cannot be specified for root class

Explanation

A remote node was specified for a root class.

Resolution

Specify the node identifier of the current node for a local class.

60020

ERROR: current node must be specified as scope

Explanation

Current node is not specified for shared class scope.

Resolution

Always specify a node set which includes current node as scope.

60021

ERROR: multi-nodes must be specified for shared class

Explanation

Only one node is specified for shared class scope.

Resolution

Multiple node identifiers must be specified in a scope.

60022

ERROR: node: unknown node

Explanation

Node *node* does not exist.

Resolution

First check the cluster system environment. Then, try again after changing the node identifier.

60023

ERROR: class: not shared by the current node

Explanation

Scope of *class* does not include the current node.
In the cluster system, *class* created in another node is specified.

Resolution

Execute command on a node sharing *class*. Or, change the specified class name.

60024

ERROR: node: current node not active in class

Explanation

Current node *node* is a standby node in shared class *class*.

Resolution

Execute the command from another active node.

60025

ERROR: too many disks in class

Explanation

There are too many disks registered with *class*.

Resolution

Create a new class.

60026

ERROR: device: not connected to node

Explanation

The physical disk *device* has not been connected to the node *node*, or *device* has not been registered with the cluster resource database.

Resolution

Check the system configuration and the resource configuration.

If *device* has not been connected to *node*, specify a correct physical device name and re-execute the command.

If *device* has not been registered with the resource database, perform automatic resource registration to register *device* with the resource database, and re-execute the command.

60027

ERROR: device: not a shared device

Explanation

Physical disk *device* is not a shared disk unit, or cluster shared disk definition may not be correct.

Resolution

Check the system configuration and execute the command.

60028

ERROR: device: already exists in class

Explanation

Physical disk *device* is already registered with *class*.

Resolution

Same physical disk cannot be registered with more than one class. Specify a correct physical disk.

60029

ERROR: device: already exists in another class

Explanation

Physical disk *device* is already registered with another class. This class is not shared by current node.

Resolution

Same physical disk cannot be registered with more than one class. Specify a correct physical disk.

60030

ERROR: object: physical disk device not found

Explanation

The physical disk *device* that composes the object *object* was not found.

Resolution

The physical disk *device* may be connected incorrectly or a disk failure may have occurred. Fix *device* connection or recover from disk failure. Alternatively, specify an object other than *object*.

60031

ERROR: physical disk device not found

Explanation

The physical disk *device* was not found.

Resolution

The physical disk *device* may be connected incorrectly or a disk failure may have occurred. Fix *device* connection or recover from disk failure.

60032

ERROR: device: no such device

Explanation

There is no physical disk *device*.

Resolution

Specify a correct physical disk name.

60033

ERROR: device: cannot open, errno=errno

Explanation

Cannot open physical disk *device*.

Resolution

Confirm that physical disk *device* is operating normally.

60034

ERROR: device: not a hard disk

Explanation

Physical disk *device* is not a hard disk.

Resolution

GDS cannot manage devices other than hard disks.

60035

ERROR: device: disk driver *driver* not supported

Explanation

Physical disk unit with the driver name *driver* is not supported.

Resolution

There is no resolution.

60036

ERROR: device: illegal format

Explanation

Physical disk *device* format is illegal.

Resolution

Check the format status.

60037

ERROR: object: device busy

Explanation

Object *object* is in use.

Resolution

Change to unused status and execute the command again.

60038

ERROR: *object*: linked to a cluster service

Explanation

Object *object* is used in a cluster application.

Resolution

Check the cluster environment settings.

60039

ERROR: *device*: configuration information exists in private slice

Explanation

Configuration information exists in the private slice of physical disk *device*, and registering *device* failed. The following events a) to d) are the possible causes.

- a) *device* is already registered with another class.
- b) *device* is registered with a class in another domain.
- c) The entire SDX disk was copied or is being copied to *device* with the disk unit's copy function.
- d) After registering *device* with the class, it was not removed normally, and the private slice and configuration information remain existent in *device* illegally.

Resolution

Check the system configuration and so on, and identify the cause among a) to d). If the cause is a) or b), specify a physical disk not registered with a class. If it is c) or d), use the `format(1M)` command and remove the slice whose tag is 33 as needed. To check slice tag values, use the `prtvtoc(1M)` command.

60040

ERROR: *device*: type cannot be specified except undef

Explanation

The class scope includes a node to which physical disk *device* is unconnected and the type can only be set to "undef" for registering *device* to the class.

Resolution

See "[Appendix D Command Reference](#)."

60041

ERROR: *class*: *device*: type cannot be specified except undef

Explanation

Class *class* includes a switch group and the type can only be set to "undef" for registering physical disk *device* to *class*.

Resolution

See "[Appendix D Command Reference](#)."

60042

```
ERROR: device: write error, errno=errno
```

Explanation

A write error occurred in physical disk *device*.

Resolution

A disk failure may have occurred. Identify the cause by referring to disk driver log messages and so on, and recover the disk. If *device* is a write-locked disk that is connected to a switch group, register writable disks first.

60043

```
ERROR: device: disk ID information differs from all disks in specified class
```

Explanation

The specified class includes no disk with the same disk ID information (class name and disk name) as that stored in the private slice of physical disk *device*. For this reason, *device* cannot be registered with the specified class.

Resolution

Check the system configuration and so on, and specify a physical disk and class name properly.

60044

```
ERROR: object: name already assigned to another object
```

Explanation

Object name *object* already exists within class. You cannot create multiple objects with the same name.

Resolution

Specify another name and execute the command again.

60045

```
ERROR: cannot connect to sdxclld
```

Explanation

Connecting to sdxclld, a GDS cluster linked daemon was unsuccessful.

Resolution

When unable to identify the cause, collect investigation material and contact your local customer support.

60046

`ERROR: physical device driver returned an error, errno=errno`

Explanation

Physical disk driver returned an error.

Resolution

Identify the cause by referring to error number and message log.

60047

`ERROR: special file operation failed, errno=errno`

Explanation

Operating special file was unsuccessful.

Resolution

Identify the cause by referring to error number and GDS log message.

60048

`ERROR: sfdsk driver returned an error, errno=errno`

Explanation

GDS driver returned an error.

Resolution

Identify the cause by referring to error number, GDS log message, and syslog message.

60049

`ERROR: sfdsk driver returned a temporary error, try again for a while`

Explanation

GDS driver returned a temporary error.

Resolution

Execute the command again after a while.

60050

`ERROR: class: class closed down`

Explanation

`class` is closed. You cannot operate objects within a closed class.

Resolution

A number of disk failures may have occurred. Identify the cause by referring to object status, GDS log message, and syslog message. For information on recovery, see "[F.1.4 Class Status Abnormality](#)."

60051

ERROR: class: class closed down on another node

Explanation

class is closed on another node. You cannot operate objects within a closed class.

Resolution

Recover closed *class*. A number of disk failures may have occurred. Identify the cause by referring to object status, GDS log message, and syslog message. For information on recovery, see "[F.1.4 Class Status Abnormality](#)."

60052

ERROR: keep disk cannot be specified for local or shared class

Explanation

Keep disk cannot be registered with local or shared class.

Resolution

Register keep disk with root class.

60053

ERROR: too many keep disks specified

Explanation

The number of specified keep disks exceeds that of undefined disks.

Resolution

When registering multiple keep disks, specify the same number or more as undefined disks at the same time.

60054

ERROR: class: already root class exists

Explanation

You tried to create a new class when there is already a root *class*. You can only create one root class within a node.

Resolution

Specify the existing root *class* or a local class.

60055

ERROR: *pslice* is not backup slice, correct slice tag and size by format(1M) command

Explanation

Physical slice *pslice* is not a backup slice.

Resolution

Change the tag of physical slice *pslice* to "backup" with the format(1M) command and change its size to that of total disk cylinders available in the entire disk.

60056

ERROR: *device* contains overlapping slices

Explanation

Physical disk *device* contains a physical slice with overlapping cylinder.

Resolution

Use the format (1M) command and correct the physical slice configuration.

60057

ERROR: *device*: one or more free slices are required, change a slice size to 0 by format(1M) command

Explanation

The number of slices created on the physical disk *device* has reached the upper limit (seven). The slice numbers are all in use and *device* cannot be registered as a keep disk.

Resolution

Use the format(1M) command to remove one or more slices from device.

60059

ERROR: *pslice*: invalid partition tag, tag=*tag_num*

Explanation

Partition tag value *tag_num* for physical slice *pslice* is invalid.

Resolution

Change the tag of *pslice* to "unassigned" or the size of *pslice* to 0 (zero) with the format(1M) command.

60060

ERROR: *device*: no enough unassigned disk space, reserve enough space by format(1M) command

Explanation

Physical disk *device* does not have sufficient free disk space.

Resolution

Use the format (1M) command and create sufficient free disk space or swap area.
For the size see "[A.2.6 Disk Size](#)."

60061

ERROR: device: no unassigned disk space nor swap space, reserve enough space by format(1M) command

Explanation

Physical disk *device* does not have free disk space or swap area.

Resolution

Use the format (1M) command and create sufficient free disk space or swap area.
For the size see "[A.2.6 Disk Size](#)."

60062

ERROR: device: no enough unassigned disk space nor swap space, reserve enough space by format(1M) command

Explanation

Physical disk *device* does not have sufficient free disk space or swap area.

Resolution

Use the format (1M) command and create sufficient free disk space or swap area.
For the size see "[A.2.6 Disk Size](#)."

60063

ERROR: pslice: too large swap device

Explanation

The size of the swap device *pslice* is too large. In 32-bit Solaris, a disk with the swap device of which size is equal to or larger than 2 GB cannot be registered with the root class as the keep disk.

Resolution

Reduce the swap disk size to a size smaller than 2GB. Alternatively, change the system to 64-bit Solaris.

60064

ERROR: device: invalid physical device, managed by driver

Explanation

Specified physical disk is managed by *driver*, and is therefore invalid.

Resolution

Confirm the I/O and cluster system configuration, and specify a correct physical disk name.

60065

ERROR: device: disk driver driver not supported for root class

Explanation

Physical disk device with driver name *driver* cannot be registered with root class.

Resolution

Register the disk with local class or shared class as needed.

60066

ERROR: device: IDE disk cannot be specified as spare disk

Explanation

The *device* is an IDE disk. It cannot be used as a spare disk.

Resolution

For spare disks, use disks other than IDE disks.

60067

ERROR: class: no such class

Explanation

Cannot find *class*.

Resolution

Check GDS configuration.

60068

ERROR: group: not a group

Explanation

group is not a group name.

Resolution

There is another object with the name *group* within class. Check configuration.

60069

ERROR: group is a lower level stripe group

Explanation

Group *group* is a stripe group connected to another group. Disks and groups cannot be connected to or disconnected from *group*.

Resolution

Disconnect *group* from the higher level group as necessary.

60070

ERROR: *group*: connected to a lower level stripe group

Explanation

Group *group* is connected to a lower level stripe group. Disks and groups cannot be connected to or disconnected from *group*.

Resolution

Disconnect the higher level stripe group of *group* from its higher level group as necessary.

60071

ERROR: too many groups in *class*

Explanation

Class *class* already has the maximum number of groups possible. A maximum of 100 groups can be created within, and a maximum of 1024 groups can be created within a local class or a shared class.

Resolution

Create a new class.

60073

ERROR: too many disks and/or groups are connected to *group*

Explanation

The maximum number of disks or lower level groups are already connected to group *group*.

Resolution

There is no resolution.

60074

ERROR: *object*: smaller than stripe width of group *group*

Explanation

Since the available size of disk or lower level group indicated by *object* is smaller than the stripe width of group *group*, *object* cannot be connected to *group*.

Resolution

Execute the command after indicating a disk or lower level group with sufficient size, or remove group *group* and adjust the stripe width to a smaller size.

60075

ERROR: class: three or more nodes exist in class scope

Explanation

The scope of class *class* includes 3 or more nodes. Classes whose scope includes 3 or more nodes cannot include switch groups.

Resolution

Specify a shared class of which scope is 2 nodes.

60076

ERROR: disk: type disk exists in class

Explanation

Class *class* includes disk *disk* of the *type* type. The attempted operation is not supported for classes that include *type* type disks.

Resolution

Delete *disk* from *class* as needed and re-execute the command.

60077

ERROR: class includes a group that cannot exist together with a switch group

Explanation

The class *class* includes one of the following groups and a switch group cannot be created.

- Mirror group
- Stripe group
- Concatenation group to which any lower level switch group is not connected

Resolution

Specify a correct class name.

60078

ERROR: active disk must be specified

Explanation

When creating a switch group, specify the active disk in the `-a actdisk=disk` option.

Resolution

Specify the option properly, referring to "[Appendix D Command Reference](#)."

60079

ERROR: disk: active disk not specified by -d option

Explanation

Active disk *disk* was not specified in the -d option.

Resolution

When creating a switch group, specify one of the disks specified in the -d option of the sdxdisk -C command as the active disk.

60080

```
ERROR: too many disks specified
```

Explanation

Too many disks specified.

Resolution

Specify a proper number of disks.

For the number of disks that can be connected to a group, see "[A.1.3 Number of Disks.](#)"

60081

```
ERROR: disk: physical scope is not included in class scope
```

Explanation

The physical scope of disk *disk* is not included in the class scope and *disk* cannot be connected to a switch group.

Resolution

Specify a correct disk name. For the physical scope of a disk, use the sdxinfo -D command and check the displayed DEVCONNECT field. For the class scope, use the sdxinfo -C command and check the displayed SCOPE field.

60082

```
ERROR: disk: physical scope must be same as class scope
```

Explanation

The physical scope of the active disk matches the class scope and therefore the physical scope of an inactive disk must match the class scope. The physical scope of disk *disk* does not match the class scope and *disk* cannot be connected to the switch group as the inactive disk.

Resolution

Specify a correct disk name. For the physical scope of a disk, use the sdxinfo -D command and check the displayed DEVCONNECT field. For the class scope, use the sdxinfo -C command and check the displayed SCOPE field.

60083

```
ERROR: disk: physical scope must include only node
```

Explanation

The physical scope of inactive disk *disk* can include only node *node*.

Resolution

Specify a correct disk name, referring to "[Appendix D Command Reference](#)."

60084

ERROR: *disk*: class scope is not included in physical scope

Explanation

The class scope is not included in the physical scope of disk *disk* and *disk* cannot be connected to a group other than a switch group.

Resolution

Specify a correct disk name. For the physical scope of a disk, use the `sdxinfo -D` command and check the displayed DEVCONNECT field. For the class scope, use the `sdxinfo -C` command and check the displayed SCOPE field.

60085

ERROR: *disk*: no such disk

Explanation

There is no *disk*.

Resolution

Check GDS configuration.

60086

ERROR: *object*: already connected to group

Explanation

Disk or group indicated by *object*, is already connected to group *group*.

Resolution

Specify a correct disk name or group name.

60087

ERROR: *disk* is a spare disk

Explanation

disk is a spare disk. You cannot connect a spare disk to a group.

Resolution

Change the disk attributes to undefined, and execute the command again.

60088

ERROR: *object* not in status status

Explanation

Object status is not *status* status.

Resolution

Confirm that object status is *status*, and execute the command again.

60089

ERROR: *object* too small

Explanation

The size of *object* is too small.

Resolution

Check the necessary object size and specify a larger object.

60090

ERROR: another disk must be connected to *group*

Explanation

You must connect another disk to *group*.

Resolution

Connect another disk and execute the command again.

60091

ERROR: invalid physical slice number, *pslice_num*

Explanation

Invalid physical slice number *pslice_num* is specified.

Resolution

For physical slice number *pslice_num*, specify an integer from 0 to 7, excluding 2.

60092

ERROR: *pslice* is backup slice

Explanation

Physical slice *pslice* is a backup slice.

Resolution

For physical slice number, specify an integer from 0 to 7, excluding 2.

60093

ERROR: *pslice* is private slice

Explanation

Physical slice *pslice* is a private slice.

Resolution

Do not specify a private slice for physical slice number.

60094

ERROR: *pslice*: corresponding volume attributes must be specified

Explanation

Volume attributes corresponding to physical slice *pslice* is not specified.

Resolution

See "[Appendix D Command Reference](#)."

60095

ERROR: *object*: invalid size

Explanation

Object size of *object* is invalid.

Resolution

See "[Appendix D Command Reference](#)."

60096

ERROR: two or more keep disks cannot be connected to a group

Explanation

More than one keep disk cannot be connected to group.

Resolution

See "[Appendix D Command Reference](#)."

60097

ERROR: two or more single disks cannot be connected to a group

Explanation

More than one single disk cannot be connected to a group.

Resolution

See "[Appendix D Command Reference](#)."

60098

ERROR: both keep and single disks cannot be connected to a group

Explanation

Both keep and single disks cannot be connected to a group at the same time.

Resolution

See "[Appendix D Command Reference](#)."

60099

ERROR: *disk*: keep disk cannot be connected to existing group

Explanation

Keep disk *disk* cannot be connected to an existing group.

Resolution

See "[Appendix D Command Reference](#)."

60100

ERROR: *disk*: single disk cannot be connected to existing group

Explanation

Single disk *disk* cannot be connected to the existing group.

Resolution

See "[Appendix D Command Reference](#)."

60101

ERROR: two or more IDE disks cannot be connected to a group

Explanation

More than one IDE disk cannot be connected to a group.

Resolution

When mirroring an IDE disk, use it with disks other than IDE disks.

60102

ERROR: two or more disks cannot be connected to a group

Explanation

If the FJSVsdxl package of PRIMECLUSTER GDS is not installed, multiple disks cannot be connected to a group.

Resolution

If FJSVsdxl has not been installed normally, reinstall it.

60103

ERROR: *group*: a lower level switch group is connected

Explanation

A switch group has been connected to the group *group* and it is impossible to connect a disk to *group* or *group* to another group.

Resolution

Specify a correct higher level group name.

60104

ERROR: *disk*: not a bootable device

Explanation

Disk *disk* cannot be booted.

Resolution

Confirm the disk configuration. When you cannot specify the cause, collect investigation material and contact your local customer support.

60105

ERROR: too few valid configuration database replicas

Explanation

There is not sufficient valid configuration database for class. This message is displayed when the majority of disks registered with class are unavailable.

Leaving it as is may cause serious problems.

Resolution

For details, see "[F.1.4 Class Status Abnormality](#)."

60106

ERROR: *group*: no ENABLE disk in group

Explanation

Group *group* does not have a disk in ENABLE status connected.

Resolution

See "[F.1.2 Disk Status Abnormality](#)" and recover disks connected to the group.

60107

ERROR: msec: invalid delay value

Explanation

Delay time *msec* is invalid.

Resolution

See "[Appendix D Command Reference](#)."

60108

ERROR: lgroup: mirror group cannot be connected to another group

Explanation

Since group *lgroup* is a mirror group, it cannot be connected to another group.

Resolution

Specify a correct lower level group name.

60110

ERROR: lgroup: same type as higher level group hgroup

Explanation

Since the type attributes for group *lgroup* and higher level group *hgroup* are the same, *lgroup* cannot be connected to *hgroup*.

Resolution

Specify a correct group name.

60111

ERROR: hgroup: same name as a lower level group

Explanation

For the higher level group name, name *hgroup*, which is the same as the lower level group was indicated.

Resolution

Different names must be indicated for higher level group and lower level group.

60112

ERROR: hgroup: any group cannot be connected to a switch group

Explanation

Group *hgroup* is a switch group and any group cannot be connected to.

Resolution

Specify a correct group name.

60114

ERROR: *group*: is a lower level concatenation group

Explanation

The group *group* is a concatenation group connected to another group. Any group cannot be connected to *group*.

Resolution

Disconnect *group* from the higher level group as needed and re-execute the command.

60115

ERROR: *lgroup*: stripe group cannot be connected to concatenation group

Explanation

The group *lgroup* is a stripe group and cannot be connected to a concatenation group.

Resolution

Specify a correct group name.

60116

ERROR: *lgroup*: switch group cannot be connected to mirror group

Explanation

The group *lgroup* is a switch group and cannot be connected to a mirror group.

Resolution

Specify a correct group name.

60117

ERROR: *lgroup*: switch group cannot be connected to stripe group

Explanation

The group *lgroup* is a switch group and cannot be connected to a stripe group.

Resolution

Specify a correct group name.

60118

ERROR: *hgroup*: disk is connected

Explanation

A disk is already connected to concatenation group *hgroup*, and a switch group cannot be connected.

Resolution

Specify a correct group name.

60119

ERROR: *size*: must be more than zero

Explanation

size must be a positive integer.

Resolution

Specify the correct size.

60120

ERROR: *size*: invalid size

Explanation

The indicated size *size* is invalid.

Resolution

See "[Appendix D Command Reference](#)."

60121

ERROR: *group*: no such group

Explanation

There is no *group*.

Resolution

Specify a correct group name.

60122

ERROR: *group*: not the highest level group

Explanation

Group *group* is not the highest level group.

Resolution

Disconnect *group* from the higher level group as necessary.

60123

ERROR: too many volumes exist in *object*

Explanation

The maximum number of volumes, already exists in the *object*, that is a class, a group, or a single disk.

Resolution

Create volumes in another group or on a single disk as needed.

60124

ERROR: too many volumes with physical slices in *object*

Explanation

In the *object*, the number of volumes which consist of physical slices has reached the maximum. *object* is a name of group or single disk.

Resolution

Change the physical slice attributes for volumes in *object* from on to off using the `sdxattr -V` command as needed, and execute the command again.

60125

ERROR: *group*: volume with physical slice(s) cannot be created

Explanation

Although a physical slice cannot be created in group *group*, an attempt was made to create a volume with one or more physical slices.

Resolution

Create a volume without any physical slice for *group* either by:

- Turning "off" the [Physical Slice] option in the Volume Configuration dialog box when creating the volume for *group* in GDS Management View
- Using the `-a pslice=off` option when creating the volume for *group* with the `sdxvolume -M` command

If *group* is a mirror group, connect one or more disks to the *group* in order to create a volume with a physical slice.

60126

ERROR: *object*: no enough space

Explanation

There is not enough space in the group or single disk indicated by *object*.

Resolution

Change the size you specify as needed.

60127

ERROR: *disk*: not a single disk

Explanation

The disk *disk* is not a single disk.

Resolution

Specify a correct disk name.

60128

ERROR: *disk*: *status* disk connected to *group*

Explanation

disk in *status* status is connected to *group*.

Resolution

Restore the status of *disk* and execute the command again.

60129

ERROR: *status* disk exists in *group*

Explanation

The *disk* in *status* status is connected to *group*, or it connected to the lower level group of *group*.

Resolution

Confirm the disk status, and please execute a command again after canceling the *status* status if necessary.

60130

ERROR: *disk*: no such disk in *group*

Explanation

disk is not connected to *group*.

Resolution

Specify a correct disk name and group name.

60131

ERROR: *disk*: *device*: device busy

Explanation

Disk *disk* (physical disk name is *device*) is in use.

Resolution

Change to unused status, and execute the command again.

60132

ERROR: too many nodes specified

Explanation

Too many nodes were specified.

Resolution

Specify the nodes included in the scope of the class.

60133

ERROR: option: cannot be specified for root nor local class

Explanation

Command option *option* cannot be specified for a root nor local class.

Resolution

See "[Appendix D Command Reference](#)."

60134

ERROR: node: not in scope, class=class

Explanation

Node *node* is not included in the scope of the class *class*.

Resolution

Confirm your GDS configuration, specify a node included in the scope of the class, and try the command again.

60135

ERROR: volume: cannot start, class closed down, node=node

Explanation

Starting volume *volume* was unsuccessful due to the class to which volume *volume* belongs was blockading it on node *node*.

Resolution

Recover the blockage of the class to which volume *volume* belongs. A number of disk failures may have occurred. Identify the cause by referring to object status, GDS log message, and syslog message.

For information on recovery, see "[F.1.4 Class Status Abnormality](#)."

60136

ERROR: volume: no such volume

Explanation

There is no *volume*.

Resolution

Specify a correct volume name.

60137

ERROR: object in status status

Explanation

object is in *status* status.

Resolution

Confirm the object status, and cancel the status *status* if necessary.

60138

ERROR: some ACTIVE volumes exist in class, node=node

Explanation

Volumes of class *class* are already active on node *node*, and volumes of *class* cannot be started or created on this node.

Resolution

Create volumes of class *class* on node *node*. Alternatively, as needed, stop all the volumes within *class* on *node* and then start or create volumes on this node.

If a cluster application to which *class* resources are registered is operating or is standby on *node*, before stopping volumes on *node*, stop the cluster application.

60139

ERROR: volume: some ACTIVE volumes exist in class, node=node

Explanation

Volumes of class *class* are already active on node *node*, and volume *volume* cannot be started on this node.

Resolution

Stop all the volumes within class *class* on node *node* as needed, start volume *volume* on this node.

If a cluster application to which *class* resources are registered is operating or is standby on *node*, before stopping volumes on *node*, stop the cluster application.

60140

ERROR: volume: active disk not connected to node

Explanation

The active disk of a switch group including switch volume *volume* is not connected to node *node* and the *volume* cannot be started on the *node*.

Resolution

If an inactive disk is connected to a switch group including *volume*, use the `sdxattr -G` command and switch the inactive disk to the active disk in order to start *volume* on *node*.

60141

```
ERROR: volume: active disk of lower level group group is not connected to node
```

Explanation

The active disk of lower level switch group *group* is not connected to node *node*, and volume *volume* cannot be started on *node*. *group* is a lower level switch group that is connected to the highest level group to which *volume* belongs.

Resolution

If an inactive disk is connected to *group*, switch the active disk to the inactive disk with the `sdxattr -G` command in order to enable *volume* startup on *node*.

60142

```
ERROR: lock is set on volume volume, node=node
```

Explanation

The "Lock volume" mode for volume *volume* on node *node* is turned on.

Resolution

Turn off the "Lock volume" mode, or use the `-e unlock` option as necessary.

60143

```
ERROR: volume: cannot stop, class closed down, node=node
```

Explanation

Stopping volume *volume* was unsuccessful due to the class to which volume *volume* belongs was blockading it on node *node*.

Resolution

Recover the blockage of the class to which volume *volume* belongs. A number of disk failures may have occurred. Identify the cause by referring to object status, GDS log message, and syslog message.
For information on recovery, see "[F.1.4 Class Status Abnormality](#)."

60144

```
ERROR: object.volume: status slice exists in object
```

Explanation

Slice *object.volume* in *status* status exists in *object*.

Resolution

Confirm the object status, and cancel the status *status* if necessary.

60145

ERROR: object in status status, node=node

Explanation

object on *node* is in *status* status.

Resolution

Confirm the object status, and cancel the status *status* if necessary.

60146

ERROR: volume: stripe type volume cannot be resized

Explanation

The volume *volume* cannot be resized since it is a volume created in a stripe group.

Resolution

See "Stripe Type Volume and Concatenation Type Volume Expansion" in "[A.2.15 Online Volume Expansion](#)."

60147

ERROR: volume: concat type volume cannot be resized

Explanation

The volume *volume* cannot be resized since it is a volume created in a concatenation group.

Resolution

See "Stripe Type Volume and Concatenation Type Volume Expansion" in "[A.2.15 Online Volume Expansion](#)."

60148

ERROR: volume: consists of multiple mirror slices

Explanation

The volume *volume* cannot be resized since it is a mirror volume that consists of multiple mirror slices.

Resolution

Disconnect disks and lower level groups from the mirror group so that *volume* is composed of only one mirror slice, and try this command again.

60149

ERROR: spare disk connected for disk

Explanation

Spare disk is connected instead of *disk*.

Resolution

First recover the *disk* status.

60150

ERROR: *object.volume* is only valid slice

Explanation

object.volume is the only valid slice within volume *volume*. You cannot continue your operation because data within *volume* may be lost if you continue your operation.

Resolution

If *volume* is a mirror volume, you will be able to continue your operation by recovering the mirroring (for example, by connecting a new disk to group).

If *volume* is not a mirror volume, remove *volume* as necessary.

60151

ERROR: *object*: the last disk or group in lower level group *group*

Explanation

Object *object* is the only disk or group connected to the lower level group *group*. You cannot disconnect *object* from *group*.

Resolution

Disconnect *group* from the higher level group as necessary.

60152

ERROR: *object*: not connected to the end of concatenation group *group*

Explanation

The object *object* is not the disk or switch group that was last connected to the concatenation group *group*. It is impossible to disconnect *object* from *group*.

Resolution

To disconnect disks and groups from *group*, perform disconnection in inverse order of connection. For the order of connecting disks and groups to *group*, use the `dxinfo -G` command and check the displayed DISKS field.

60153

ERROR: *object*: disk space is assigned to volume *volume*

Explanation

Disk space of disk or lower switch group *object* is assigned to volume *volume*, and *object* cannot be disconnected from the concatenation group.

Resolution

If needed, remove *volume* first.

60154

ERROR: group: inactive disk is connected

Explanation

The group *group* has a connected inactive disk and physical disk swap or disconnection cannot be performed for the active disk of the *group*.

Resolution

As needed, switch the active disk with the `sdxattr -G` command and perform physical disk swap or disconnection for the previous active disk.

60155

ERROR: disk: not inactive disk

Explanation

The disk *disk* is the active disk of a lower level switch group, and physical disk swap cannot be performed.

Resolution

As needed, switch the active disk with the `sdxattr -G` command, and perform physical disk swap on the previous active disk.

60156

ERROR: lgroup: no such group in hgroup

Explanation

Group *lgroup* is not connected to group *hgroup*.

Resolution

Specify a correct group name.

60157

ERROR: one or more volumes exist in group

Explanation

One or more volumes exist in group *group*.

Resolution

First, remove the volume if necessary.

60158

ERROR: *disk* connected to *group*

Explanation

disk is connected to *group*.

Resolution

First, disconnect the disk from group if necessary.

60159

ERROR: *disk*: The last ENABLE disk in class cannot be removed

Explanation

When there is a disk in SWAP or DISABLE status in class, you cannot remove the last disk in ENABLE status.

Resolution

First recover the disk in SWAP or DISABLE status. Or, you can register a new disk with class.

60160

ERROR: *disk*: cannot be removed to avoid class closing down

Explanation

Disk *disk* stores the configuration database and removal of *disk* will result in class closure. For this reason, the disk cannot be removed.

Resolution

Recover or remove a failed disk in the class that includes *disk* and then remove *disk*.

60161

ERROR: one or more volumes exist in *disk*

Explanation

Volume exists in *disk*.

Resolution

First, remove the volume as needed.

60162

ERROR: one or more groups exist in *class*

Explanation

Group exists in *class*.

Resolution

First, remove the group as needed.

60163

ERROR: *disk: status disk exists in class*

Explanation

disk in *status* status exists in *class*.

Resolution

First, restore the *disk* in *status* status.

60164

ERROR: *volume: status volume exists in class, node=node*

Explanation

volume in *status* status on *node* exists in *class*.

Resolution

First, change the volume status as needed.

60165

ERROR: *disk: no such disk*

Explanation

There is no disk *disk*.

Resolution

Specify a correct disk name.

60166

ERROR: *volume: not associated with object*

Explanation

You cannot specify slice by combining disk or group indicated by *object*, and volume indicated by *volume*.

Resolution

See "[Appendix D Command Reference](#)" and specify a correct disk or group name and volume name.

60167

ERROR: *volume is single volume*

Explanation

Volume is a single volume. Slices of single volumes cannot be detached.

Resolution

Specify a correct volume name.

60168

```
ERROR: volume: not a mirror volume
```

Explanation

Volume *volume* is not a mirror volume.

Resolution

Specify a correct volume name.

60169

```
ERROR: object: not connected to the highest level group
```

Explanation

Disk or group indicated by *object* is not connected to the highest-level group.

Resolution

See "[Appendix D Command Reference](#)" and specify a correct disk name or group name.

60170

```
ERROR: disk: The last ENABLE disk in class cannot be swapped
```

Explanation

You cannot swap the last ENABLE status disk in class.

Resolution

Depending on the configuration, use a different method to avoid such situations. For example, you can register a new disk with class.

60171

```
ERROR: disk: keep disk cannot be swapped out
```

Explanation

Keep disk *disk* cannot be swapped.

Resolution

Change the type attribute of *disk*, and execute the command again.

60172

ERROR: disk: volume in status status

Explanation

There is a volume related to disk in *status* status.

Resolution

First, recover the volume in *status* status.

60173

ERROR: disk: the highest level group is not a mirror group

Explanation

The highest level group of disk *disk* is not a mirror group.
You cannot swap *disk*.

Resolution

Disconnect *disk* from the group, and execute the command in a situation where it is not connected to any group.

60174

ERROR: disk: cannot be swapped to avoid class closing down

Explanation

The disk *disk* stores the configuration database and swapping *disk* will close down the class. Therefore, swapping the disk is prevented.

Resolution

Replace a failed disk included in the class to which *disk* is registered and swap *disk*.

60175

ERROR: disk: device: device busy on node node

Explanation

Disk *disk* (the physical disk name is *device*) is in use on *node*.

Resolution

Change the status to unused and execute the command again.

60176

ERROR: disk: device: cannot open, errno=errno

Explanation

Cannot open physical disk *device*.

Resolution

Confirm that physical disk *device* is operating normally.

60177

ERROR: disk: device: not a hard disk

Explanation

Physical disk *device* is not a hard disk.

Resolution

Specify a hard disk.

60178

ERROR: disk: device: illegal format

Explanation

Physical disk *device* format is illegal.

Resolution

Check the format of the physical disk.

60179

ERROR: disk busy - /dev/rdisk/*pslice*

Explanation

Physical disk *pslice* is in use.

Resolution

Change to unused status and execute the command again.

60180

ERROR: disk: device: linked to a cluster service

Explanation

Physical disk *device* is used in a cluster application.

Resolution

Check the settings for cluster environment.

60181

ERROR: disk: device: not enough size

Explanation

Physical disk size is too small.

Resolution

Specify a physical disk with sufficient size.

60182

```
ERROR: disk: device: not connected to node
```

Explanation

Physical disk *device* is not connected to *node*. Or, cluster shared disk definition may not be correct.

Resolution

Check the system configuration and execute the command.

60183

```
ERROR: disk: device: invalid disk, managed by driver
```

Explanation

Specified disk (corresponding physical disk name is *device*) is managed by *driver*, and is therefore invalid.

Resolution

Confirm the I/O and cluster system configuration, and specify a correct physical disk name.

60184

```
ERROR: object.volume: status slice exists in group
```

Explanation

Slice *object.volume* in *status* status already exists in group *group*.

Resolution

Check the object status and cancel *status* status if necessary.

60185

```
ERROR: disk : device : invalid disk on node node, managed by driver
```

Explanation

On node *node*, physical disk *device* of disk *disk* has been managed by driver *driver* and cannot be handled.

Resolution

Check I/O configurations and cluster system configurations, and specify a correct physical disk name.

60186

`ERROR: disk : device : cannot open on node node, errno=errno`

Explanation

Physical disk *device* cannot be opened on node *node*.

Resolution

Check whether or not the physical disk *device* is operating normally.

60187

`ERROR: disk : device : node : not a hard disk`

Explanation

Physical disk *device* on node *node* is not a hard disk.

Resolution

Specify a hard disk.

60188

`ERROR: disk : device : node : illegal format`

Explanation

The format of physical disk *device* on node *node* is incorrect.

Resolution

Check the format.

60189

`ERROR: disk busy on node node - /dev/rdisk/pslice`

Explanation

Physical disk *pslice* on node *node* is in use.

Resolution

Change the status to unused and retry the command.

60190

`ERROR: disk : device : node : linked to a cluster service`

Explanation

Physical disk *device* on node *node* is being used by cluster applications.

Resolution

Check the cluster environment configurations.

60191

```
ERROR: disk : device : node: not enough size
```

Explanation

The size of physical disk *device* on node *node* is too small.

Resolution

Replace it with a physical disk of sufficient size.

60192

```
ERROR: disk : device : read error, errno=errno
```

Explanation

A read error occurred on physical disk *device*.

Resolution

The possible cause is a disk failure. Identify the cause based on disk driver log messages and so on and restore the disk.

60193

```
ERROR: disk : device : node: read error, errno=errno
```

Explanation

A read error occurred on physical disk *device* on node *node*.

Resolution

The possible cause is a disk failure. Identify the cause based on disk driver log messages and so on and restore the disk.

60194

```
ERROR: object: read error, errno=errno
```

Explanation

Read error occurred on the disk indicated by *object*, or on a disk connected to group indicated by *object*, or on a disk connected to a lower level of the group indicated by *object*.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and promptly recover the disk. Otherwise, indicate another disk or group that is normal.

60195

```
ERROR: object.volume : no such slice
```

Explanation

You cannot specify slice by combining disk or group indicated by *object*, and volume indicated by *volume*.

Resolution

See "[Appendix D Command Reference](#)" and specify a correct disk or group name and volume name.

60196

ERROR: *disk*: not active disk

Explanation

Disk *disk* is not the active disk.

Resolution

Specify the active disk.

60197

ERROR: *disk*: I/O error not occur

Explanation

An I/O error has not been occurred on the disk *disk*. There is no need to restore *disk* using the `sdxfix -D` command.

Resolution

There is no resolution.

60198

ERROR: *class*: not close down

Explanation

Class *class* is not closed. Recovery *class* with the `sdxfix -C` command is not required.

Resolution

There is no resolution.

60199

ERROR: no valid configuration database

Explanation

No valid configuration database of the class is found. If all (or most of) the disks that are registered with the class are unavailable, this message is output.

Resolution

See "[F.1.4 Class Status Abnormality](#)."

60200

ERROR: class: closed down on all nodes in class scope

Explanation

During recovery of class *class* from the closed status, *class* was closed on all the nodes within the class scope.

Resolution

Re-execute the `sdxfix -C` command.

60201

ERROR: disk: connected to switch group

Explanation

Disk *disk* is connected to a switch group and an I/O error status of the *disk* cannot be removed.

Resolution

Make *disk* exchangeable with the `sdxswap -O` command or [Operation]:[Swap Physical Disk] in GDS Management View to remove an I/O error status of *disk*. After this procedure, swap the physical disks of *disk* according to need and then make the *disk* usable again with the `sdxswap -I` command or [Operation]:[Restore Physical Disk] in GDS Management View.

60202

ERROR: volume: cannot restart to copy, cancel current interrupted copy operation by sdxcopy command with -C option

Explanation

Copying process within volume could not be resumed.

Resolution

Cancel the interrupted copying process according to need and resume copying.

60203

ERROR: attribute=value: cannot modify type attribute of root class

Explanation

The root class type attribute cannot be changed.

Resolution

Specify a correct class name.

60204

ERROR: class: class names must be unique within a domain

Explanation

Class *class* with the same name as class already exists within the cluster domain.

Resolution

When renaming a class resulted in this message, specify another class name.

When expanding the class scope resulted in this message, see "[Important Point 2](#)" in "[A.2.26 Changing Over from Single Nodes to a Cluster System](#)."

60205

ERROR: class: volume minor numbers must be unique within a domain

Explanation

A volume with the same minor number as that of a volume in the class *class* was found within the cluster domain.

Resolution

See "[Important Point 2](#)" in "[A.2.26 Changing Over from Single Nodes to a Cluster System](#)."

60206

ERROR: one or more disks not connected to node

Explanation

Disks that are not connected to node *node* or for which disk resources are not created yet exist in the class.

Resolution

Check the hardware configuration and the disk configuration within the class. If disk resources are not created yet, create those through the automatic resource registration.

For the automatic resource registration, see "[PRIMECLUSTER Cluster Foundation \(CF\) Configuration and Administration Guide](#)."

60207

ERROR: disk: IDE disk cannot be specified as spare disk

Explanation

The *disk* is an IDE disk. It cannot be used as a spare disk.

Resolution

For spare disks, use disks other than IDE disks.

60208

ERROR: object: no such object

Explanation

There is no *object*.

Resolution

Specify a correct object name.

60209

ERROR: class: shared objects information not yet available, try again for a while

Explanation

Share object information is not available yet.

Resolution

Wait until cluster control is activated, and execute the command again.

60210

ERROR: node: node in stopped status

Explanation

You cannot proceed with operation since *node* is in stopped status.

Resolution

Start the *node* and try again.

60211

ERROR: node: node in abnormal status

Explanation

node is in an abnormal status and cannot proceed with operation.

Resolution

Confirm the normal startup of *node* and try again.

60212

ERROR: cluster communication failure

Explanation

Cannot proceed with operation since communication failed with cluster.

Resolution

Check that the cluster system and GDS are operating normally.
After recovering, try again.

60213

ERROR: cluster communication failure, sdxerrno=sdxerrno

Explanation

Cannot proceed with operation since communication failed with cluster.

Resolution

Check that the cluster system and GDS are operating normally.

After recovering, try again.

60214

```
ERROR: cluster communication failure, remote-node=node, sdxerrno=sdxerrno
```

Explanation

Cluster communication with remote node *node* failed. Operation cannot be performed.

Resolution

Check that the cluster system and GDS are operating normally.

After recovering, try again.

60215

```
ERROR: class: not a root class
```

Explanation

Class *class* is not a root class. The possible causes are as follows.

- a) The specified class name is wrong.
- b) An attempt to use a function that is only for a root class was made for a local class or a shared class.
- c) In a system where the FJSVsdxsl package of PRIMECLUSTER GDS is not installed normally, an attempt to create a local class group or a shared class group was made.
- d) In a system where the FJSVsdxsl package of PRIMECLUSTER GDS Snapshot is not installed normally, an attempt to perform proxy operation was made for a local class or a shared class.

Resolution

If the cause is a) or b), see "[Appendix D Command Reference](#)" and specify a correct class name. If it is c) or d), install FJSVsdxsl or FJSVsdxsl normally.

60216

```
ERROR: disk: not a keep disk
```

Explanation

Disk *disk* is not a keep disk.

Resolution

See "[Appendix D Command Reference](#)" and specify a correct disk name.

60217

```
ERROR: disk: not connected to any group
```

Explanation

Disk *disk* is not connected to a group.

Resolution

See "[Appendix D Command Reference](#)" and specify a correct disk name.

60218

```
ERROR: volume: status volume exists in group
```

Explanation

There is *volume* in *status* status in *group*.

Resolution

Recover volume status, and execute the command again.

60219

```
ERROR: disk: not a system disk
```

Explanation

Disk *disk* is not a system disk.

Resolution

Specify a correct disk name, and execute the command again.

60221

```
ERROR: device: mandatory system disk must be registered to class
```

Explanation

The system disk *device* is not registered with the class *class*. *device* contains the slice currently operating as / (root), /usr, or /var, and to perform system disk setting, it is necessary to register *device* with *class*.

Resolution

Complete preparations properly referring to "[5.2.1 System Disk Settings](#)" or "[Appendix D Command Reference](#)" and execute the command again.

60223

```
ERROR: disk: mandatory system disk must be specified
```

Explanation

The system disk *disk* was not specified. *disk* contains the slice currently operating as / (root), /usr, or /var, and to cancel disk mirroring, it is necessary to specify *disk*.

Resolution

Specify all system disks with the slices currently operating as / (root), /usr, or /var, and execute the command again.

60224

```
ERROR: disk: two or more disks connected to group
```

Explanation

group to which *disk* is connected has two or more disks connected.

Resolution

See "[Appendix D Command Reference](#)" and complete preparation correctly. Then execute the command again.

60225

```
ERROR: root file system not mounted on volume
```

Explanation

Root file system is not mounted on volume.

Resolution

Confirm the configuration, and see "[Appendix D Command Reference](#)."

60226

```
ERROR: illegal slice name
```

Explanation

Slice name includes a "." (period).

Resolution

Specify the correct slice name.

60227

```
ERROR: disk.volume cannot be operated on the current node, take over by sdxslice command with -T option
```

Explanation

Slice *disk.volume* cannot be operated on the current node.

Resolution

Take over the slice by executing `sdxslice -T` command.

60228

```
ERROR: volume: physical slice attribute value is off
```


Explanation

The physical slice attribute value of volume *volume* is "off." A slice in a volume without physical slices cannot be detached.

Resolution

Retry the command after turning the physical slice attribute of *volume* to be "on" according to need.

60229

ERROR: *object*: device busy on node *node*

Explanation

object is in use on *node*.

Resolution

Change to unused status and execute the command again.

60230

ERROR: *class*: not a shared class

Explanation

class is not a shared class.

Resolution

Specify a shared class.

60231

ERROR: *param*: invalid parameter name

Explanation

Parameter name *param* is invalid.

Resolution

See "[Appendix D Command Reference](#)."

60232

ERROR: *param =val*: invalid parameter value

Explanation

Parameter value *val* is invalid.

Resolution

See "[Appendix D Command Reference](#)."

60233

ERROR: param: parameter name duplicated

Explanation

The same parameter name *param* has already been specified.

Resolution

You can only use a parameter name once.

60234

ERROR: copy_concurrency=val: value more than or equal to the number of actually running copy operations must be specified

Explanation

Value smaller than the number of copying currently in process has been specified for *copy_concurrency* parameter.

Resolution

Set a value more than or equal to the number of copying currently in process for *copy_concurrency* parameter.

60235

ERROR: mode=string: access mode duplicated

Explanation

Multiple access modes have been specified.

Resolution

Specify only one access mode.

60236

ERROR: mode=string: invalid access mode

Explanation

Access mode value *string* is invalid.

Resolution

See "[Appendix D Command Reference](#)."

60237

ERROR: volume: already started with different access mode, node=node

Explanation

On the node *node*, an attempt was made to start the volume *volume* specifying the access mode, but the volume *volume* is already active in an access mode other than the specified mode.

Resolution

Stop the volume *volume* and restart it as needed.

60238

ERROR: volume: related to proxy volume proxy

Explanation

Volume *volume* is a master volume related to proxy volume *proxy*.

Resolution

Cancel the relationship of master volume *volume* and proxy volume *proxy* according to need and execute again.

60239

ERROR: volume: related to master volume master

Explanation

Volume *volume* is a proxy volume related to master volume *master*.

Resolution

Cancel the relationship of master volume *master* and proxy volume *volume* according to need and execute again.

60240

ERROR: volume: related to proxy volume proxy with EC

Explanation

EC session exists between volume *volume* and proxy volume *proxy*.

Resolution

Cancel the relationship of master volume *volume* and proxy volume *proxy* according to need, and execute again.

60241

ERROR: volume: related to master volume master with EC

Explanation

EC session exists between volume *volume* and master volume *master*.

Resolution

Use the `sdxproxy Cancel` command and cancel the EC session between master volume *master* and proxy volume *volume* according to need. Alternatively, as needed, cancel the relationship of master volume *master* and proxy volume *volume*, and execute again.

60242

ERROR: volume: related to proxy volume proxy with TimeFinder

Explanation

There is a BCV pair between the volume *volume* and the proxy volume *proxy*.

Resolution

Break the relation between the master volume *volume* and the proxy volume *proxy* according to need, and try the command again.

60243

ERROR: volume: related to master volume master with TimeFinder

Explanation

There is a BCV pair between the volume *volume* and the master volume *master*.

Resolution

Break the relation between the master volume *master* and the proxy volume *volume* according to need, and try the command again.

60244

ERROR: volume: related to proxy volume proxy with SRDF

Explanation

There is an SRDF pair between the volume *volume* and the proxy volume *proxy*.

Resolution

Break the relation between the master volume *volume* and the proxy volume *proxy* according to need, and try the command again.

60245

ERROR: volume: related to master volume master with SRDF

Explanation

There is an SRDF pair between the volume *volume* and the master volume *master*.

Resolution

Break the relation between the master volume *master* and the proxy volume *volume* according to need, and try the command again.

60246

ERROR: proxy : no parted proxy volume in proxy group

Explanation

No parted proxy volume exists in proxy group *proxy*.

Resolution

According to need, part the proxy and retry the command.

60248

ERROR: volume: parted proxy volume

Explanation

Volume *volume* is a proxy volume parted from master volume.

Resolution

Rejoin the proxy volume to the master volume or break the relationship between them according to need and try this command again.

60249

ERROR: group: related to proxy group proxy

Explanation

Group *group* is a master group related to proxy group *proxy*.

Resolution

Cancel the relationship between master group *group* and proxy group *proxy* according to need, and execute again.

60250

ERROR: group: related to master group master

Explanation

Group *group* is a proxy group related to master group *master*.

Resolution

Cancel the relationship between master group *master* and proxy group *group* according to need, and execute again.

60251

ERROR: volume: related to master or proxy volume

Explanation

Volume *volume* is related to either master volume or proxy volume.

Resolution

Cancel the relationship between master and proxy according to need, and execute again.

60252

ERROR: volume: joined to master volume master

Explanation

Volume *volume* is a proxy volume joined to master volume *master*.

Resolution

Part volume *volume* from master volume *master*, or cancel the relationship with master volume according to need, and execute again.

60253

ERROR: volume: copying from master volume master

Explanation

Data is being copied from master volume *master* to volume *volume*.

Resolution

After the copying process is complete, execute again.

60254

ERROR: volume: copying from proxy volume proxy

Explanation

Data is being copied from proxy volume *proxy* to volume *volume*.

Resolution

After the copying process is complete, execute again.

60255

ERROR: class is root class

Explanation

The class *class* is the root class. The attempted operation is not supported for the root class.

Resolution

Specify a correct class name.

60256

ERROR: object: not volume nor group

Explanation

The object *object* is neither a volume nor a group.

Resolution

Check the GDS configuration and specify the correct volume name or group name. Execute again.

60257

ERROR: different types of objects, master=master, proxy=proxy

Explanation

Different types of objects *master* and *proxy* were specified as the master and proxy.

Resolution

As the master and proxy, specify one of the following combinations.

- Two volumes of the same type
- A mirror volume and a single volume
- Two mirror groups

For details, see "[D.14 sdxproxy - Proxy object operations.](#)"

60258

ERROR: *object*: same name as master

Explanation

The object name specified for proxy is the same as the master's name, *object*.

Resolution

Different object names must be specified for master and proxy.

60259

ERROR: *group*: not a mirror group

Explanation

Group *group* is not a mirror group.

Resolution

Specify a correct group name.

60261

ERROR: no volume exists in *group*

Explanation

There are no volumes in group *group*.

Resolution

Use `sdxvolume -M` command to create a volume within group *group*, and execute again.

60262

ERROR: too many proxy volumes are related to *master*

Explanation

There are too many proxy volumes related to master object *master*.

Resolution

Check the GDS configuration.

For details about the number of proxy volumes, see "[A.1.9 Number of Proxy Volumes.](#)"

60263

ERROR: master: corresponding proxy volume name must be specified

Explanation

There is no corresponding proxy volume name specified for master volume *master*.

Resolution

See "[D.14 sdxproxy - Proxy object operations.](#)"

60264

ERROR: proxy: proxy volume name duplicated

Explanation

The same proxy volume name *proxy* is specified for more than one master volume.

Resolution

Specify a different name, and execute again.

60265

ERROR: volume: no such volume in group

Explanation

Volume *volume* does not exist in group *group*.

Resolution

Specify a correct volume name or group name, and execute again.

60266

ERROR: object: object name duplicated

Explanation

The duplicate object name *object* was specified, or the group to which the volume *object* belongs and the volume *object* itself were specified at the same time.

Resolution

You can specify an object only once.

60267

ERROR: proxy: already parted

Explanation

Proxy volume *proxy* is already parted.

Resolution

Nothing needs to be done.

60268

ERROR: one point copy not available

Explanation

The OPC (One Point Copy) function is unavailable and proxy operations cannot be performed.

Resolution

See "(1) The Advanced Copy function cannot be used in master-proxy copying." in "F.1.7 Proxy Object Abnormality."

60269

ERROR: proxy: already joined

Explanation

Proxy object *proxy* is already joined.

Resolution

Nothing needs to be done.

60270

ERROR: proxy: not joined to master

Explanation

Proxy volume *proxy* is not joined to master volume.

Resolution

Join the proxy volume *proxy* to master volume according to need, and execute again.

60271

ERROR: proxy: no such proxy object

Explanation

Proxy object *proxy* cannot be found.

Resolution

Specify a correct proxy object name.

60272

ERROR: master: no such master object

Explanation

Master object *master* cannot be found.

Resolution

Specify a correct master object name.

60273

ERROR: volume: exists in proxy group

Explanation

Volume *volume* is a proxy volume in proxy group.

Resolution

Specify a proxy group and execute again.

60274

ERROR: group: not a proxy group

Explanation

group is not a proxy group group name.

Resolution

Check the GDS configuration, and specify a proxy group group name.

60275

ERROR: volume: copying with EC

Explanation

Volume *volume* is in the process of EC copy.

Resolution

After the copying process is complete, execute again.

60276

ERROR: volume: copying with OPC

Explanation

Volume *volume* is in the process of OPC copy.

Resolution

After the copying process is complete, execute again.

60277

ERROR: volume: copying with TimeFinder

Explanation

Volume *volume* is being copied with TimeFinder.

Resolution

Execute the command again after the copying process is complete.

60278

ERROR: volume: copying with SRDF

Explanation

Volume *volume* is being copied with SRDF.

Resolution

Execute the command again after the copying process is complete.

60279

ERROR: volume: related to same master volume as proxy proxy, master=master

Explanation

Volume *volume* is related to the same master volume *master* as proxy volume *proxy*.

Resolution

Specify a correct volume name. For details, see "[Appendix D Command Reference](#)."

60280

ERROR: master and proxy exist in same group group

Explanation

The specified master volume and proxy volume exist in the same group *group*.

Resolution

Specify volumes that exist in different groups.

60281

ERROR: proxy: joined to master with EC, rejoin them by soft copy and try again

Explanation

Slices cannot be swapped since EC session in process between proxy *proxy* and master.

Resolution

Cancel the EC session by either of the following operations according to need and try this command again.

- Part the master and the proxy once and rejoin them with the `sdxproxy Rejoin -e softcopy` command.
- Break the relationship between the master and the proxy once and rejoin them with the `sdxproxy Join -e softcopy` command.

60282

```
ERROR: proxy: joined to master with TimeFinder, rejoin them by soft copy and try again
```

Explanation

Slices cannot be swapped because there is a BCV pair between the proxy *proxy* and the master.

Resolution

Cancel the BCV pair by either of the following operations according to need and try this command again.

- Part the master and the proxy once and rejoin them with the `sdxproxy Rejoin -e softcopy` command.
- Break the relationship between the master and the proxy once and rejoin them with the `sdxproxy Join -e softcopy` command.

60283

```
ERROR: proxy: joined to master with SRDF, rejoin them by soft copy and try again
```

Explanation

Slices cannot be swapped because there is an SRDF pair between the proxy *proxy* and the master.

Resolution

Cancel the SRDF pair by either of the following operations according to need and try this command again.

- Part the master and the proxy once and rejoin them with the `sdxproxy Rejoin -e softcopy` command.
- Break the relationship between the master and the proxy once and rejoin them with the `sdxproxy Join -e softcopy` command.

60284

```
ERROR: volume: proxy volume cannot be specified when using TimeFinder
```

Explanation

You cannot specify *volume* when you perform parting, rejoining or restoring because there is a BCV pair between the proxy volume *volume* and the master.

Resolution

If you wish to perform parting, rejoining, or restoring, you must specify the group to which *volume* belongs when executing the command.

60285

ERROR: volume: proxy volume cannot be specified when using SRDF

Explanation

You cannot specify *volume* when you perform parting, rejoining or restoring because there is an SRDF pair between the proxy volume *volume* and the master.

Resolution

If you wish to perform parting, rejoining, or restoring, you must specify the group to which *volume* belongs when executing the command.

60286

ERROR: proxy: failed to start soft copy

Explanation

An error occurred when synchronization copying with the soft copy function was started between the proxy volume *proxy* and the master volume.

Resolution

Collect investigation material and contact your local customer support.

60288

ERROR: OPC not available

Explanation

The OPC function is unavailable.

Resolution

See "(1) The Advanced Copy function cannot be used in master-proxy copying." in "F.1.7 Proxy Object Abnormality."

60289

ERROR: EC not available

Explanation

The EC function is unavailable.

Resolution

See "(1) The Advanced Copy function cannot be used in master-proxy copying." in "F.1.7 Proxy Object Abnormality."

60290

ERROR: proxy: too many EC/OPC sessions

Explanation

The number of EC or OPC sessions within the physical disk (LU) or the disk array body has reached the upper limit of supported concurrent sessions. For this reason, a new EC or OPC session cannot be started.

Resolution

To make copying by EC or OPC available, wait until the running session ends, and try this command again. Alternatively, according to need, cancel the running session using the `sdxproxy Cancel` command, the `sdxproxy Break` command, or `[Operation]:[Proxy Operation]:[Break]` and try this command again.

60291

```
ERROR: proxy: offset is different from master volume master
```

Explanation

Master volume *master* and proxy volume *proxy* have different top block (sector) numbers. The top block number is not a physical block number that indicates the offset on a physical disk, but is a logical block number that indicates the offset within a group (or a single disk) to which the volume belongs. The layout (offsets and sizes) of volumes within a master group and a proxy group must be consistent.

Resolution

For the proxy group, choose a group in which the layout of volumes are consistent with those of the master group. For the top block (sector) numbers and the sizes of volumes, use the `sdxinfo` command and check the `1STBLK` field and the `BLOCKS` field of the displayed volume information.

60292

```
ERROR: proxy: number of volumes is different from master group master
```

Explanation

Master group *master* and proxy group *proxy* include different numbers of volumes. The layout (offsets and sizes) of volumes within a master group and a proxy group must be consistent.

Resolution

For the proxy group, choose a group in which the layout of volumes is consistent with those of the master group.

60293

```
ERROR: proxy: failed to start OPC, source=disk.volume, target=disk.volume, class=class
```

Explanation

An error occurred between the proxy volume *proxy* and the master volume when OPC started.

"source" indicates the copy source slice name, "target" indicates the copy destination slice name, and "class" indicates the class to which the copy source and the copy destination slices belong.

Resolution

Identify the cause by referring to GDS log messages, disk driver log messages, and Fujitsu Storage Systems ETERNUS log messages that were output right before the error occurrence, and recover the status.

60294

ERROR: proxy: failed to start EC session, source=disk.volume, target=disk.volume, class=class

Explanation

An error occurred between the proxy volume *proxy* and the master volume when the EC session started.

"source" indicates the copy source slice name, "target" indicates the copy destination slice name, and "class" indicates the class to which the copy source and the copy destination slices belong.

Resolution

Identify the cause by referring to GDS log messages, disk driver log messages, and Fujitsu Storage Systems ETERNUS log messages that were output right before the error occurrence, and recover the status.

60295

ERROR: copy region exceeds volume region, volume=volume, offset=blkno, size=size

Explanation

The region specified with the start block number *blkno* and the size *size* exceeds the region of the volume *volume*.

Resolution

Check the size of the volume *volume* using the `sdxinfo` command, and modify the copy region file.

60296

ERROR: cannot open copy region file region_file, errno=errno

Explanation

Opening the copy region file *region_file* failed.

Resolution

Collect investigation material and contact your local customer support.

60297

ERROR: cannot read copy region file region_file, errno=errno

Explanation

Reading the copy region file *region_file* failed.

Resolution

Collect investigation material and contact your local customer support.

60298

ERROR: syntax error in copy region file region_file, line=line

Explanation

The copy region file *region_file* contains an error at the line indicated by *line*.

Resolution

Check contents of *region_file*.

60299

ERROR: too many copy regions in copy region file *region_file*

Explanation

The copy region file *region_file* specifies too many copy regions.

Resolution

Modify *region_file*.

60300

ERROR: copy region file *region_file* contains overlapping regions, line=*line1*, line=*line2*

Explanation

Copy regions file *region_file* specified in the lines indicated by *line1* and *line2* overlap.

Resolution

Check contents of *region_file*.

60301

ERROR: OPC not available with multiple extents

Explanation

The disk unit may not support the OPC copy function with multiple extents or the function may not currently be available for some reason.

Resolution

Check the disk unit's hardware configuration. According to need, specify one copy region at a time and execute the command for every region.

60303

ERROR: cannot get configuration information, *sdxinfo(1)* command failed

Explanation

The *sdxinfo(1)* command failed, and GDS configuration information could not be obtained.

Resolution

Remove the error after identifying the cause by referring to the message for the *sdxinfo(1)* command that was output right before the occurrence of the error, and try the *sdxproxy Root* command again.

60304

ERROR: volume: alternative volume *altvol* already specified for volume *curvol*

Explanation

Another alternative volume *volume* was assigned to the volume *curvol* for which the alternative volume *altvol* has already been specified.

Resolution

You can specify only one alternative volume for a volume.

60305

ERROR: volume: corresponding volume not specified in /etc/vfstab

Explanation

The volume *volume* is not the alternative volume of a volume described in /etc/vfstab as a file system or a swap area.

Resolution

Specify a volume that fits one of the following requirements.

- Proxy volume of a master volume described in the /etc/vfstab file
- Master volume of a proxy volume described in the /etc/vfstab file
- Another proxy volume related to the master volume of a proxy volume described in the /etc/vfstab file

60306

ERROR: volume: read only volume

Explanation

The access mode of the volume *volume* is ro (read only). Since a read-only volume cannot be mounted, this cannot be configured for an alternative boot environment.

Resolution

Change the access mode attribute of the volume *volume* to rw (read and write) using the `sdxattr -V` command, and try this command again.

60307

ERROR: alternative root volume must be specified

Explanation

The alternative root volume has not been specified.

Resolution

You must specify the alternative root volume.

60308

ERROR: volume: file system cannot be checked or repaired, fsck(1M) command failed, exit-status=exitstat

Explanation

The fsck(1M) command for the volume *volume* failed, and a file system on *volume* could not be inspected or repaired. The exit status of the fsck(1M) command is *exitstat*. Inconsistency may have arisen in the file system on *volume*.

Resolution

Remove the error after identifying the cause by referring to the message for the fsck(1M) command that was output right before the occurrence of the error and to the manual for the fsck(1M) command. If required, restore the data on *volume*, for example, using the backed up data, and try the sdxproxy Root command again.

60309

ERROR: volume: file system cannot be mounted, mount(1M) command failed, exit-status=exitstat
details

Explanation

The mount(1M) command for the volume *volume* failed, and the file system on *volume* could not be mounted. The exit status of the mount(1M) command is *exitstat*. *details* is the error message for the mount(1M) command.

Resolution

Remove the error after identifying the cause based on *details*, and try the sdxproxy Root command again.

60310

ERROR: volume: cannot change boot-device parameter, eeprom(1M) command failed, exit-status=exitstat
details

Explanation

The eeprom(1M) command failed, and the boot-device parameter could not be changed to the device for the volume *volume*. The exit status of the eeprom(1M) command is *exitstat*. *details* is the error message for the eeprom(1M) command.

Resolution

Remove the error after identifying the cause based on *details*, and try the sdxproxy Root command again.

60312

ERROR: volume: file system cannot be unmounted, umount(1M) command failed, exit-status=exitstat
details

Explanation

The umount(1M) command for the volume *volume* failed, and the file system on *volume* temporarily mounted could not be unmounted. The exit status of the umount(1M) command is *exitstat*. *details* is the error message for the umount(1M) command.

Resolution

Remove the error after identifying the cause based on *details*, unmount *volume*, and try the `sdxproxy Root` command again.

60313

ERROR: class: not a shadow class

Explanation

The class *class* is not a shadow class.

Resolution

See "[Appendix D Command Reference](#)," and specify a proper command name and class name.

60314

ERROR: device: no configuration information

Explanation

Since no configuration information resides in the private slice on the physical disk *device* or no private slice exists on *device*, *device* cannot be registered with a shadow class.

The possible causes are as follows.

- a) *device* is not registered with a class in another domain yet, and the private slice has not been copied from the SDX disk to *device* with the disk unit's copy function. In addition, the disk is not a disk previously removed from a class with the `sdxconfig Remove -e keepid` command.
- b) *device* is already registered with a class in another domain, but has not been enabled.
- c) The private slice has been copied from the SDX disk to *device* with the disk unit's copy function, but the copy source SDX disk has not been enabled.

Resolution

Check on the system configuration and so on, and identify which cause among a), b) and c) applies. If the cause is a), see "[Appendix D Command Reference](#)" and specify a correct command name and physical disk name. If it is b), restore *device* in another domain. If it is c), restore the copy source SDX disk.

60315

ERROR: device: registered with illegal class in another domain

Explanation

While the physical disk *device* is already registered with a different class in another domain, registering it with the same shadow class was attempted.

Resolution

You should register a disk registered with the same class in another domain with one shadow class.

60316

ERROR: device: disk: not same as disk name *diskname* in another domain

Explanation

The physical disk *device* is already registered as *diskname* with a class in another domain. It cannot be registered with a shadow class as *disk* that is another disk name.

Resolution

Specify the same disk name *diskname* as in another domain.

60317

```
ERROR: device: private slice size not same as another disk in class
```

Explanation

Since the physical disk *device* has the private slice that is unequal in size compared to the private slices on other disks registered with the shadow class *class*, it cannot be registered with *class*.

Resolution

Check on the system configuration and so on, and specify a correct disk name and shadow class name.

60319

```
ERROR: no license
```

Explanation

The command cannot be executed. The possible causes are as follows.

- a) In a system where the FJSVsdxl package of PRIMECLUSTER GDS is not installed normally, the `sdxconfig` command was used.
- b) In a system where the FJSVsdxsl package of PRIMECLUSTER GDS Snapshot is not installed normally, the `sdxshadowdisk` command was used.

Resolution

To use the `sdxconfig` command, install FJSVsdxl normally. To use the `sdxshadowdisk` command, install FJSVsdxsl normally.

60320

```
ERROR: output file already exists
```

Explanation

The specified output file is an existing file.

Resolution

Specify a nonexistent file name. To overwrite an existing file, use the `-e` update option.

60321

```
ERROR: failed to create configuration file
```

Explanation

Configuration file creation failed.

Resolution

Check whether the specified path to the configuration file is correct.

60322

ERROR: class: failed to get configuration information

Explanation

Acquisition of *class* configuration information failed.

Resolution

Collect investigation material and contact your local customer support.

60323

ERROR: proxy: proxy volume exists in class

Explanation

Proxy volume *proxy* exists in *class*.

Resolution

Break the relationship of proxy volume *proxy* to the master as needed.

60324

ERROR: group: switch group exists in class

Explanation

Class *class* includes switch group *group*. The attempted operation is not supported for classes that include switch groups.

Resolution

Delete the *group* as needed and re-execute the command.

60325

ERROR: failed to output configuration table

Explanation

Sending the configuration table to standard output or to a configuration file failed.

Resolution

Collect investigation material and contact your local customer support.

60326

ERROR: input file not found

Explanation

The specified input file does not exist.

Resolution

Specify a correct file name.

60327

ERROR: *class* : not same as class name *name* in configuration table

Explanation

The specified class name *class* is different from the class name *name* in the configuration table.

Resolution

Specify a correct class name, or change the class name in the configuration table to *class* with the `sdxcnfig Convert` command.

60328

ERROR: *disk* : no such disk in configuration table

Explanation

The configuration table does not contain disk *disk*.

Resolution

Specify a correct disk name.

60329

ERROR: *device* : no such device in configuration table

Explanation

The configuration table does not contain physical disk *device*.

Resolution

Specify a correct physical disk name.

60330

ERROR: *group* : no such group in configuration table

Explanation

The configuration table does not contain group *group*.

Resolution

Specify a correct group name.

60331

ERROR: *object* : exists in *type* group in configuration table

Explanation

In the configuration table *object* that is a disk or a lower level group is connected to a *type* type group. Disks connected to concat or stripe type groups and lower level groups connected to stripe type groups cannot be removed from configuration tables.

Resolution

Specify a correct disk name or group name.

60332

ERROR: at least one object must remain in configuration table

Explanation

Removing the specified objects from the configuration table will result in no object in the configuration table. A configuration table needs to contain a minimum of one object.

Resolution

Specify correct object names.

60333

ERROR: *object.volume* is only valid slice in configuration table

Explanation

In the configuration table slice *object.volume* is an only valid slice that comprises mirror volume *volume*. For this reason, *object* that is a disk or a lower level group cannot be removed from the configuration table.

Resolution

Specify a correct disk name or group name.

60334

ERROR: *class* : not a local class

Explanation

Class *class* is not a local class.

Resolution

Specify a local class.

60335

ERROR: file name too long

Explanation

The specified file name contains too many characters.

Resolution

Specify a correct file name.

60336

ERROR: failed to open input file, errno=errno

Explanation

Input file open failed.

Resolution

Identify the cause based on the error number *errno*.

60337

ERROR: configuration table corrupted, sdxfunc=sdxfunc, sdxerrno=sdxerrno

Explanation

Contents of the configuration table are invalid.

Resolution

Collect investigation material and contact your local customer support.

60338

ERROR: class : already exists

Explanation

Class *class* already exists.

Resolution

Change the class name in the configuration table with the sdxconfig Convert command and re-execute the command.

60339

ERROR: class : already exists in another node

Explanation

Class *class* already exists in another node.

Resolution

Change the class name in the configuration file with the sdxconfig Convert command and re-execute the command.

60340

ERROR: *device* : assigned to *disk1* and *disk2* in configuration table

Explanation

In the configuration table one physical disk *device* is assigned to disks *disk1* and *disk2*.

Resolution

Change the physical disk assigned to *disk1* or *disk2* in the configuration table with the `sdxconfig Convert` command and re-execute the command.

60341

ERROR: *device* : failed to get physical disk information

Explanation

Acquisition of geometry information, VTOC, or the device number of physical disk *device* failed.

Resolution

Check whether physical disk *device* is normally operating.

60342

ERROR: *device* : size must be *size* blocks

Explanation

The physical disk *device* must have *size* blocks.

Resolution

Replace physical disk *device* with a physical disk that has *size* blocks. Alternatively, change physical disk *device* in the configuration table to another physical disk that has *size* blocks with the `sdxconfig Convert` command.

60343

ERROR: *device* : private slice size must be *size* blocks

Explanation

The size of the private slice of physical disk *device* is different from that described in the configuration table and therefore the class object configuration cannot be restored. The private slice (slice 0) of *device* must have *size* blocks.

Resolution

Exchange physical disk *device* for a physical disk whose private slice (slice 0) size is *size* blocks. Alternatively, change physical disk *device* in the configuration table to another physical disk whose private slice size is *size* blocks with the `sdxconfig Convert` command.

60344

ERROR: mismatch of class names in private slices on *device1* and *device2*

Explanation

The class names of physical disks *device1* and *device2* stored in the private slices do not match and therefore *device1* and *device2* cannot be registered with one class.

Resolution

Check the system configuration and edit the physical disks described in the configuration table with the `sdxconfig Convert` command.

60345

```
ERROR: device : mismatch of disk names, disk1 in private slice, disk2 in configuration table
```

Explanation

The disk name *disk1* stored in the private slice does not match the disk name *disk2* described in the configuration table and therefore physical disk *device* cannot be registered as disk *disk2* with a class.

Resolution

Check the system configuration and edit the physical disks described in the configuration table with the `sdxconfig Convert` command.

60346

```
ERROR: class : restoration based on configuration file failed
```

Explanation

Restoration of the *class* configuration based on configuration file *file* failed.

Resolution

Collect investigation material and contact your local customer support.

60347

```
ERROR: class: some node trying to get class master privilege
```

Explanation

The current node or some other node is trying to obtain master privileges of shared class *class*.

Resolution

If necessary, try the command again after a while.

60348

```
ERROR: class: class master not found
```

Explanation

The class master of shared class *class* is not found and no operation is possible for *class*.

Resolution

If necessary, try the command again after a while.

60349

```
ERROR: failed to update NVRAM configuration parameter,  
parameter='param'
```

Explanation

Updating the parameter *param* value in device's NVRAM (nonvolatile memory) failed.

Resolution

Collect investigation material and contact your local customer support.

60368

```
ERROR: device: device is MPxIO disabled
```

Explanation

MPxIO is disabled for the physical disk *device*.

Resolution

MPxIO should be enabled for the *device*. Register the disk with the class after enabling MPxIO.

60369

```
ERROR: disk: cannot be operated by sdxswap command
```

Explanation

disk cannot be operated by the sdxswap command.

Resolution

Remove *disk* from the class, change the disk unit, and re-register the new disk with the class.

60372

```
ERROR: device: failed to get information of physical device
```

Explanation

Failed to get the information of the specified physical disk.

Resolution

See if the multipath configuration of the physical disk is correct. If you still cannot determine the cause of the error, collect investigation material and contact your local customer support.

60418

```
ERROR: class: same class id assigned to another class in scope
```

Explanation

The same class ID as the class ID of the class *class* is used in another class on the node to be added to the scope.

Resolution

Take the following steps to remove the class *class*, and after recovering the status, change the scope.

1. Back up the object configuration information of the class in which a node is added to the scope. (sdxcfg Backup command)
2. Stop the volume of the class in which a node is added to the scope. (sdxvolume -F command)
3. For the shared class, change the class in which a node is added to the scope to the local class. (sdxattr -C command)
4. Remove the object configuration information of the class in which a node is added to the scope. (sdxcfg Remove command)
5. Restore the object configuration information of the class in which a node is added to the scope. (sdxcfg Restore command)
6. Add a node to the scope in the restored local class. (sdxattr -C command)

E.4.2 Warning Messages (62000-62099)

62000

```
WARNING: spare disk disk too small
```

Explanation

The size of *disk* is too small and may not function as a spare disk.

Resolution

Specify a larger disk, and execute the command again.

62001

```
WARNING: device: write error, errno=errno
```

Explanation

A write error occurred in physical disk *device*.

Resolution

If *device* is a write-locked disk that is connected to a switch group, no action is required. Otherwise, remove *device* from the class and check the status of *device*.

62002

```
WARNING: group: free blocks are reduced
```

Explanation

Free blocks on *group* were reduced.

Resolution

You may not be able to create a volume with sufficient capacity. Execute commands as needed and attempt recovery.

62003

WARNING: another disk must be connected to group

Explanation

You must connect another disk to *group*.

Resolution

Connect another disk.

62004

WARNING: object: copying not completed successfully

Explanation

Synchronization copying did not complete successfully.

Resolution

Disk failure may have occurred. Identify the cause by referring to GDS log message and syslog message.

62005

WARNING: object: gave up wait for the completion of copying by a cancel request

Explanation

Synchronization copying was canceled before completion.

Resolution

Check the status of the object *object*. If the copying is in progress, there is no need to work around. If it is not in progress, re-execute the copying where it is necessary.

For details, see "[Appendix D Command Reference](#)."

62006

WARNING: cannot set use-nvramrc? parameter to true, NVRAMRC not empty

Explanation

Since a command is already registered with NVRAMRC, use-nvramrc? parameter cannot be set to "true."

Resolution

For details, see "[A.2.9 System Disk Mirroring](#)."

62007

WARNING: group: no spare disk available

Explanation

There is no valid spare disk in *group*.

Resolution

Define a spare disk as needed.

62008

WARNING: *object.volume*: cannot attached due to in *status* status

Explanation

Slice *object.volume* could not be attached since it is in *status* status.

Resolution

Check the slice status and cancel *status* status as needed.

62009

WARNING: *volume*: no need to resize volume

Explanation

There was no need to resize the volume *volume*.

Resolution

No particular resolution required.

62010

WARNING: *disk.volume*: special file(s) not found

Explanation

Special file for slice *disk.volume* could not be found.

Resolution

No particular resolution required.

62011

WARNING: *node*: node in stopped status

Explanation

node is in STOP status.

Resolution

No particular resolution required. However, promptly activating the node is recommended.

62012

WARNING: node: node in abnormal status

Explanation

node is in abnormal status.

Resolution

No particular resolution required. However, promptly recovering the node and activating it normally is recommended.

62013

WARNING: object: already in status status

Explanation

object is already in *status* status.

Resolution

No particular resolution required.

62014

WARNING: disk : device : write error, errno=errno

Explanation

A write error occurred on physical disk *device*.

Resolution

If *device* is a write-locked disk that is connected to a switch group, no action is required. In other situations, have *device* exchangeable and check its status.

62015

WARNING: disk : device : node: write error, errno=errno

Explanation

A write error occurred on physical disk *device* on node *node*.

Resolution

If *device* is a write-locked disk that is connected to a switch group, no action is required. In other situations, have *device* exchangeable and check its status.

62016

WARNING: object: no need to update attribute value

Explanation

There was no need to change *object* attributes.

Resolution

No particular resolution required.

62017

WARNING: *pslice*: entry in /etc/vfstab not updated, unsupported file system type *fstype*

Explanation

Entry in the /etc/vfstab file relevant to the physical slice *pslice* was not updated since *fstype* is a file system type not supported by the `sdxroot` command.

Resolution

Edit the /etc/vfstab file or change the file system configuration according to need.
If *fstype* is `sfxfs` or `sfvfs`, see "[A.2.11 GFS Partition on a Keep Disk](#)."

62018

WARNING: correct /etc/vfstab and file system configuration before rebooting, or system may not be booted

Explanation

Invalid entry is included in the /etc/vfstab file. If such file system entry as /(/root), /usr, and /var essential to system start-up is invalid, the system cannot be started when rebooted.

Resolution

Refer to a GDS's WARNING message output just before this message and pinpoint the invalid entry in the /etc/vfstab file. Be sure to repair the /etc/vfstab file and the file system configuration as appropriate prior to rebooting the system.
If entry in a GFS file system is invalid, see "[A.2.11 GFS Partition on a Keep Disk](#)."

62019

WARNING: *volume*: entry in /etc/vfstab not updated, unsupported file system type *fstype*

Explanation

Since file system type *fstype* is not supported, entry in /etc/vfstab file for *volume* could not be updated.

Resolution

Directly edit /etc/vfstab as necessary.

62020

WARNING: ignored parameter *param* for EC

Explanation

Since EC function will be used for copying, parameter *param* was ignored.

62021

WARNING: ignored parameter *param* for OPC

Explanation

Since OPC function will be used for copying, parameter *param* was ignored.

62022

WARNING: ignored parameter *param* for TimeFinder

Explanation

Parameter *param* was ignored for copying with TimeFinder.

62023

WARNING: ignored parameter *param* for SRDF

Explanation

Parameter *param* was ignored for copying with SRDF.

62024

WARNING: *proxy*: no session

Explanation

No copy session exists between the proxy *proxy* and the master

Resolution

No special action is required.

E.4.3 Information Messages (64000-64099)

64000

INFO: waiting for a response from sdxservd daemon...

Explanation

Awaiting response from sdxservd daemon.

64001

INFO: *class*: created class

Explanation

class was created.

64002

INFO: *disk*: created *disk*

Explanation

disk was registered.

64003

INFO: *device*: disabled access to physical special files

*/dev/rdisk/devices**

*/dev/dsk/devices**

Explanation

Physical special files can no longer be accessed.

64004

INFO: *class*: removed *class*

Explanation

class was removed.

64005

INFO: *group*: created *group*

Explanation

group was created.

64006

INFO: *disk*: connected *disk* to *group group*

Explanation

disk was connected to *group*.

64007

INFO: *lgroup*: connected *group* to another *group hgroup*

Explanation

Group *lgroup* was connected to another group *hgroup*.

64008

INFO: *group*: removed group

Explanation

group was removed.

64009

INFO: *object*: waiting for the completion of copying...

Explanation

Waiting for synchronization copying to complete.

64010

INFO: *object*: copying completed successfully

Explanation

Synchronization copying completed successfully.

64011

INFO: *volume*: created volume

Explanation

volume was created.

64012

INFO: *volume*: started volume on node *node*

/dev/sfdsk/class/rdsk/volume

/dev/sfdsk/class/dsk/volume

Explanation

volume was started on *node*. You can now access via special file.

64013

INFO: *volume*: stopped volume on node *node*

Explanation

volume was stopped on *node*.

64014

INFO: *volume*: removed volume

Explanation

volume was removed.

64015

INFO: *volume*: resized volume

Explanation

The volume *volume* was resized.

64016

INFO: *disk*: disconnected disk from group *group*

Explanation

disk was disconnected from *group*.

64017

INFO: *lgroup*: disconnected group from group *hgroup*

Explanation

Group *lgroup* was disconnected from *hgroup*.

64018

INFO: *group*: free blocks are increased

Explanation

Free blocks on *group* increased.

64019

INFO: *disk*: removed disk

Explanation

disk was removed from class.

64020

INFO: *device*: enabled access to physical special files

*/dev/rdisk/devices**

*/dev/dsk/devices**

Explanation

You can now access to physical special files.

64021

INFO: *disk.volume*: temporarily detached slice from volume *volume*

Explanation

Slice *disk.volume* was temporarily disconnected from *volume*.

64022

INFO: *disk.volume*: enabled slice

/dev/sfdsk/class/rdsk/disk.volume

/dev/sfdsk/class/dsk/disk.volume

Explanation

The slice *disk.volume* can now be accessed independently of the volume. Use the special file displayed.

64023

INFO: *disk.volume*: disabled slice

Explanation

Slice *disk.volume* cannot be accessed any more.

64024

INFO: *disk.volume*: reattached slice to volume *volume*

Explanation

Slice *disk.volume* is attached to *volume* again.

64025

INFO: *disk: device*: ready for swapping

Explanation

The disk *disk* is now exchangeable.

For a file unit device supporting hot swap, you may demount the physical disk *device* at this point for swapping disks.

If active swap is not supported, shut down the system and turn off the power before removing the physical disk.

64026

INFO: *disk: device: ready for swapping on node node*

Explanation

On the node *node*, the physical disk *device* of the disk *disk* is now exchangeable.

If the disk unit supports hot swap, you may simply demount *device* connected to *node* and swap disks.

If it does not support hot swap, shot down *node* and turn off the power, then demount *device* and swap disks.

64027

INFO: *disk: restored disk*

Explanation

disk restored its original status.

64028

INFO: *class: class recovery completed successfully*

Explanation

Recovery of the class *class* from the closed status succeeded.

64029

INFO: *volume: recovered volume with disk disk*

Explanation

volume was recovered based on *disk*.

64030

INFO: *volume: recovered volume with lower level group lgroup*

Explanation

Volume *volume* was recovered based on lower level group *lgroup*.

64031

INFO: *volume : recovered volume in group group*

Explanation

Volume *volume* in group *group* has been recovered.

64032

INFO: *disk*: recovered disk

Explanation

disk *disk* was recovered.

64033

INFO: no unsynchronized slice in *volume*

Explanation

There is no slice that needs synchronization copying to be started, or resumed in *volume*.

64034

INFO: *volume*: copying canceled

Explanation

Synchronization copying of *volume* has been cancelled.

64035

INFO: no copying slice exists in *volume*

Explanation

There is no slice in the process of copying in *volume*.

64036

INFO: *volume*: copying interrupted

Explanation

Synchronization copying on *volume* was canceled.

64037

INFO: *class*: updated attribute value of class

Explanation

Attribute value of *class* was updated.

64039

INFO: *disk*: updated attribute value of *disk*

Explanation

Attribute value of *disk* was updated.

64040

INFO: *group*: updated attribute value of *group*

Explanation

Attribute value of *group* was updated.

64041

INFO: *volume*: updated attribute value of *volume*

Explanation

Attribute value of *volume* was updated.

64042

INFO: *slice*: updated attribute value of *slice*

Explanation

Attribute value of *slice* was updated.

64043

INFO: updated */etc/vfstab* file

Explanation

/etc/vfstab file has been updated.

64044

INFO: updated */kernel/drv/sfdsk.conf* file

Explanation

/kernel/drv/sfdsk.conf file has been updated.

64046

INFO: updated */etc/system* file

Explanation

/etc/system file has been updated.

64048

INFO: completed definitions of system disk(s) mirroring

Explanation

Definition for system disk mirroring is complete.

64049

INFO: system must be rebooted:

```
shutdown -y -g0 -i6
```

Explanation

Reboot the system by executing the following command.

```
shutdown -y -g0 -i6
```

64050

INFO: completed definitions of system disk(s) mirroring cancellation

Explanation

Canceling definition for system disk mirroring is complete.

64051

INFO: disk.volume: started slice on node node

```
/dev/sfdsk/class/rdisk/disk.volume
```

```
/dev/sfdsk/class/dsk/disk.volume
```

Explanation

The slice *disk.volume* was activated on *node*. You may access the slice by using the displayed special file.

64052

INFO: disk.volume: stopped slice on node node

Explanation

Slice *disk.volume* was stopped on *node*. You cannot access slice until it is activated again.

64053

INFO: *disk.volume*: disabled slice on remote node *node*

Explanation

Slice *disk.volume* has become unavailable on remote node *node*.

64054

INFO: *disk.volume*: current node *node* took over and started slice

/dev/sfdsk/class/rdisk/disk.volume

/dev/sfdsk/class/dsk/disk.volume

Explanation

The current node *node* took over the slice *disk.volume* and activated it. You may access the slice by using the displayed special file.

64055

INFO: *volume*: copying parameter modification completed

Explanation

Modification of parameter related to the copying process of *volume* is complete.

64056

INFO: configuration parameters modification completed

Explanation

Modification of configuration parameter is complete.

64057

INFO: *proxy*: related and joined to master *master*

Explanation

Proxy *proxy* has been related and joined to master *master*.

64058

INFO: *proxy*: created proxy volume

Explanation

Proxy volume *proxy* has been created.

64059

INFO: *proxy*: broken relation with master

Explanation

Relationship between master and proxy *proxy* has been cancelled.

64060

INFO: *proxy*: forced to break relation with master

Explanation

Relationship between master and proxy *proxy* was forcibly cancelled.

64061

INFO: *proxy*: must be recovered by `sdxfix(1M)` command

Explanation

In order to restore data of proxy *proxy*, volume must be restored with `sdxfix(1M)` command.

64062

INFO: *volume*: started volume on node *node*,...

`/dev/sfdsk/class/rdsk/volume`

`/dev/sfdsk/class/dsk/volume`

Explanation

The volume *volume* started on the nodes *node*,.... It can be accessed with the special file.

64063

INFO: *proxy*: parted from master

Explanation

Proxy *proxy* was parted from master.

64064

INFO: *proxy*: parted from master instantly

Explanation

Proxy *proxy* was immediately parted from master.

64065

INFO: *proxy*: rejoined to master *master*

Explanation

Proxy *proxy* has been rejoined to master *master*.

64066

INFO: *master*: restored data from proxy *proxy*

Explanation

Copying from proxy *proxy* to master *master* started.

64067

INFO: *master*: restored data from proxy *proxy* instantly

Explanation

Copying from proxy *proxy* to master *master* started. Even if the copying is in process, you may start master volumes and access valid data.

64068

INFO: swapped slice components between master and proxy

Explanation

Slice comprising proxy and slice comprising master have been swapped.

64069

INFO: *proxy*: session canceled

Explanation

A session of the disk unit's copy function that resided between proxy *proxy* and the master was canceled.

64070

INFO: *proxy*: related to master *master*

Explanation

Master *master* and proxy *proxy* were related and parted.

64071

INFO: *proxy*: updated data with master *master*

Explanation

Data of proxy *proxy* was updated with data of master *master*.

64072

INFO: proxy: updated data with master master instantly

Explanation

Data of proxy *proxy* was instantly updated with data of master *master*.

64073

INFO: proxy: specified regions updated

Explanation

Data in the specified regions on the proxy *proxy* was updated.

64074

INFO: master: specified regions restored from proxy proxy

Explanation

Based on the proxy *proxy*, data in the specified regions on the master *master* was restored.

64075

INFO: volume: executing fsck(1M) command:

details

Explanation

The file system on the alternative root volume *volume* will be inspected and repaired with the fsck(1M) command. *details* is the message for the fsck(1M) command.

64076

INFO: updated /etc/vfstab file on alternative root volume volume

Explanation

The /etc/vfstab file on the alternative root volume *volume* was updated.

64077

INFO: volume will be used as dump device in alternative boot environment

Explanation

The volume *volume* is used as a dump device in an alternative boot environment. *volume* is made available as a dump device when the alternative boot environment is enabled.

64078

INFO: updated /etc/system file on alternative root volume *volume*

Explanation

The /etc/system file on the alternative root volume *volume* was updated.

64080

INFO: updated boot-device parameter:

```
old=old-boot-device ...
```

```
new=new-boot-device ...
```

Explanation

The boot-device parameter *old-boot-device* ... was replaced with *new-boot-device*

Be sure to take a copy of *old-boot-device* ... and *new-boot-device* ...

To cancel the switching to the alternative boot environment, execute the following command.

```
# eeprom boot-device="old-boot-device ..."
```

Alternatively, execute the following command in the OpenBoot environment.

```
ok setenv boot-device old-boot-device ...
```

In general, the number of boot device names for *old-boot-device* ... and *new-boot-device* ... corresponds to the number of disks in the current boot environment and the alternative boot environment respectively. However, GDS may not recognize device names depending on some controller types. If that is so, two device names for one boot disk are displayed, but it has no effect on the operation.

64082

INFO: completed definitions of switching boot environment

Explanation

The boot environment switching is now defined.

64083

INFO: completed definitions of alternative boot environment:

```
current-boot-device=current-boot-device ...
```

```
alternative-boot-device=alternative-boot-device ...
```

Explanation

The alternative boot environment is now defined.

The boot device names in the current boot environment are *current-boot-device ...* and those in the alternative boot environment are *alternative-boot-device ...*

Be sure to take a copy of *current-boot-device ...* and *alternative-boot-device ...*

Executing the following command in the OpenBoot environment can switch to the alternative boot environment.

```
ok setenv boot-device alternative-boot-device ...
```

In general, the number of boot device names for *current-boot-device ...* and *alternative-boot-device ...* corresponds to the number of disks in the current boot environment and the alternative boot environment respectively. However, GDS may not recognize device names depending on some controller types. If that is so, two device names for one boot disk are displayed, but it has no effect on the operation.

64088

```
INFO: class: configuration backup created  
file
```

Explanation

The object configuration of class *class* was backed up in object table format to configuration file *file*.

64089

```
INFO: configuration table converted successfully  
file
```

Explanation

The configuration table was converted successfully. *file* indicates a configuration file to which the configuration table was output.

64090

```
INFO: device: disk ID information not removed
```

Explanation

Physical disk *device* retains unremoved disk identification information.

64091

```
INFO: class: waiting for restoration...
```

Explanation

Waiting for restoration of the object configuration of class *class*.

64092

```
INFO: class: restored based on configuration file  
file
```

Explanation

The *class* configuration was restored based on configuration file *file*.

64093

```
INFO: class: current node node got class master privilege
```

Explanation

Current node *node* obtained master privileges of shared class *class*.

64094

```
INFO: NVRAM configuration parameter has been updated,  
parameter='param'
```

Explanation

The parameter *param* value in device's NVRAM (nonvolatile memory) was updated.

64095

```
INFO: no need to update NVRAM configuration parameter,  
parameter='param'
```

Explanation

The parameter *param* value in device's NVRAM (nonvolatile memory) was not updated.

E.4.4 Fix Messages (66000)

66000

```
TO FIX: Usage: usage
```

Resolution

Follow the *usage* syntax and execute the command.

E.4.5 Internal Error Messages (68000)

68000

```
command: severity: module: internal error, details
```


Explanation

An internal error occurred.

details gives you the cause of the error, *command* the command name, *severity* the severity of message, and *module* the module name (usually the internal function name) where the event was detected.

The following message "a" may be output if the `sdxcopy -C` command is executed at node startup or during GDS command execution. If the `sdxcopy -C` command has ended normally (the exit status is 0), there is no problem.

a) `sdxcopy: ERROR: SDX_COPY_STATUS_WAIT2: internal error, sdxerrno=145`

Resolution

In the event of "a", if the exit status of the `sdxcopy -C` command is 0, the message "a" may be ignored. For the `sdxcopy -C` command's exit status, refer to the following message recorded in the `sdxservd` daemon's log file `/var/opt/FJVSdx/msglog/sdxservd.log`.

```
SDX:sdxservd: INFO: command exited, exit-status=exitstat:  
sdxcopy -C -c class -v volume,...
```

In other events, collect investigation material and contact your local customer support.

E.5 Operation Management View Messages

This section describes the GDS Management View messages.

E.5.1 Error Messages (5000-5099)

5000

Command line output: *details*

Explanation

An error given in *details* occurred when executing the command.

Resolution

Identify the cause referring to the command line output. When you cannot identify the cause, contact your local customer support.

5001

Executing the command was unsuccessful. Continue?

Explanation

Executing the command was unsuccessful.

Resolution

Click <Yes> to continue, or <No> to cancel.

5002

The following exceptions were received: *details*

Explanation

Exceptions given in *details* occurred.

Resolution

Contact your local customer support.

5003

The input object name is already defined. Change the object name.

Explanation

The input object name is already defined.

Resolution

Change the object name.

5004

An error occurred in the log monitoring mechanism on node *node*. Monitoring tasks will stop. Check the status of the node and update to the latest information.

Explanation

An error such as a node crash was detected in the log monitoring mechanism.

Resolution

Verify that the node is in operation, and update to the latest information.

5005

Reading GDS configuration information was unsuccessful on node *node*. Contact the system administrator.

Explanation

Reading GDS configuration information was unsuccessful on node *node*.

Resolution

Contact your local customer support.

5006

An invalid class name was specified. Input another name.

Explanation

An invalid character was specified for the class name.

Resolution

Click <OK> and type a class name again.

For information on characters available for class names, see "[A.1.1 Object Name](#)."

5007

An invalid group name was specified. Input another name.

Explanation

An invalid character was specified for the group name.

Resolution

Click <OK> and type a group name again.

For information on characters available for group names, see "[A.1.1 Object Name.](#)"

5008

An invalid volume name was specified. Input another name.

Explanation

An invalid character was specified for the volume name.

Resolution

Click <OK> and type a volume name again.

For information on characters available for volume names, see "[A.1.1 Object Name.](#)"

5009

An invalid disk name was specified. Input another name.

Explanation

An invalid character was specified for the disk name.

Resolution

Click <OK> and type a disk name again.

For information on characters available for disk names, see "[A.1.1 Object Name.](#)"

5010

The file system type or mount point has not been set. Select a file system type and/or input a mount point.

Explanation

<Create> was selected before specifying the file system type or mount point.

Resolution

Select a file system type and/or input a mount point, and click <Create>.

5011

The input mount point is already in use. Change the mount point.

Explanation

The input mount point is already in use with a different file system.

Resolution

Change the mount point, and click <Create>.

5012

A system disk cannot be removed from group configuration disks.

Explanation

In system disk setting, deletion of a system disk was attempted.

Resolution

System disks cannot be deleted from group configuration disks.

5013

You cannot add a physical disk. Up to 8 disks can be added to group configuration disks.

Explanation

In system disk setting, adding the ninth disk to group configuration disks was attempted.

Resolution

A maximum of 8-way multiplex mirroring is supported with system disks. You must limit the group configuration disks to eight.

5014

There are no free slices in the system disk. Secure a free slice before setting the system disk.

Explanation

The number of slices on the system disk has reached the upper limit (seven) and system disk setting cannot be performed.

Resolution

Remove one or more system disk's slices, for example, with the format (1M) command before performing system disk setting.

5015

There are unexpected device ID on node *node*. Exit.

Explanation

Starting GDS Management View failed due to unsupported GDS configuration information on node *node*.

Resolution

Contact your local customer support.

5016

The order of removing disks and groups composing the group is wrong.
Remove them in the proper order.

Explanation

A disk or group not removable was selected.

Resolution

When you remove disks and groups composing a group, you must select them in the inverse order against the added order.

5017

The stripe width setting is wrong. Specify an integer that is the power of 2.

Explanation

An invalid character or value was specified for the stripe width.

Resolution

Click <OK>, and input a valid stripe width again.

5018

The specified name is already in use. Input another name.

Explanation

The specified name is in use.

Resolution

Click <OK> and type an unused class name, and then click <Next>.

5019

An invalid volume size was specified. Input a valid size.

Explanation

The specified volume size is invalid.

Resolution

Click <OK> and type a valid volume size.

When the maximum size is unknown, you can apply the maximum size to [Volume Size] by clicking <Maximum Size>.

5020

The GDS Management View processing was aborted.

Explanation

The GDS Management View processing was aborted. A command the GDS Management View issued may be timed out.

Resolution

Collect investigation material and contact your local customer support.

5021

`The status gathering was aborted.`

Explanation

The process to gather object statuses was aborted. The status gathering command the GDS Management View issued may be timed out.

Resolution

Collect investigation material and contact your local customer support.

E.5.2 Warning Messages (5000,5100-5199)

5000

Command line output: *details*

Explanation

An error given in *details* occurred when executing the command.

Resolution

Identify the cause referring to the command line output. When you cannot identify the cause, contact your local customer support.

5100

`The input class name is already defined. Register the disk with this class?`

Explanation

The input class name is already defined.

Resolution

Click <Yes> to add the disk to the specified class, or <No> to register the disk with another class.

5101

`The input group name is already defined. Connect the disk with this group?`

Explanation

The input group name is already defined.

Resolution

Click <Yes> to additionally connect the disk with the specified group, or <No> to connect the disk with another group.

5102

If you register a physical disk with this class, all the data contained in the disk will be lost. Continue?

Explanation

All of the disk data will be lost if a physical disk is registered with the class.

Resolution

Click <Yes> to register the physical disk, or <No> to cancel.

5103

If you remove the volume, all the volume data will be lost. Continue?

Explanation

All of the volume data will be lost if you remove the volume.

Resolution

Click <Yes> to remove the volume, or <No> to cancel.

5104

The volume will be recovered. Continue?

Explanation

A volume that cannot be accessed will be recovered.

Resolution

Click <Yes> to recover the volume, or <No> to cancel.

5105

The selected disk is normal. Swap disks?

Explanation

A normal disk was selected for swapping.

Resolution

Click <Yes> to swap the disks, or <No> to cancel.

5106

For automatic configuration settings at least two mounted I/O adapters are necessary. This system configuration does not support automatic configuration settings. Go to the Class Configuration screen.

Explanation

With automatic configuration settings, disks connected to different I/O adapters can be mirrored, and on that account automatic configuration settings cannot be adopted when all the selected disks are connected to one I/O adapter.

Resolution

From the Settings menu, select the desired configuration settings (Class Configuration, Group Configuration, Volume Configuration) and configure.

5107

All the data contained in the selected disk will be lost. Continue?

Explanation

Data in the disk you selected for automatic configuration will be lost.

Resolution

Click <Yes> to continue, or <No> to cancel.

5108

The physical disk will be excluded from GDS. To swap the disks, the relative slices will stop. Detaching the disk will affect data reliability. Continue?

Explanation

When physical disks are swapped, the relative slices will stop operating until the restore is complete, and the reliability of the volume will fall.

Resolution

Click <Yes> to swap physical disks, or <No> to cancel.

5109

The physical disk will be included into GDS. All the data in the physical disk will be lost. Continue?

Explanation

When a physical disk is restored, data in the physical disk will be lost.

Resolution

Click <Yes> to restore the physical disk, or <No> to cancel.

5110

The slice will be detached from the volume. Detaching the slice will affect data reliability. Continue?

Explanation

Detaching a slice will lower reliability.

Resolution

Click <Yes> to detach the slice, or <No> to cancel.

5111

The detached slice will be attached to the volume. All the data in the slice will be lost. Continue?

Explanation

When a slice is attached, data in the slice will be overwritten.

Resolution

Click <Yes> to attach the slice, or <No> to cancel.

5112

The volume will be stopped. Continue?

Explanation

The volume will be stopped.

Resolution

Click <Yes> to stop volume, or <No> to cancel.

5113

Removing all disks configuring a class will result in removing the class as well. Continue?

Explanation

Removing all the disks in a class will result in removing the class as well.

Resolution

Click <Yes> to remove class, or <No> to cancel.

5114

Removing all disks configuring a group will result in removing the group as well. Continue?

Explanation

Removing all the disks in a group will result in removing the group as well.

Resolution

Click <Yes> to remove group, or <No> to cancel.

5115

A file system will be created on the selected volume. Continue?

Explanation

A file system will be created on the selected volume. If you continue, all the data in the volume will be lost.

Resolution

Click <Yes> to create a file system, or <No> to cancel.

5116

The selected volume will be unmounted and mount information for the volume will be deleted. It will be no longer available as a file system. Continue?

Explanation

The selected volume will be unmounted and mount information for the volume will be deleted. If it is deleted, the volume can no longer be used as a file system.

Resolution

Click <Yes> to remove the file system, or <No> to cancel.

5117

Changes on GDS configuration information were detected. Update to the latest configuration information?

Explanation

Commands other than those from GDS Management View changed the GDS configuration information.

Resolution

Click <Yes> to update the GDS Management View display, or <No> to cancel.

5118

The physical disk will be excluded from GDS. Continue?

Explanation

In order to swap physical disks, an unused disk or a spare disk will be detached. This will not affect the volume.

Resolution

Click <Yes> to swap the physical disks, or <No> to cancel.

5119

A file system has been configured on the selected volume. Remove the file system from the File System Configuration.

Explanation

The volume cannot be removed because of a file system configured on the volume.

Resolution

Remove the file system from the [File System Configuration] screen and then remove the volume.

5120

Physical disks have not been swapped. Verify that swapping took place. Are you sure you want to attach the disk?

Explanation

Physical disks may not have been swapped.

Resolution

Click <Yes> to attach the physical disk, or <No> to cancel.

5121

The status of physical disks connected to the system will be updated. Updating to the latest information may consume time. Continue?

Explanation

The status of physical disks connected to the system will be updated. During the update, operations from GDS Management View are unavailable.

Resolution

Click <Yes> to update the physical disk information, or <No> to cancel.

5122

A file system is created in a volume within the class. You cannot change attributes because it will change the path. Change the attributes after removing the file system.

Explanation

Renaming the class containing a volume with a file system was attempted.

Resolution

Renaming the class will result in changing the access path. First delete the file system.

5123

A file system is created in the volume. You cannot change attributes because it will change the path. Change the attributes after removing the file system.

Explanation

Renaming the volume with a file system was attempted.

Resolution

Renaming the volume will result in changing the access path. First delete the file system.

5124

No disk is available as a system disk after unmirroring. Recover the defective disk before unmirroring the system disk.

Explanation

Unmirroring a system disk was attempted when the mirrored disk was abnormal.

Resolution

First swap the abnormal physical disk for another. After recovering the physical disk, unmirror the system disk.

5125

Unmirroring the system disk will affect system reliability. Are you sure you want to unmirror?

Explanation

<Unmirror> was selected in the [Unmirror System Disk] screen.

Resolution

Click <Yes> to unmirror the system disk, or <No> to cancel.

5126

Since synchronization copying is in process, configuring system disk mirroring is not complete. After synchronization copying is complete, select System Disk Settings again.

Explanation

While configuring system disk mirroring, a volume in process of synchronization copying was detected.

Resolution

Click <OK> and close the message. From the GDS Management View screen, monitor the root class volume in process of synchronization copying. Verify that the volume status changes from "warnings" to "active," and try configuring mirroring for the system disk again.

5127

Configuring system disk mirroring is not complete. You must rewrite the mount information. Proceed with the system disk mirroring configuration?

Explanation

System disk mirroring has not been configured yet.

Resolution

Click <Yes> to rewrite the mount information and finish the configuration, or <No> to cancel.

5128

The slice will become inaccessible. Continue?

Explanation

The detached slice will be made inaccessible.

Resolution

Click <Yes> to restrict access to the slice, or <No> to allow access.

5129

The slice will become accessible. Continue?

Explanation

The slice being inaccessible as a result of stopping the slice will be made accessible.

Resolution

Click <Yes> to allow access to the slice, or <No> to restrict access.

5130

Copying of the slice will be cancelled. It cannot be resumed. Continue?

Explanation

The resynchronization copying of the slice will be cancelled. The copying cannot be resumed from the point where it stops.

Resolution

Click <Yes> to cancel resynchronization copying, or <No> not to cancel.

5131

Copying of the slice will start (resume). Continue?

Explanation

Copying of the slice will be performed to synchronize the data of the slices of the volume. Copying of an interrupted slice will resume as well.

Resolution

Click <Yes> to start (resume) resynchronization copying, or <No> not to start (resume).

5132

Remove a group in the 'root' class?

Explanation

A group in the "root" class is going to be removed.

Resolution

Click <Yes> to remove the group, or <No> to cancel.

5133

An error was detected on *node*. Display *node* information?

Explanation

An object error was detected on another node *node* configuring the cluster system.

Resolution

Click <Yes> to display the information about the node with the error, or <No> not to display.

5134

An error was detected on *node*. Check the node for the status after the current operation ends.

Explanation

An object error was detected on another node *node* configuring the cluster system.

Resolution

Check the node on which the error was detected for the status after the current operation ends.

5135

Obtaining the file system type list failed. File system cannot be created.

Explanation

Creating a file system failed since available supported file systems cannot be found.

Resolution

Install supported file systems, and then create a file system.

5136

All physical disks are registered.

Explanation

All of the local disks are already registered with the class.

Resolution

There is no resolution.

5138

This physical disk cannot be included into GDS because it is being handled at maintenance work. See the maintenance state of the physical disk.

Explanation

The physical disk is being handled at maintenance work and cannot be included into GDS management.

Resolution

See the maintenance state of the physical disk. For details, see "Machine Administration Guide."

5139

The proxy will be joined. The joined proxy will lose existing data. Do you want to continue?

Explanation

If the proxy is joined, its contents are overwritten by data of the master volume.

Resolution

Click <Yes> to join the proxy, or <No> to cancel the process.

5140

The proxy will be joined. Do you want to continue?

Explanation

The master group and the proxy group will be joined.

Resolution

Click <Yes> to join the proxy, or <No> to cancel the process.

5141

The proxy will be related. Do you want to continue?

Explanation

The master and the proxy will be related.

Resolution

Click <Yes> to relate the proxy, or <No> to cancel the process.

5142

The relationship between the master and the proxy will be canceled. Do you want to continue?

Explanation

The relationship between the master and the proxy will be canceled.

Resolution

Click <Yes> to break the relationship, or <No> to cancel the process.

5143

The master and the proxy will be rejoined. The rejoined proxy will lose existing data. Do you want to continue?

Explanation

If the proxy is rejoined, its contents are overwritten by data of the master.

Resolution

Click <Yes> to rejoin the proxy, or <No> to cancel the process.

5144

`Slices of the master and the proxy will be swapped. Do you want to continue?`

Explanation

Slices comprising the master and the proxy are switched.

Resolution

Click <Yes> to switch the slices, or <No> to cancel the process.

5145

`Reboot the system to enable the system disk settings.`

Explanation

System disk setting ended normally. It is necessary that the system be rebooted to enable system disk settings.

Resolution

Click <OK> to close the message, and reboot the system immediately.

5146

`Reboot the system in multi-user mode to complete unmirroring.`

Explanation

System disk unmirroring ended normally. It is necessary that the system be rebooted in multi-user mode to complete system disk unmirroring.

Resolution

Click <OK> to close the message, and reboot the system in multi-user mode immediately.

5147

`The system disk includes a master or a proxy volume. Break the relationship of the master and proxy before unmirroring the system disk.`

Explanation

The group to which the system disk belongs includes a master or a proxy volume.

Resolution

Cancel the proxy related to the system disk and then unmirror the system disk.

E.5.3 Information Messages (5200-5299)

5200

The slice was attached and the volume will be restored.

Explanation

The slice was attached. In order to restore the volume, resynchronization copying will start.

Resolution

Click <OK> and close the message. Check the completion of resynchronization copying in volume information field on the main screen. For the resynchronization copying progress, check the slice information field in the main screen.

5201

The volume restore is complete.

Explanation

The volume was restored.

Resolution

Click <OK> and close the message.

5202

The slice was detached. The following access path is available.
Access path: /dev/sfdsk/classname/rdisk/slicename

Explanation

The slice was detached to make it accessible.

Resolution

Pay attention to the displayed access path and click <OK>.

5203

The disk was excluded from GDS. Swap physical disks. Ask your local customer support for physical disk swapping.

Explanation

For swapping, a physical disk was excluded from GDS.

Resolution

Click <OK>, and contact your local customer support to have physical disks swapped.

5204

The physical disk was included into GDS. The volume will be restored.

Explanation

The swapped physical disk was placed under GDS.

Resolution

Click <OK> and close the message.

5205

An optimum disk pair cannot be found. Create a volume from the Class Configuration screen.

Explanation

The number of disks or disk size was inappropriate and therefore could not perform optimum mirroring with automatic configuration.

Resolution

From the [Settings] menu, click the configuration settings (Class Configuration, Group Configuration, Volume Configuration) and set up.

5206

A file system was created.

Explanation

A file system was created and mount information was added.

Resolution

Click <OK> and close the message.

5207

The file system was removed.

Explanation

The file system and mount information was removed.

Resolution

Click <OK> and close the message.

5209

Status monitoring will be updated every *interval* minutes.

Explanation

The object monitoring interval was changed.

Resolution

Click <OK> and close the message.

5210

Exit Global Disk Services?

Explanation

GDS Management View will be exited.

Resolution

Click <Yes> to exit, or <No> to continue operation.

5211

The following files were updated to complete mirroring system disk.

```
/etc/system  
/etc/vfstab  
/kernel/drv/sfdsk.conf
```

Explanation

The system disk setting process completed mirroring definition.

Resolution

Click <OK> and close the message.

5212

System disk unmirroring is complete.

Explanation

The system disk unmirroring process completed unmirroring.

Resolution

Click <OK> and close the message.

5213

The slice was made inaccessible. Activate the slice in order to gain access.

Explanation

The detached slice was made inaccessible.

Resolution

Click <OK> and close the message.

5214

The slice is now accessible.

Explanation

The detached slice was made accessible.

Resolution

Click <OK> and close the message.

5215

The selected slice is neither temp nor temp-stop status. You cannot perform the operation on the selected slice.

Explanation

The [Stop/Activate Slice] operation was attempted on a slice not in the temp or temp-stop status.

Resolution

Click <OK> and close the message.

5216

Copying of the slice was cancelled.

Explanation

The resynchronization copying of the slice was cancelled.

Resolution

Click <OK> and close the message.

5217

Copying of the slice was started (resumed).

Explanation

The resynchronization copying was started (resumed).

Resolution

Click <OK> and close the message.

5218

The selected volume is not in the copy status. You cannot perform the operation on the selected volume.

Explanation

The copying operation was attempted while the slices in the volume are not in the copy status.

Resolution

Click <OK> and close the message.

5219

The selected slice is neither active nor stop status. You cannot perform the operation on the selected slice.

Explanation

The [Detach Slice] operation was attempted on a slice not in the active or stop status.

Resolution

Click <OK> and close the message.

5220

There is a slice in the copy status in the selected volume. You cannot perform the operation on the selected volume.

Explanation

The [Start Copying] operation was attempted on a volume with a slice in process of synchronization copying.

Resolution

Click <OK> and close the message.

5221

There are no slices in the invalid status in the selected volume. You cannot perform the operation on the selected volume.

Explanation

The [Start Copying] operation was attempted on a volume without a slice in the invalid status.

Resolution

Click <OK> and close the message.

5222

There is no available proxy for the selected master. Create a volume that can be joined first.

Explanation

There is no volume that can be joined as a proxy volume to the selected master volume.

Resolution

Create a volume to be joined to the master volume first conforming to the following rules.

- Belonging to the class to which the master volume belongs
- Belonging to a mirror group or to a single disk

- Not belonging to the group to which the master volume belongs
- Not belonging to the single disk to which the master volume belongs
- Not related to another volume as a master or a proxy
- Equivalent to the master volume in size

For details, see "[A.1.8 Proxy Configuration Preconditions](#)" and the description of the Join option in "[D.14 sdxproxy - Proxy object operations](#)."

5223

There is no available proxy for the selected master. Create a group that can be joined first.

Explanation

There is no group that can be joined as a proxy group to the master group.

Resolution

Create a group to be joined to the master group first conforming to the following rules.

- Belonging to the class to which the master group belongs
- Mirror type
- Not related to another group as a master or a proxy
- Including no volume

For details, see "[A.1.8 Proxy Configuration Preconditions](#)" and the description of the Join option in "[D.14 sdxproxy - Proxy object operations](#)."

5224

Joining the proxy succeeded. Synchronization copying will start.

Explanation

The master and the proxy were joined. Synchronization copying from the master to the proxy will start to synchronize them.

Resolution

Click <OK> to close the message. Check the synchronization copying status in the slice information field in the main screen.

5225

There is no available proxy for the selected master. Create a volume that can be related first.

Explanation

There is no volume that can be related to the selected master volume.

Resolution

Create a volume to be related to the master volume conforming to the following rules.

- Belonging to the class to which the master volume belongs
- Belonging to a mirror group or to a single disk
- Not belonging to the group to which the master volume belongs

- Not belonging to the single disk to which the master volume belongs
- Not related to another volume as a master or a proxy
- Equivalent to the master volume in size

For details, see "[A.1.8 Proxy Configuration Preconditions](#)" and the description of the Relate option in "[D.14 sdxproxy - Proxy object operations](#)."

5226

There is no available proxy for the selected master. Create a group that can be related first.

Explanation

There is no group that can be related as a proxy group to the selected master group.

Resolution

Create a group to be related to the master group conforming to the following rules.

- Belonging to the class to which the master group belongs
- Mirror type
- Not related to another group as a master or a proxy
- Whose volume configuration (offset and size) is same as that of the master group

For details, see "[A.1.8 Proxy Configuration Preconditions](#)" and the description of the Relate option in "[D.14 sdxproxy - Proxy object operations](#)."

5227

Relating the proxy succeeded.

Explanation

The master and the proxy were related.

Resolution

Click the <OK> to close the message.

5228

Breaking the proxy succeeded.

Explanation

The relationship between the master and the proxy was broken.

Resolution

Click the <OK> to close the message.

5229

Parting the proxy succeeded.

Explanation

The proxy was parted from the master.

Resolution

Click <OK> to close the message.

5230

Rejoining the proxy succeeded. The master and the proxy will be resynchronized

Explanation

The master and the proxy were rejoined. Resynchronization copying from the master to the proxy will start to resynchronize them.

Resolution

Click <OK> to close the message.

Check the resynchronization copying status in the volume Information field in the main screen. For the resynchronization copying progress, check the slice information field in the main screen.

5231

The update of the proxy started.

Explanation

Copying from the master to the proxy started.

Resolution

Click <OK> to close the message. Check the completion of the resynchronization copying in the volume information field in the main screen. For the copying status, check the slice information field in the main screen.

5232

Updating the proxy succeeded. Start up the proxy to access.

Explanation

Copying from the master to the proxy started. Before the copying ends, you may start the proxy for access. Contents of the proxy are overwritten by the data of the master at the moment when the copying started.

Resolution

Click <OK> to close the message. Check the copying status. in the slice information field on the main screen.

5233

The restoration of the master started.

Explanation

Copying from the proxy to the master started.

Resolution

Click <OK> to close the message. Check the completion of the copying in the volume information field in the main screen. For the copying progress, check the slice information field on the main screen.

5234

Restoring the master succeeded. Start up the master to access.

Explanation

Copying from the proxy to the master started. Before the copying ends, you may start the master for access. Contents of the master are overwritten by the data of the proxy at the moment when the copying started.

Resolution

Click <OK> to close the message. Check the copying status in the volume information field on the main screen.

5235

Rejoining the proxy succeeded, and the restoration of the master started.

Explanation

The master and the proxy were rejoined and synchronization copying from the proxy to the master started.

Resolution

Click <OK> to close the message. Check the completion of the resynchronization copying in the volume information field on the main screen. For the resynchronization copying progress, check the slice information field in the main screen.

5236

Rejoining the proxy and restoring the master succeeded. Start up the master to access.

Explanation

The master and the proxy were rejoined and resynchronization copying from the proxy to the master started. Before the resynchronization copying ends, you may start the master for access. Contents of the master are overwritten by the data of the proxy at the moment when the copying started.

Resolution

Click <OK> to close the message. For the resynchronization copying status, check the slice information field on the main screen.

5237

Swapping the slices of the master and the proxy succeeded.

Explanation

Slices comprising the master and the proxy were switched.

Resolution

Click <OK> to close the message.

E.6 Error Codes (90000-90099)

This section explains the error codes output in ACM's inter-server replication function messages.

90005

ERROR: class name too long

Explanation

The specified class name is too long.

Resolution

Specify a correct class name.

90007

ERROR: volume name too long

Explanation

The specified volume name is too long.

Resolution

Specify a correct volume name.

90009

ERROR: no enough address space

Explanation

The OS malloc function returned an error.

Resolution

Check whether memory and a swap area are sufficient.

90034

ERROR: cannot start required number of copy sessions

Explanation

The requested number of copy sessions cannot be started.

Resolution

Check the configuration and status of the copy source and the copy destination volumes.

If normal slices composing the copy source or copy destination volume are less than the number of sessions to be started, restore the slice status or change the volume configuration as needed. Alternatively, specify a lower number of sessions not to exceed the number of normal slices of each volume.

90035

ERROR: copy session not found

Explanation

No copy session exists between the copy source and the copy destination volumes.

Resolution

Confirm that there is no error in the configuration and status of the copy source and the copy destination volumes.

90036

ERROR: copy sessions are in inconsistent statuses

Explanation

The statuses of copy sessions between the copy source and the copy destination volumes are inconsistent.

Resolution

With the relevant ACM commands, suspend or cancel the EC or REC copying, and start the copying again as needed.

Appendix F Troubleshooting

Since GDS automatically attempts recovery when various abnormalities are detected, users do not have to be aware of most abnormalities. However, depending on the type of abnormality, or when multiple abnormalities occur at the same time, automatic recovery may not be possible.

This section describes resolutions to failures on objects and physical disks managed by GDS for systems administrators.



As a general rule, all operations explained here must be executed with superuser authority.

F.1 Resolving Problems

F.1.1 Slice Status Abnormality

If the slice status is one of the following statuses, take action as indicated for the relevant situation.

- (1) Mirror slice configuring the mirror volume is in **INVALID** status.
- (2) The copy destination slice was made **INVALID** due to an I/O error generated on the copy source slice during synchronization copying.
- (3) Slice configuring the volume is in **TEMP** status.
- (4) Slice configuring volume is in **TEMP-STOP** status.
- (5) Slice configuring the volume is in **COPY** status.
- (6) Slice configuring the volume is in **NOUSE** status.

(1) Mirror slice configuring the mirror volume is in **INVALID** status.

Explanation

You can check the status of the slice configuring the volume as follows.

```
# sdxinfo -S -o Volume1

OBJ      CLASS  GROUP  DISK      VOLUME  STATUS
-----
slice    Class1 Group1  Object1   Volume1 ACTIVE
slice    Class1 Group1  Object2   Volume1 INVALID
```

In this example, among the slices that exist in volume Volume1, the slice within Object2 is in **INVALID** status, as shown in the **STATUS** field. Object2 is a disk or lower level group connected to the highest level mirror group Group1.

The following five events could possibly cause the **INVALID** status of the mirror slice Volume1.Object2.

- An I/O error occurred on the mirror slice Volume1.Object2.

(Cause a)

A disk component relevant to Object2 failed to operate properly, and an I/O error occurred on the mirror slice Volume1.Object2.

(Cause b)

A component other than disks relevant to Object2 (such as an I/O adapter, an I/O cable, an I/O controller, a power supply, and a fan) failed to operate properly, and an I/O error occurred on the mirror slice Volume1.Object2.

- An I/O error occurred on the mirror slice Volume1.Object1 during synchronization copying to the mirror slice Volume1.Object2.

(Cause a')

A disk component relevant to Object1 failed to operate properly during synchronization copying to the mirror slice Volume1.Object2, and an I/O error occurred on the copy source slice Volume1.Object1.

(Cause b')

A component other than disks relevant to Object1 (such as an I/O adapter, an I/O cable, an I/O controller, a power supply and a fan) failed to operate properly during synchronization copying to the mirror slice Volume1.Object2, and an I/O error occurred on the mirror slice Volume1.Object1 that is the copy source.

- Others

(Cause c)

Synchronization copying to the mirror slice Volume1.Object2 was cancelled as the result of a cause such as [Cancel Copying] selection from the GDS Management View, sdxcopy command execution, or a power outage.

Resolution

- 1) Identify the physical disk name of a faulty disk using the sdxcinfo command.

(Example A1)

```
# sdxcinfo -G -o Volume1
OBJ  NAME  CLASS  DISKS  BLKS  FREEBLKS  SPARE
-----
group  Group1  Class1  Object1:Object2  17596416  17498112  0

# sdxcinfo -D -o Volume1 -e long
OBJ  NAME  TYPE  CLASS  GROUP  DEVNAM  DEVBLKS  FREEBLKS  DEVCONNECT  STATUS  E
-----
disk  Object1  mirror  Class1  Group1  c1t1d0  17596416  *  node1
ENABLE  *
disk  Object2  mirror  Class1  Group1  c2t3d0  17596416  *  node1
ENABLE  *
```

In this example, Object2 is a disk connected with the highest level group Group1. As indicated in the E field, an I/O error occurred on the disk Object2, and the possible cause is (Cause a) or (Cause b). The physical disk name of the disk Object2 is c2t3d0 as shown in the DEVNAM field.

In example A1, if the value 0 is in the E field of the disk Object2 including a slice in the INVALID status and if the value 1 is in the E field of the disk Object1 that is mirrored with the disk Object2, it indicates an I/O error occurred on the disk Object1 and the possible cause is (Cause a') or (Cause b'). In such a case, see "(2) The copy destination slice was made INVALID due to an I/O error generated on the copy source slice during synchronization copying." and perform restoration.

If the E field of any disk does not contain the value 1 in example A1, the possible cause is (Cause c).

(Example B1)

```
# sdxcinfo -G -o Volume1
OBJ  NAME  CLASS  DISKS  BLKS  FREEBLKS  SPARE
-----
group  Group1  Class1  Object1:Object2  35127296  35028992  0
```

```

group Object1 Class1 Disk1:Disk2          35127296      * *
group Object2 Class1 Disk3:Disk4          35127296      * *

# sdxinfo -D -o Volume1 -e long
OBJ  NAME  TYPE  CLASS  GROUP  DEVNAM  DEVBLKS  FREEBLKS  DEVCONNECT  STATUS  E
-----
disk  Disk1  stripe Class1  Object1 c1t1d0  17596416      * node1
ENABLE 0
disk  Disk2  stripe Class1  Object1 c1t2d0  17596416      * node1
ENABLE 0
disk  Disk3  stripe Class1  Object2 c2t3d0  17682084      * node1
ENABLE 1
disk  Disk4  stripe Class1  Object2 c2t4d0  17682084      * node1
ENABLE 0

```

In this example, Object2 is a lower level group connected with the highest level group Group1. As indicated in the E field, an I/O error occurred on the disk Disk3 connected with Object2 and the possible cause is (Cause a) or (Cause b). The physical disk name corresponding to Disk3 is c2t3d0 as shown in the DEVNAM field.

In example B1, if the 0 value is in the E field of the disks (Disk3 and Disk4) connected with the lower level group Object2 including a slice in the INVALID status and if the value 1 is in the E field of the disk (Disk1 or Disk2) connected with the lower level group Object1 that is mirrored with the lower level group Object2, it indicates an I/O error occurred on the disk connected with Object1 and the possible cause is (Cause a) or (Cause b). In such a case, see "(2) The copy destination slice was made INVALID due to an I/O error generated on the copy source slice during synchronization copying." and perform restoration.

In example B1, if there are no disks with "1" in the E field, the possible cause of the INVALID status is (Cause c).

2) Refer to disk driver log messages and check the physical disk abnormalities.

The causes of disk hardware failures can be failures or defects in components such as I/O adapters, I/O cables, I/O controllers, power supplies, and fans other than the disks.

Contact your local customer support and specify which component failed, or might be defective.

If there are no failures or defective components, the possible cause of the INVALID is (Cause c).

The resolution procedure is illustrated below for each of the three causes a, b, and c.

a. For (Cause a)

a1) Perform the following operations before and after disk swapping. For the procedures for swapping disks from Operation Management View, see "5.3.4 Disk Swap."

Before swapping the disks, execute the following command.

(Example A1)

```
# sdxswap -O -c Class1 -d Object2
```

In the example, disk Object2 connected to the highest level group Group1 will be swapped.

(Example B1)

```
# sdxswap -O -c Class1 -d Disk3
```

In the example, disk Disk3 will be swapped. Disk3 is a disk connected to lower level group Object2, which is a lower level group of the highest level group Group1.

a2) Swap disks.

a3) After swapping the disks, execute the following command.

(Example A3)

```
# sdxswap -I -c Class1 -d Object2
```

or

(Example B3)

```
# sdxswap -I -c Class1 -d Disk3
```

a4) Check the slice status according to step 3).

b. For (Cause b)

b1) Shut down the system once, repair the disabled part, and reboot the system. Consequently, synchronization copying is performed and the mirroring status is restored.

b2) Check the slice status according to step 3).

c. For (Cause c)

c1) Perform synchronization copying of mirror volume.

```
# sdxcopy -B -c Class1 -v Volume1
```

c2) Check the slice status according to step 3).

3) You can confirm the recovery of the slice configuring the volume, as shown below.

```
# sdxinfo -S -o Volume1
OBJ    CLASS  GROUP  DISK    VOLUME  STATUS
-----
slice  Class1 Group1  Object1 Volume1 ACTIVE
slice  Class1 Group1  Object2 Volume1 ACTIVE
```

In this example, the slices within Object1 and Object2 are both in ACTIVE status. This indicates that the recovery process was completed successfully.

(2) The copy destination slice was made INVALID due to an I/O error generated on the copy source slice during synchronization copying.

Explanation

When an I/O error occurs on the copy source slice during synchronization copying, the copy destination slice becomes INVALID while the source slice is still ACTIVE.

The following two events are possible causes.

(Cause a)

A disk component of the copy source failed to operate properly, and an I/O error occurred on the copy source slice.

(Cause b)

A component other than the copy source disk (such as an I/O adapter, an I/O cable, an I/O controller, a power supply, and a fan) failed to operate properly during synchronization copying, and an I/O error occurred on the copy source slice.

For details on determining whether the status is relevant to one of these events and identifying the physical disk name of a faulty disk, see [Explanation] and step 1) of [Resolution] described in "(1) Mirror slice configuring the mirror volume is in INVALID status."

Resolution

First examine the physical disk abnormalities referring to disk driver log messages and so on. Then contact your local customer support and locate the disabled or faulty part.

When the possible cause is (Cause b), shut down the system once, repair the disabled part, and reboot the system. Consequently, synchronization copying is performed and the mirroring status is restored.

When the possible cause is (Cause a), follow the procedures below and repair the slice. The procedure is illustrated for each of the following three situations.

A. For /(root), /usr, or /var

B. For the swap area

C. For others (other than /(root), /usr, /var, swap)

The following illustrates restoration procedures when the class name is Class1, the volume name is Volume1, the name of a faulty disk of the copy source is Disk1, and the name of a disk of the copy destination is Disk2 as examples.

```
# sdxinfo -S -o Volume1
OBJ    CLASS  GROUP  DISK    VOLUME  STATUS
-----
slice  Class1  Group1 Disk1    Volume1 ACTIVE
slice  Class1  Group1 Disk2    Volume1 INVALID

# sdxinfo -G -o Volume1
OBJ    NAME    CLASS  DISKS          BLKS    FREEBLKS  SPARE
-----
group  Group1  Class1 Disk1:Disk2    17596416 17498112 0

# sdxinfo -D -o Volume1 -e long
OBJ    NAME    TYPE  CLASS  GROUP  DEVNAM  DEVBLKS  FREEBLKS  DEVCONNECT  STATUS  E
-----
disk   Disk1   mirror Class1  Group1 c1t1d0 17596416          * node1
ENABLE 1
disk   Disk2   mirror Class1  Group1 c2t3d0 17596416          * node1
ENABLE 0
```

A. For /(root), /usr, or /var

Contact your local customer support.

B. For the swap area

B.1) Remove the volume from the swap area.

```
# swap -d /dev/sfdsk/Class1/dsk/Volume1
```

Depending on the part or severity of failure in disks that constitute the volume, the swap -d command may fail due to an I/O error. In this event, remove the volume from the swap area performing steps B.1.1) through B.1.3).

B.1.1) Comment out the /dev/sfdsk/Class1/dsk/Volume1 line in the /etc/vfstab file to prevent use of the volume as a swap area after the system is rebooted.

```
# vi /etc/vfstab
Before edit:
/dev/sfdsk/Class1/dsk/Volume1 - - swap - no -

After edit:
#/dev/sfdsk/Class1/dsk/Volume1 - - swap - no
-
```


B.1.2) To remove the dump device after the system is rebooted, edit the definition of DUMPADM_DEVICE in the /etc/dumpadm.conf file.

```
# vi /etc/dumpadm.conf

Before edit:
DUMPADM_DEVICE=/dev/sfdsk/Class1/dsk/Volume1

After edit:
DUMPADM_DEVICE=swap
```

B.1.3) Reboot the system.

```
# shutdown -y -i6 -g0
```

B.2) Stop the volume.

```
# sdxvolume -F -c Class1 -v Volume1
```

B.3) Restore the status of the copy destination slice in INVALID status.

```
# sdxfix -v -c Class1 -d Disk2 -v Volume1
```

B.4) Verify that the restored copy destination slice is in STOP status and the copy source slice is in INVALID status now.

```
# sdxinfo -S -o Volume1
OBJ    CLASS  GROUP  DISK    VOLUME  STATUS
-----
slice  Class1 Group1 Disk1    Volume1 INVALID
slice  Class1 Group1 Disk2    Volume1 STOP
```

B.5) Start the volume.

```
# sdxvolume -N -c Class1 -v Volume1
```

B.6) Add the volume to the swap area again.

```
# swap -a /dev/sfdsk/Class1/dsk/Volume1
```

When step B.1.1) was performed, undo the edit that was made in the /etc/vfstab file.

```
# vi /etc/vfstab
Before edit:
#/dev/sfdsk/Class1/dsk/Volume1 -      -      swap -      no      -

After edit:
/dev/sfdsk/Class1/dsk/Volume1 -      -      swap -      no      -
```

B.7) Remove the faulty copy source disk from GDS management to give it a replaceable status.

```
# sdxswap -O -c Class1 -d Disk1
```

B.8) Swap the faulty copy source disk.

B.9) Put the swapped disk back in control of GDS management to make it available.

```
# sdxswap -I -c Class1 -d Disk1
```

C. For others (other than /(root), /usr, /var, and swap)

C.1) Exit applications using the volume.

C.2) Unmount the file system on the volume when it has been mounted.

```
# umount /dev/sfdsk/Class1/dsk/Volume1
```

Depending on the part or severity of failure in disks that constitute the volume, the umount command may fail due to an I/O error. In this event, execute the umount command again with the -f option.

C.3) Stop the volume.

```
# sdxvolume -F -c Class1 -v Volume1
```

C.4) Restore the status of the copy destination slice in INVALID status.

```
# sdxfix -v -c Class1 -d Disk2 -v Volume1
```

C.5) Verify that the restored copy destination slice is in STOP status and the copy source slice is in INVALID status now.

```
# sdxinfo -S -o Volume1
OBJ    CLASS  GROUP  DISK    VOLUME  STATUS
-----
slice  Class1 Group1  Disk1   Volume1 INVALID
slice  Class1 Group1  Disk2   Volume1 STOP
```

C.6) Start the volume.

```
# sdxvolume -N -c Class1 -v Volume1
```

C.7) The consistency of volume data may have been lost. Restore the backup data or perform repair using the fsck(1M) command if necessary.

 **Note**

.....
If the I/O error occurred on the copy source slice during resynchronization copying after the system went down, restoration may possibly be performed with the fsck(1M) command.
.....

C.8) Remove the faulty copy source disk from GDS management to make it a replaceable status.

```
# sdxswap -O -c Class1 -d Disk1
```

C.9) Swap the faulty copy source.

C.10) Put the swapped disk back in the control of GDS management to make it available.

```
# sdxswap -I -c Class1 -d Disk1
```

(3) Slice configuring the volume is in TEMP status.

Explanation

The slice was not attached after it has been detached with the `sdxslice` command. Or else, you have not performed [Attach Slice] after performing [Detach Slice] from Operation Management View.

Resolution

Attach the slice again with the `sdxslice` command, or perform [Attach Slice] from Operation Management View as necessary.

(4) Slice configuring volume is in TEMP-STOP status.

Explanation

The slice was not activated after it has been stopped with the `sdxslice` command, or the detached node is not current node. Or else, you have not performed [Stop Slice] after performing [Activate Slice] from Operation Management View.

Resolution

Activate slice or take over slice with the `sdxslice` command as needed. Or, perform [Activate Slice] from Operation Management View.

(5) Slice configuring the volume is in COPY status.

Explanation

In order to attach a slice, synchronization copying is currently in process. Or, synchronization copying is in process between master and proxy.

Resolution

Wait until synchronization copying is complete. Note that a slice in the process of synchronization copying will not restrict you from accessing an active volume.

(6) Slice configuring the volume is in NOUSE status.

Explanation

When the status of disk related to slice is either in DISABLE or SWAP status, the slice becomes NOUSE to inhibit slice operation.

Resolution

Recover disk in DISABLE or SWAP status. For details, see "[F.1.2 Disk Status Abnormality](#)."

F.1.2 Disk Status Abnormality

If the disk status is one of the following statuses, take action as indicated for the relevant situation.

- [\(1\) Disk is in DISABLE status.](#)
- [\(2\) Disk is in SWAP status.](#)
- [\(3\) Disk is in I/O error status.](#)

(1) Disk is in DISABLE status.

Explanation

Disk ID information is checked when booting the system. If the ID information is regarded invalid at the time, the disk will become DISABLE.

There are three reasons that may cause this DISABLE status.

(Cause a)

You changed the I/O cable connection and booted the system.

(Cause b)

You swapped the disk and booted the system.

(Cause c)

Disk failure occurred.

Resolution

- a) If the possible cause is (Cause a), shut down the system, connect the I/O cables correctly and boot again.
- b) If the possible cause is (Cause b), shut down the system, return the original disk and boot again. If disk swap is necessary, follow the procedures and swap the disks.
For the procedures for swapping disks, see "[D.8 sdxswap - Swap disk](#)" and "[5.3.4 Disk Swap](#)."
- c) If the possible cause is (Cause c), follow the procedures and swap the disks.

However, when the cause is (Cause c) in one of the following situations, the class to which the disk in DISABLE status belongs must be forcibly removed, so contact your local customer support.

- A single disk is in DISABLE status and an INVALID single volume exists on the single disk.
- A disk in DISABLE status is connected to the group and the group includes an INVALID volume.

With a single node system, in this situation or in a situation where the disk in DISABLE status is connected to a lower level group, do not perform any operation, such as volume creation or lower level group connection, that changes the configuration on the class to which the disk in DISABLE status belongs. If this type of operation is performed, configurations of the class may come to conflict.

(2) Disk is in SWAP status.

Explanation

If the sdxswap -O command is executed, or [Swap Physical Disk] is performed from Operation Management View, the disk will be in SWAP status.

Resolution

Complete the disk swap and restore the disk by either executing the sdxswap -I command, or by performing [Restore Physical Disk] from Operation Management View.

(3) Disk is in I/O error status.

Explanation

If an I/O error occurs, 1 (one) is given as a value in the E field displayed with the sdxinfo -e long command and the disk icon in GDS Management View turns red. Here, the disk is in ENABLE (enabled) status.

Resolution

Check the statuses of volumes and slices relative to the disk with an I/O error, and perform restoration, referring to "[F.1.3 Volume Status Abnormality](#)" or "[F.1.1 Slice Status Abnormality](#)". The I/O error status is removed, during the restoration by:

- Executing [Swap Physical Disk] in GDS Management View for the relevant disk
- Executing the sdxswap -O command for the relevant disk
- Normally ending synchronization copying of volumes relative to the relevant disk

If none of the operations above are performed during the recovery and the I/O error remains unremoved, after restoring hardware and such, use the sdxfix -D command to remove the I/O error status.

```
# sdxfix -D -c class name -d disk name -e online
```

F.1.3 Volume Status Abnormality

If the volume status is one of the following statuses, take action as indicated for the relevant situation.

- (1) Mirror volume is in INVALID status.
- (2) Single volume is in INVALID status.
- (3) Stripe volume or volume in concatenation group is in INVALID status.
- (4) Master volume is in INVALID status.
- (5) Proxy volume is in INVALID status.
- (6) Volume is in STOP status.
- (7) I/O error occurs although mirror volume is in ACTIVE status.
- (8) An I/O error occurs on a single volume.
- (9) An I/O error occurs on a stripe volume or a volume in a concatenation group.

(1) Mirror volume is in INVALID status.

Explanation

You can confirm the status of the volume as shown below.

```
# sdxinfo -v -o Volume1
OBJ   NAME   CLASS  GROUP  SKIP  JRM  1STBLK  LASTBLK  BLOCKS  STATUS
-----
volume *      Class1 Group1 *    *      0    65535   65536  PRIVATE
volume Volume1 Class1 Group1 off   on    65536 17596415 17530880 INVALID
```

In this example, volume Volume1 that exists in the highest level group Group1 is in INVALID status, as shown in the STATUS field.

If none of the mirror slices consisting the mirror volume contains valid data (ACTIVE or STOP), the mirror volume becomes INVALID. You cannot start a volume in INVALID status.

There are two reasons that may cause this INVALID status.

(Cause a)

Disk is in DISABLE status.

(Cause b)

Master-proxy relationship was cancelled forcibly while master data was being copied to proxy.

Resolution

1) Confirm that there is a disk in DISABLE status within the group with which the volume is associated as follows.

(Example A1)

```
# sdxinfo -G -o Volume1
OBJ   NAME   CLASS  DISKS          BLKS    FREEBLKS  SPARE
-----
group Group1  Class1  Disk1:Disk2    17596416  0 0

# sdxinfo -D -o Volume1
OBJ   NAME   TYPE   CLASS  GROUP  DEVNAM  DEVBLKS  DEVCONNECT  STATUS
```

```

-----
disk  Disk1  mirror Class1  Group1  c1t1d0  17596416  node1          ENABLE
disk  Disk2  mirror Class1  Group1  c2t3d0  17596416  node1          DISABLE
-----

```

In this example, disks Disk1 and Disk2 are connected to the highest level mirror group Group1. The disk Disk2 is in DISABLE status as shown in the STATUS field.

(Example B1)

```

# sdxinfo -G -o Volume1
OBJ      NAME      CLASS    DISKS                BLKS      FREEBLKS  SPARE
-----
group    Group1    Class1   Group2:Group3        35127296  17530880  0
group    Group2    Class1   Disk1:Disk2          35127296      *  0
group    Group3    Class1   Disk3:Disk4          35127296      *  0

# sdxinfo -D -o Volume1
OBJ      NAME      TYPE     CLASS    GROUP    DEVNAM    DEVBLKS    DEVCONNECT    STATUS
-----
disk     Disk1     stripe  Class1   Group2   c1t1d0    17596416    node1          ENABLE
disk     Disk2     stripe  Class1   Group2   c1t2d0    17596416    node1          DISABLE
disk     Disk3     stripe  Class1   Group3   c2t3d0    17682084    node1          ENABLE
disk     Disk4     stripe  Class1   Group3   c2t4d0    17682084    node1          ENABLE

```

In this example, lower level stripe groups, Group2 and Group3 are connected to the highest level mirror group Group1. Disk Disk2 which is connected to Group2 is in DISABLE status as shown in the STATUS field.

2) If the possible cause is (Cause a), restore the disk by following the procedures in "F.1.2 Disk Status Abnormality."

3) From the disks and lower level groups connected to the highest level mirror group, determine the disk or lower level group to which the slice you will use to recover data belongs. Then, execute the sdxfix command to recover data.

(Example A3)

```
# sdxfix -v -c Class1 -d Disk1 -v Volume1
```

In this example, Volume1 is recovered after a slice in disk Disk1.

(Example B3)

```
# sdxfix -v -c Class1 -g Group3 -v Volume1
```

In this example, Volume1 is recovered after a slice in lower level stripe group Group3.

4) Start the volume.

```
# sdxvolume -N -c Class1 -v Volume1 -e nosync
```

5) Access Volume1 and check its contents. Restore backup data or run fsck to regain data integrity as necessary.

6) Perform synchronization copying on volume.

```
# sdxcopy -B -c Class1 -v Volume1
```

(2) Single volume is in INVALID status.

Explanation

You can confirm the status of the volume as shown below.

```
# sdxinfo -V -o Volume1
OBJ      NAME      CLASS  GROUP  SKIP  JRM  1STBLK  LASTBLK  BLOCKS  STATUS
-----
volume *      Class1  Disk1  *      *      0      32767   32768  PRIVATE
volume Volume1 Class1  Disk1  off   on    32768   65535   32768  INVALID
volume *      Class1  Disk1  *      *      65536  8421375 8355840 FREE
```

In the example, the single volume Volume1 that exists on single disk Disk1 is in INVALID status, as shown in the STATUS field.

You cannot start a volume in INVALID status.

There are two reasons that may cause this INVALID status.

(Cause a)

Single disk is in DISABLE status. In this case, the single slice becomes NOUSE status.

(Cause b)

Master-proxy relationship was cancelled forcibly while master data was being copied to proxy.

Resolution

1) Confirm that the single disk is in DISABLE status as shown below.

```
# sdxinfo -D -o Volume1
OBJ      NAME      TYPE   CLASS  GROUP  DEVNAM  DEVBLKS  DEVCONNECT  STATUS
-----
disk     Disk1     single Class1  *      c1t11d0 8355840  node1       DISABLE
```

In this example, the single disk Disk1 is in DISABLE status, as shown in the STATUS field.

2) If the possible cause is (Cause a), restore the disk by following the procedures given in section "[F.1.2 Disk Status Abnormality.](#)"

3) Execute the sdxfix command to recover the single volume's data.

```
# sdxfix -V -c Class1 -d Disk1 -v Volume1
```

4) Start the volume.

```
# sdxvolume -N -c Class1 -v Volume1
```

5) Access Volume1 and check its content. Restore backup data or run the fsck command to regain data integrity as necessary.

(3) Stripe volume or volume in concatenation group is in INVALID status.

Explanation

You can confirm the status of the volume as shown below.

```
# sdxinfo -v -o Volume1
OBJ      NAME      CLASS  GROUP  SKIP JRM 1STBLK  LASTBLK  BLOCKS  STATUS
-----
volume *      Class1  Group1 *    *      0    65535   65536  PRIVATE
volume Volume1 Class1  Group1 off  on    65536 17596415 17530880 INVALID
```

In this example, volume Volume1 that exists in the highest level group Group1 is in INVALID status, as shown in the STATUS field.

If any of the disks related to volume is in DISABLE status, the slices consisting that volume become NOUSE status, and the volume becomes INVALID. You cannot start a volume in INVALID status.

Resolution

1) You can confirm the status of the disk related to the volume as shown below.

```
# sdxinfo -G -o Volume1 -e long
OBJ      NAME      CLASS  DISKS          BLKS    FREEBLKS SPARE MASTER TYPE  WIDTH
-----
group Group1  Class1  Group2:Group3  70189056 65961984 *    *    stripe 32
group Group2  Class1  Disk1:Disk2    35127296 *    *    *    concat *
group Group3  Class1  Disk3:Disk4    35127296 *    *    *    concat
*

# sdxinfo -D -o Volume1
OBJ      NAME      TYPE  CLASS  GROUP  DEVNAM  DEVBLKS  DEVCONNECT  STATUS
-----
disk Disk1    concat Class1  Group2  c1t1d0 17596416 node1        ENABLE
disk Disk2    concat Class1  Group2  c1t2d0 17596416 node1        DISABLE
disk Disk3    concat Class1  Group3  c2t3d0 17682084 node1        ENABLE
disk Disk4    concat Class1  Group3  c2t4d0 17682084 node1        ENABLE
```

In this example, the lower level concatenation groups Group2 and Group3 are connected to the highest level stripe group Group1, and Disk2 connected to Group2 is in DISABLE status as shown in the STATUS field.

2) Follow the procedures in "[F.1.2 Disk Status Abnormality](#)" and restore the disk status.

3) Execute the sdxfix command to recover the volume's data. With -g option, indicate the highest level group name (in this example, Group1).

```
# sdxfix -v -c Class1 -g Group1 -v Volume1
```

4) Start the volume.

```
# sdxvolume -N -c Class1 -v Volume1
```

5) Access Volume1 and check its content. Restore backup data or run the fsck command to regain data integrity as necessary.

(4) Master volume is in INVALID status.

Explanation

If the copying process fails while copying data from the proxy volume to the master volume because of an I/O error or such, the status of the master volume to which the data is being copied becomes INVALID.

Resolution

1) Check if there is a DISABLE status disk in the group to which the volume belongs with the following command.

```
# sdxinfo -D -o Volume1
```

OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	DEVCONNECT	STATUS
disk	Disk1	mirror	Class1	Group1	c1t1d0	8421376	*	ENABLE
disk	Disk2	mirror	Class1	Group1	c1t2d0	8421376	*	DISABLE

In this example, Disk2 is in DISABLE status.

If there is a disk in DISABLE status, see section "(1) Disk is in DISABLE status." in "F.1.2 Disk Status Abnormality," and check which of the causes listed in that section applies. If the possible cause is (Cause a) or (Cause b), follow the procedures and restore the disk.

2) Follow the procedures given in section "(1) Mirror slice configuring the mirror volume is in INVALID status." in "F.1.1 Slice Status Abnormality," and check if there is a disk hardware abnormality. If there is, identify the faulty part.

When the abnormality is caused by a failed or defective non-disk component, repair the faulty part.

3) Procedures to restore the data for different scenarios are given below.

- Non-disk component failure
 - When recovering data using the proxy volume:
 - > Follow steps a) to restore.
 - When recovering data using backup data on media such as tapes:
 - > Follow steps b) to restore.
- When caused by a disk component failure:
 - When the disk does not belong to master group:
 - some disks connected to the group have failed
 - > Follow steps c) to restore.
 - all disks connected to the group have failed
 - > Follow steps d) to restore.
 - When the disk belongs to master group:
 - > Follow steps e) to restore.

a) Procedures to recover master volume data using proxy volume.

a1) In order to check if the proxy volume that will be used to recover data is separated from the master volume, execute the `sdxinfo -V -e` long command, and check the PROXY field.

a2) If the proxy volume is not separated, execute the following command.

```
# sdxproxy Part -c Class1 -p Volume2
```

a3) Exit all applications accessing the proxy volume. When using the proxy volume as a file system, execute unmount. When using the proxy volume as a file system, execute unmount.

a4) If the proxy volume is started, execute the following command.

```
# sdxvolume -F -c Class1 -v Volume2
```

a5) Recover master volume data using the proxy volume's data.

```
# sdxproxy RejoinRestore -c Class1 -p Volume2
```

b) Procedures to recover data using backup data.

b1) When the volume is in INVALID status, you must first change it to STOP status. Decide on the disk (slice) you wish to use to recover data, and execute the sdxfix command.

```
# sdxfix -V -c Class1 -d Disk1 -v Volume1
```

In this example, Volume 1 is restored after a slice in Disk 1.

b2) When the volume to be restored is stopped, start it with the following command.

```
# sdxvolume -N -c Class1 -v Volume1 -e nosync
```

b3) Access the volume to be restored and check its contents. Restore backup data or run fsck to regain data integrity as necessary.

b4) When mirroring is configured with the volume, perform synchronization copying.

```
# sdxcopy -B -c Class1 -v Volume1
```

c) Procedures to swap some disks connected to the group.

c1) If you restore the INVALID master volume later using data of a proxy volume related to the master volume, or use data of proxy volumes related to the master volume after restoring it, part the proxy volumes using the sdxproxy Part command.

```
# sdxproxy Part -c Class1 -p Volume2
```

c2) When there is a volume in INVALID status in the group, change it to STOP status with the sdxfix -V command. -d option indicates the disk without abnormality.

```
# sdxfix -V -c Class1 -d Disk1 -v Volume1
```

c3) Follow the procedures and swap the disks. For details, see "[D.8 sdxswap - Swap disk](#)" and "[5.3.4 Disk Swap](#)."

c4) Recover the master volume data. If data will be recovered using the proxy volume, follow procedures described in a). If data will be recovered using backup data on media such as tapes, follow procedures described in b).

d) Procedures to swap all disks connected to the group.

d1) Exit all applications accessing the master volume and the proxy volume that will be used to recover data. When using the proxy volume or the master volume as a file system, execute unmount.

d2) Stop the master volume and proxy volume in d1).

```
# sdxvolume -F -c Class1 -v Volume1
# sdxvolume -F -c Class1 -v Volume2
```

d3) Execute the `sdxproxy RejoinRestore` command and restore the master volume data using proxy volume in d1). If the command terminates normally and the master volume is not in `INVALID` status, restoration process is complete, and you do not need to perform steps d4) and after.

```
# sdxproxy RejoinRestore -c Class1 -p Volume2
```

d4) Execute the `sdxproxy Swap` command and swap the slices of the master volume with the proxy volume in d1).

```
# sdxproxy Swap -c Class1 -p Volume2
```

d5) By performing step d4), the status of master volume will not be in `INVALID` status, and the status of the proxy volume becomes `INVALID`. Follow the procedures given in section "[\(5\) Proxy volume is in INVALID status.](#)" in "[F.1.3 Volume Status Abnormality](#)," and restore the proxy volume in `INVALID` status.

d6) Execute the `sdxproxy Swap` command and swap the slices of the master volume and the proxy volume you swapped in step d4).

```
# sdxproxy Swap -c Class1 -p Volume2
```

e) Procedures to swap disks connected to the master group.

e1) Exit all applications accessing the master group, and volumes in the proxy group that will be used to recover data. When using the volume as a file system, execute `umount`.

e2) Stop all volumes in the master group and the proxy group in e1).

```
# sdxvolume -F -c Class1 -v Volume1
# sdxvolume -F -c Class1 -v Volume2
```

e3) Execute the `sdxproxy RejoinRestore` command and restore the master group data using the proxy group in e1). If the command terminates normally and all master volumes are not in `INVALID` status, restoration process is complete, and you do not need to perform steps e4) and after.

```
# sdxproxy RejoinRestore -c Class1 -p Volume2
```

e4) Execute the `sdxproxy Swap` command and swap the slices of the master group and the proxy group in e1).

```
# sdxproxy Swap -c Class1 -p Group2
```

e5) By performing step e4), the master volume will not be in `INVALID` status and the status of the proxy volume becomes `INVALID`. Follow the procedures given in section "[\(5\) Proxy volume is in INVALID status.](#)" in "[F.1.3 Volume Status Abnormality](#)," and restore the proxy volume in `INVALID` status.

e6) Execute the `sdxproxy Swap` command and swap the slices of the master group and the proxy group you swapped in step e4).

```
# sdxproxy Swap -c Class1 -p Group2
```

(5) Proxy volume is in `INVALID` status.

Explanation

If the copying process fails while copying data from the master volume to the proxy volume because of an I/O error or such, the status of the proxy volume to which the data is being copied becomes INVALID.

Resolution

1) Check if there is a DISABLE status disk in the group to which the volume belongs with the following command.

```
# sdxinfo -D -o Volume1
```

OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	DEVCONNECT	STATUS
disk	Disk1	mirror	Class1	Group1	c1t1d0	8421376	*	ENABLE
disk	Disk2	mirror	Class1	Group1	c1t2d0	8421376	*	DISABLE

In this example, Disk2 is in DISABLE status.

If there is a disk in DISABLE status, see section "(1) Disk is in DISABLE status." in "F.1.2 Disk Status Abnormality," and check which of the causes (There are three causes a, b and c listed.) listed in that section applies. If it is due to (Cause a) or (Cause b), follow the procedures and restore the disk.

2) Follow the procedures given in "(1) Mirror slice configuring the mirror volume is in INVALID status." in "F.1.1 Slice Status Abnormality," and check if there is a disk hardware abnormality. If there is, identify the faulty part. When the abnormality was caused by a failed or defective non-disk component, repair the faulty part.

3) Procedures to restore the data for different scenarios are given below.

- When caused by a non-disk component failure:
 - > Follow steps a) to restore.
- When caused by a disk component failure:
 - disk does not belong to proxy group
 - some disks connected to the group have failed
 - > Follow steps b) to restore.
 - When all disks connected to the group have failed:
 - > Follow steps c) to restore.
- When the disk belongs to proxy group:
 - > Follow steps d) to restore.

a) Procedures to recover proxy volume data using the master volume.

a1) In order to check if the proxy volume is separated from the master volume, execute the `sdxinfo -V -e` long command, and check the PROXY field.

a2) If the proxy volume is not separated, execute the following command.

```
# sdxproxy Part -c Class1 -p Volume2
```

a3) Rejoin the proxy volume with the master volume.

```
# sdxproxy Rejoin -c Class1 -p Volume2
```

b) Procedures to swap some disks connected to the group.

b1) Cancel the relationship with master volume using the `sdxproxy Break` command.

```
# sdxproxy Break -c Class1 -p Volume2
```

b2) Separate the volumes that are in INVALID status in the group with the `sdxfix -V` command, and change them to STOP status. `-d` option indicates the disk without abnormality.

```
# sdxfix -V -c Class1 -d Disk1 -v Volume2
```

b3) Follow the procedures and swap the disks. For details, see "D.8 sdxswap - Swap disk," or section "5.3.4 Disk Swap."

b4) Join the master and the proxy again with the `sdxproxy Join` command.

```
# sdxproxy Join -c Class1 -m Volume1 -p Volume2
```

c) Procedures to swap all disks connected to the group.

c1) Cancel the relationship with the master using the `sdxproxy Break` command.

```
# sdxproxy Break -c Class1 -p Volume2
```

c2) Exit all applications accessing the volume in the group. When using the volume as a file system, execute `umount`.

c3) Stop all volumes in the group.

```
# sdxvolume -F -c Class1 -v Volume2
```

c4) Check the volume configuration of the group (such as volume names and sizes) with the `sdxinfo` command, and keep a note of it.

c5) Remove all volumes in the group.

```
# sdxvolume -R -c Class1 -v Volume2
```

c6) Follow the procedures and swap the disks. For details, see "D.8 sdxswap - Swap disk," or section "5.3.4 Disk Swap."

c7) Create the volume that you removed in step c5) again.

```
# sdxvolume -M -c Class1 -g Group2 -v Volume2 -s size
```

c8) Stop the volume you created in step c7).

```
# sdxvolume -F -c Class1 -v Volume2
```

 **Note**

When the volume is in INVALID status, the error message "ERROR: disk: volume in INVALID status" will be displayed. Ignore the message and proceed to the next step.

c9) Join the master volume and the proxy volume again, with the `sdxproxy Join` command.

```
# sdxproxy Join -c Class1 -m Volume1 -p Volume2
```

d) Procedures to swap disks connected to the proxy group.

d1) Cancel the relationship with the master using the `sdxproxy Break` command.

```
# sdxproxy Break -c Class1 -p Group2
```

d2) Exit all applications accessing the volume in the group. When using the volume as a file system, execute `umount`.

d3) Stop all volumes in the group.

```
# sdxvolume -F -c Class1 -v Volume2
```

 **Note**

When the volume is in INVALID status, the error message "ERROR: *disk: volume* in INVALID status" will be displayed. Ignore the message and proceed to the next step.

d4) Remove all volumes in the group.

```
# sdxvolume -R -c Class1 -v Volume2
```

d5) Follow the procedures and swap the disks. For details, see "[D.8 sdxswap - Swap disk](#)," or section "[5.3.4 Disk Swap](#)."

d6) Join the master group and the proxy group again with the `sdxproxy Join` command.

```
# sdxproxy Join -c Class1 -m Group1 -p Group2 -a Volume1=Volume2:on
```

(6) Volume is in STOP status.

Explanation

Normally, volumes automatically start when the system is booted and become ACTIVE. The volume status will change to STOP when the volume is stopped with the Stop Volume menu in the GDS Management View or the `sdxvolume -F` command.

In a cluster system, among volumes within GDS shared classes registered with cluster applications, volumes other than proxy volumes start or stop according to the cluster application modes. If a cluster application is in Offline mode, volumes other than proxy volumes are in STOP status.

Accessing a volume in STOP status will result in an EIO error (I/O error) or an ENXIO error (No such device or address).

 **See**

For the problem in a cluster system that volumes in a shared class not registered with a cluster application do not start at node startup, see "[\(4\) The GFS Shared File System is not mounted on node startup](#)." in "[F.1.9 Cluster System Related Error](#)."

Resolution

Start the volumes with the Start Volume menu in GDS Management View or the `sdxvolume -N` as necessary.

To start volumes within a GDS shared class registered with a cluster application, change the cluster application mode to Online.

(7) I/O error occurs although mirror volume is in ACTIVE status.

Explanation

A mirror volume consists of multiple slices, and in an event of an I/O error, the crashed slice will be detached. Therefore, accessing the volume will complete normally.

However, when an I/O error occurs when only one slice is ACTIVE amongst those configuring the volume, accessing the volume will result in an error. At such time, the status of the slice and the volume remains ACTIVE.

Probable situations resulting in such a problem will be described using a two-way multiplex mirroring configuration, where two disks or two lower level groups are connected to a group. As an example, means to circumvent such problems will also be described.

(Situation 1)

One of the slices was detached with the `sdxslice -M` in order to backup volume data. While accessing the volume, an I/O error occurred with the other slice.

(Prevention 1)

Before executing the `sdxslice -M` command, connect a reserved disk and temporarily configure a three-way multiplex mirroring, or make the mirrored volume available for backup.

(Situation 2)

While restoring a slice with an I/O error, an I/O error also occurred on another slice.

(Prevention 2)

By securing a spare disk within the class, effects due to delay in restoring the slice will be avoided to a certain degree

Resolution

Identify the cause of I/O error occurrence in the last ACTIVE slice, by referring to the disk driver log message.

Resolutions are described below assuming the following three circumstances:

- a. Error occurred due to a disk component failure. Will attempt recovery using backup data.
- b. Error occurred due to a disk component failure. Will attempt data recovery from a slice in INVALID status.
- c. Error occurred due to a failed or a defective non-disk component failure.

a. When the error cause is a disk component failure and recovery is performed using backup data

a1) When the error was caused by a disk component failure, no slice with valid data exists. Restore data from the backup data following the procedures given below.

a2) Exit the application accessing the volume. When the volume is used as a file system, execute the unmount command. When I/O error occurs on the unmount command, execute -f option of the unmount command.

a3) Stop the volume with the `sdxvolume` command.

```
# sdxvolume -F -c Class1 -v Volume1
```

a4) If there is a TEMP status slice within the volume, attempt recovery following the procedures given in "[F.1.1 Slice Status Abnormality](#)."

a5) If there is a NOUSE status slice within the volume, attempt recovery following the procedures given in "[F.1.1 Slice Status Abnormality](#)."

a6) Record the volume size which can be checked as follows.

```
# sdxinfo -v -o Volume1
OBJ    NAME    CLASS  GROUP  SKIP  JRM  1STBLK  LASTBLK  BLOCKS  STATUS
-----
volume *      Class1 Group1 *    *      0      32767    32768  PRIVATE
volume Volume1 Class1 Group1 off  on    32768  4161535  4128768 STOP
```

In this example, the volume size would be 4128768 blocks given in Volume1 BLOCKS field.

a7) Remove the volume with the sdxvolume command.

```
# sdxvolume -R -c Class1 -v Volume1
```

a8) Swap disks following the procedures given in "5.3.4 Disk Swap" and "D.8 sdxswap - Swap disk."

a9) Create a volume with the sdxvolume command again. For the *number_of_blocks*, use the size recorded in a6), in this example, 4128768.

```
# sdxvolume -M -c Class1 -g Group1 -v Volume1 -s number_of_blocks
```

a10) Finally, restore the backup data to Volume1.

b. When the error cause is a disk component failure and data is restored from a slice in INVALID status

b1) When the error was caused by a disk failure, and when no backup data exists, or even if it did, the data is too old, restore data from the detached INVALID status slice, following the procedures given below.

b2) Exit the application accessing the volume. When the volume is used as a file system, execute the unmount command. When I/O error occurs on the unmount command, execute -f option of the unmount command.

b3) Stop the volume with the sdxvolume command.

```
# sdxvolume -F -c Class1 -v Volume1
```

b4) If there is a TEMP status slice within the volume, attempt recovery following the procedures given in "F.1.1 Slice Status Abnormality."

b5) If there is a NOUSE status slice within the volume, attempt recovery following the procedures given in "F.1.1 Slice Status Abnormality."

b6) Determine the original mirror slice after the volume is recovered. Then, execute the sdxfix command.

(Example 1)

```
# sdxfix -v -c Class1 -d Disk2 -v Volume1
```

In this example, data is recovered from a mirror slice in the disk Disk2 which is connected to the highest level mirror group.

(Example 2)

```
# sdxfix -v -c Class1 -g Group2 -v Volume1
```

In this example, data is recovered from a mirror slice in the lower level group Group2 which is connected to the highest level mirror group.

b7) Start the volume.

```
# sdxvolume -N -c Class1 -v Volume1 -e nosync
```

b8) Create backup of Volume1 and regain data integrity by running fsck as necessary.

b9) Lastly, swap disks following the procedures given in "5.3.4 Disk Swap" and "D.8 sdxswap - Swap disk."

c. When the error cause is a non-disk component failure or defect

The slice with valid data exists within the disk, and shut down the system once, recover the failed component, and then reboot the system. Synchronization copying is automatically performed and the mirroring status will be recovered.

(8) An I/O error occurs on a single volume.

Explanation

Since a single volume consists of only one slice, accessing the volume at the time of an I/O error will result in an error. However, the status of slice and volume will remain ACTIVE.

Resolution

Identify the cause of I/O error occurrence by referring to the disk driver log message.

How to resolve the problem is described in two cases:

- a. When the error cause is a disk component failure and recovery is performed using backup data
- b. When the error cause is a non-disk component failure or defect

a. When the error cause is a disk component failure and recovery is performed using backup data

a1) In the event of a disk component failure, there will be no slice with valid data. Follow the procedures below and restore the data using the backup data. In this example, Disk1 (c1t11d0) has a failure.

```
# sdxinfo -D -o Disk1
OBJ    NAME    TYPE    CLASS  GROUP  DEVNAM  DEVBLKS  DEVCONNECT  STATUS
-----
disk   Disk1   single Class1  *      c1t11d0 8493876  node1       ENABLE
```

a2) Search the volumes within the faulty disk using the sdxinfo command. And record the volume size,

```
# sdxinfo -v -o Disk1
OBJ    NAME    CLASS  GROUP  SKIP  JRM  1STBLK  LASTBLK  BLOCKS  STATUS
-----
volume *      Class1 Disk1  *      *      0      32767   32768  PRIVATE
volume Volume1 Class1 Disk1  off   on    32768  65535   32768  ACTIVE
volume Volume2 Class1 Disk1  off   on    65536  4194303 4128768 ACTIVE
volume *      Class1 Disk1  *      *      4194304 8421375 4227072 FREE
```

In this example, Volume1 and Volume2 are within the faulty Disk1. The size of Volume1 would be 32,768 blocks as shown in the BLOCKS field. The size of Volume2 would be 4,128,768 blocks as shown in the BLOCKS field.

a3) Exit the application accessing the volume. When the volume is used as a file system, execute the unmount command. When I/O error occurs on the unmount command, execute -f option of the unmount command.

a4) Stop the volume with the sdxvolume command.

```
# sdxvolume -F -c Class1 -v Volume1,Volume2
```

a5) Remove the volumes with the sdxvolume command.

```
# sdxvolume -R -c Class1 -v Volume1
# sdxvolume -R -c Class1 -v Volume2
```

a6) Before swapping the disks, execute the following command.

```
# sdxswap -O -c Class1 -d Disk1
```

Note

If the disk is the only remaining disk in the disk class, the command results in an error as shown below. In that event, follow the steps a6'), a7') and a8').

SDX:sdxswap: ERROR: Disk1: The last ENABLE disk in class cannot be swapped

a7) Swap the disks.

a8) After swapping the disks, execute the following command.

```
# sdxswap -I -c Class1 -d Disk1
```

a6') Before swapping the disks, execute the following command.



If no error is output in a6), the steps a6'), a7'), and a8') are not required.

```
# sdxdisk -R -c Class1 -d Disk1
```

a7') Swap the disks.

a8') After swapping the disks, execute the following command.

```
# sdxdisk -M -c Class1 -d c1t11d0=Disk1:single
```

a9) Create volumes with the sdxvolume command again. For the -s option, use the size recorded in 2a), in this example.

```
# sdxvolume -M -c Class1 -d Disk1 -v Volume1 -s 32768
# sdxvolume -M -c Class1 -d Disk1 -v Volume2 -s 4128768
```

a10) Finally, restore the backup data to Volume1 and Volume2.

b. When the error cause is a non-disk component failure or defect

Shut down the system once, recover the failed component, and then reboot the system. Slice data is valid and there is no need to restore the data.

(9) An I/O error occurs on a stripe volume or a volume in a concatenation group.

Explanation

Since a stripe volume or a volume within a concatenation group consists of only one slice, accessing the volume at the time of an I/O error will also result in an error. However, the status of slice and volume will remain ACTIVE.

Resolution

Identify the cause of I/O error occurrence by referring to the disk driver log message.

You can confirm the error status of the disk related to volume and the physical disk name as shown below.

```
# sdxinfo -D -o Volume1 -e long
OBJ  NAME  TYPE  CLASS  GROUP  DEVNAM  DEVBLKS  FREEBLKS  DEVCONNECT  STATUS  E
-----
disk  Disk1  concat  Class1  Group2  c1t1d0  17596416  *  node1
ENABLE  0
disk  Disk2  concat  Class1  Group2  c1t2d0  17596416  *  node1
ENABLE  1
disk  Disk3  concat  Class1  Group3  c2t3d0  17682084  *  node1
```

```

ENABLE      0
disk  Disk4  concat Class1  Group3  c2t4d0  17682084      * node1
ENABLE      0

```

In this example, an I/O error occurs on Disk2, as shown in the E field. The physical disk name corresponding to Disk2 is c1t2d0, as shown in the DEVNAM field.

How to resolve the problem is described in two cases:

- a. When the error cause is a disk component failure and recovery is performed using backup data
- b. When the error cause is a non-disk component failure or defect

a. When the error cause is a disk component failure and recovery is performed using backup data

a1) In the event of a disk component failure, there will be no slices with valid data. Follow the procedures below and restore the data using the backup data.

a2) Record the configuration information of the group that was related to the failed disk using the `sdxinfo` command.

```

# sdxinfo -G -o Disk2 -e long
OBJ      NAME      CLASS  DISKS              BLKS      FREEBLKS SPARE MASTER TYPE  WIDTH
-----
group    Group1    Class1  Group2:Group3      70189056  65961984 *      *      stripe 32
group    Group2    Class1  Disk1:Disk2        35127296          * *      *      concat *
group    Group3    Class1  Disk3:Disk4        35127296          * *      *      concat
*

```

In this example, the lower level concatenation groups Group2 and Group3 are connected to the highest level stripe group Group1. The disks Disk1 and Disk2 are connected to Group2, and the disks Disk3 and Disk4 are connected to Group3. The stripe width for Group1 is 32 blocks.

a3) Search the volumes that exist in the highest level group that are related to the faulty disk using the `sdxinfo` command.

```

# sdxinfo -V -o Disk2
OBJ      NAME      CLASS  GROUP  SKIP JRM 1STBLK  LASTBLK  BLOCKS  STATUS
-----
volume  *          Class1  Group1 *    *    0      65535   65536  PRIVATE
volume  Volume1    Class1  Group1 *    *    65536  98303  32768  ACTIVE
volume  Volume2    Class1  Group1 *    *    98304  4227071 4128768 ACTIVE
volume  *          Class1  Group1 *    *    4227072 70189055 65961984 FREE

```

In this example, Volume1 and Volume2 exist in the highest level group Group1, that is related to the faulty disk Disk2. The size of Volume1 is 32768 blocks, and the size of Volume2 is 4128768 blocks as shown in the BLOCKS field.

a4) Exit the application accessing the volume. When the volume is used as a file system, execute `umount` command. When I/O error occurs on `umount` command, execute `-f` option of `umount` command.

a5) Stop the volume with the `sdxvolume` command.

```

# sdxvolume -F -c Class1 -v Volume1,Volume2

```

a6) Remove the volumes with the `sdxvolume` command.

```

# sdxvolume -R -c Class1 -v Volume1
# sdxvolume -R -c Class1 -v Volume2

```

a7) Disconnect the faulty disk from the group. If the group is in a hierarchical structure, disconnect from the higher group in descending order.

```
# sdxgroup -D -c Class1 -h Group1 -l Group2
# sdxdisk -D -c Class1 -g Group2 -d Disk2
```

In this example, the faulty disk Disk2 is connected to Group2, and Group2 is connected to Group1. Therefore, you should disconnect Group2 first, and then Disk2.

a8) Before swapping the disks, execute the following command.

```
# sdxswap -O -c Class1 -d Disk2
```

Note

If the disk is the only remaining disk in the disk class, the command results in an error as shown below. In that event, follow the steps a8'), a9') and a10').

```
SDX:sdxswap: ERROR: Disk2: The last ENABLE disk in class cannot be swapped
```

a9) Swap the disks.

a10) After swapping the disks, execute the following command.

```
# sdxswap -I -c Class1 -d Disk2
```

a8') Before swapping the disks, execute the following command.

Note

If no error is output in a8), the steps a8'), a9'), and a10') are not required.

```
# sdxdisk -R -c Class1 -d Disk2
```

a9') Swap the disks.

a10') After swapping the disks, execute the following command.

```
# sdxdisk -M -c Class1 -d c1t2d0=Disk2
```

a11) Connect the swapped disk to the group, referring to the group information recorded in a2). If the groups were in a hierarchical structure, connect the groups in an ascending order.

```
# sdxdisk -C -c Class1 -g Group2 -d Disk2
# sdxgroup -C -c Class1 -h Group1 -l Group2 -a type=stripe,width=32
```

a12) Create volumes with the sdxvolume command again. For the -s option, use the size recorded in a3), in this example, 32768 and 4128768.

```
# sdxvolume -M -c Class1 -g Group1 -v Volume1 -s 32768 -a pslice=off
# sdxvolume -M -c Class1 -g Group1 -v Volume2 -s 4128768 -a pslice=off
```

a13) Finally, restore the backup data to Volume1 and Volume2.

b. When the error cause is a non-disk component failure or defect

Shut down the system once, recover the failed component, and then reboot the system. Slice data is valid and there is no need to restore the data.

F.1.4 Class Status Abnormality

If the class status is one of the following statuses, take action as indicated for the relevant situation.

- (1) Class becomes closed status during operation.

(1) Class becomes closed status during operation.

Explanation

The class becomes closed when the number of configuration databases which store information on object configuration and object status within a class is insufficient, or when the communication error between nodes occurs in a cluster environment. All objects within a closed class are inaccessible.

An Insufficient number of configuration databases will occur under the following conditions:

1. When there are no disks that can be accessed normally, if there are two or less disks in ENABLE status.
2. When there are one or less disks that can be accessed normally, if there are three to five disks in ENABLE status.
3. When there are two or less disks that can be accessed normally, if there are more than six disks in ENABLE status.

However, in the event of root class, the class will not be closed unless there are no accessible disks.

GDS configuration databases cannot be stored in BCV devices and target (R2) devices since the devices are overwritten by data in copy source disks. Therefore, GDS does not regard BCV devices and target (R2) devices as "disks that can be accessed normally" described in the above conditions.

Resolution

1) You can check whether or not a class was closed during operation as follows. Do not reboot the system or restart sdxservd daemon, as it will make the checking impossible.

```
# /etc/opt/FJSVsdx/bin/sdxcdown
CLASS   DOWN REASON NDK  NEN  NDB  NLDB  DEVNAM
-----
-----
Class1  no    -      10  10   8    0
c1t1d0:c1t2d0:c1t3d0:c1t4d0:c2t1d0:c2t2d0:c2t3d0:c2t4d0
Class2  yes   Comm   10  10   8    0
c3t1d0:c3t2d0:c3t3d0:c3t4d0:c4t1d0:c4t2d0:c4t3d0:c4t4d0
Class3  yes   FewDB  10  10   1    7  c5t1d0
Class4  yes   NoDB   10  10   0    8  -
```

In this example, Class2, Class3, and Class4 with "yes" in the DOWN field are closed. The cause shown in the REASON field are as follows.

(Cause 1)

Comm Communication failure between nodes.

(Cause 2)

FewDB Insufficient number of valid configuration databases.

(Cause 3)

NoDB No valid configuration database.

2) Depending on specific causes, recovery may be difficult. First, collect the investigation material.

For information on how to collect the investigation material, see "[F.2 Collecting Investigation Material](#)."

Resolutions are described for the following two cases:

- a. Closed due to a communication error
- b. Closed due to an insufficient number of configuration databases

3a) In the event of (Cause 1), contact your local customer support.

3b) In the event of (Cause 2) or (Cause 3), all (or the majority) of the disks registered with class have abnormalities.

You can check the disks registered with class as follows.

```
# sdxinfo -D -c Class3
```

OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	DEVCONNECT	STATUS
disk	Disk31	mirror	Class3	Group1	c1t1d0	8847360	*	ENABLE
disk	Disk32	mirror	Class3	Group1	c2t1d0	8847360	*	ENABLE
disk	Disk33	mirror	Class3	Group2	c1t2d0	8847360	*	ENABLE
disk	Disk34	mirror	Class3	Group2	c2t2d0	8847360	*	ENABLE
disk	Disk35	mirror	Class3	Group3	c1t3d0	17793024	*	ENABLE
disk	Disk36	mirror	Class3	Group3	c2t3d0	17793024	*	ENABLE
disk	Disk37	mirror	Class3	Group4	c1t4d0	17793024	*	ENABLE
disk	Disk38	mirror	Class3	Group4	c2t4d0	17793024	*	ENABLE
disk	Disk39	spare	Class3	Group1	c1t5d0	17793024	*	ENABLE
disk	Disk40	spare	Class3	*	c2t5d0	17727488	*	ENABLE

In this example, ten disks from Disk31 to Disk40 are registered with Class3. Physical disk names are shown in the DEVNAM field. Identify the cause of abnormality with these physical disks by referring to disk driver log messages. The cause of abnormality could be either of the following:

(Failure 1)

Failed or defective non-disk component.

(Failure 2)

Failed disk component.

4b) In the event of (Failure 1), recover the failed or defective non-disk component (such as I/O adapter, I/O cable, I/O controller, power supply, and fan).

5b) For a local class or a shared class, execute the sdxfix command to restore the class status.

```
# sdxfix -C -c Class3
SDX:sdxfix: INFO: Class3: class recovery completed successfully
```

- If the sdxfix command ends normally, skip steps 6b) through 9b) and go on to step 10b).
- If the sdxfix command does not end normally, go on to step 6b).
- For the root class, go on to step 6b).

6b) Open the GDS configuration parameter file with an editor.

```
# vi /etc/opt/FJSVsdx/sdx.cf
```

Add the following one line in the end of the file.

```
SDX_DB_FAIL_NUM=0
```

7b) Reboot the system.

8b) Confirm that objects within the class are accessible.

```
# sdxinfo -c Class3
```

If nothing is displayed, recovery was unsuccessful. You will have to contact your local customer support. If information is displayed normally, proceed with the following procedures.

9b) In the event of (Failure 2), where a disk component has failed, follow the procedures in ["5.3.4 Disk Swap,"](#) or ["D.8 sdxswap - Swap disk,"](#) and swap the disks.

10b) After completing the recovery for both (Failure 1) and (Failure 2), check the number of valid configuration databases as described below.

```
# /etc/opt/FJSVsdx/bin/sdxcdowndown
CLASS   DOWN REASON NDK  NEN  NDB  NLDB DEVNAM
-----
-----
Class1  no   -      10  10   8   0
c1t1d0:c1t2d0:c1t3d0:c1t4d0:c2t1d0:c2t2d0:c2t3d0:c2t4d0
Class2  no   -      10  10   8   0
c3t1d0:c3t2d0:c3t3d0:c3t4d0:c4t1d0:c4t2d0:c4t3d0:c4t4d0
Class3  no   -      10  10   8   0
c5t1d0:c5t2d0:c5t3d0:c5t4d0:c6t1d0:c6t2d0:c6t3d0:c6t4d0
Class4  no   -      10  10   8   0
c7t1d0:c7t2d0:c7t3d0:c7t4d0:c8t1d0:c8t2d0:c8t3d0:c8t4d0
```

NLDB field gives the insufficient number of configuration databases. If this value is "0," the problem is resolved. If this value is "1" or more, there are still disks that have not been recovered. In the above example, all NLDB fields display "0," indicating the successful recovery.

When step 6b) was not performed, the following procedures are not required.

11b) Open the GDS configuration parameter file with an editor.

```
# vi /etc/opt/FJSVsdx/sdx.cf
```

Remove the following one line added in step 6b).

```
SDX_DB_FAIL_NUM=0
```

12b) Reboot the system.

If you cannot perform the recovery with the described procedures, contact your local customer support.

F.1.5 System Disk Abnormality

The following description shows resolutions to problems relevant to system disks on which the file systems / (root), /usr, and /var operates.

In one of the following circumstances, take action as indicated for the relevant situation.

- (1) [Create backup of root file system.](#)
- (2) [System can be booted, but the system disk data is invalid.](#)
- (3) [System cannot be booted. \(Failure of original boot disk\)](#)
- (4) [System cannot be booted. \(Boot disk data damage\)](#)
- (5) [System cannot be booted. \(Incorrect boot disk device name\)](#)
- (6) [The system cannot be booted. \(ufs logging errors\)](#)

(1) Create backup of root file system.

Explanation

Mirroring the system disk will protect data in case one of the physical disks crashes. However, to restore data damaged due to a critical fault as a result of multiple failures and so on or by an operational mistake, create backup data in advance.

Resolution

See "[6.1.1 Backing Up.](#)"

(2) System can be booted, but the system disk data is invalid.

Explanation

For some reason, the system disk data is invalid. You must restore data from the backup data created in advance.

Resolution

See "[6.1.2 Restoring \(When the System Can Be Booted\).](#)"

(3) System cannot be booted. (Failure of original boot disk)

Explanation

Even when you have the system disk mirrored, you may experience unsuccessful booting. For example, you may encounter such a problem when the file accessed during the booting process in the boot disk is physically damaged, or when a disk is swapped improperly.

If the console message of the unsuccessful boot disk indicates such is the case, try booting from the other boot disk, i.e. the mirror disk.

Resolution

1) Confirm the name of the mirror disk. Use the OpenBoot ok prompt shown below.

```
ok printenv boot-device
/pci@1f,4000/scsi@3/disk@0,0:a /pci@1f,4000/scsi@3/disk@1,0:a
```

In this example, the original disk name is /pci@1f,4000/scsi@3/disk@0,0:a, which is indicated on the left. The mirror disk is /pci@1f,4000/scsi@3/disk@1,0:a. When mirroring is configured with three-way or more multiplexing, multiple mirror disks exist.

2) Boot from the mirror disk.

```
ok boot /pci@1f,4000/scsi@3/disk@1,0:a
```


3) If booting is successful, find out what caused the original disk to fail, and recover by swap the disk.



For details on disk swap, see "[D.8 sdxswap - Swap disk](#)" and resolutions described in "[5.3.4 Disk Swap](#)."

When the procedures above cannot resolve the problem, causes other than data errors in the original boot disk unit may exist, or data errors may also exist in the mirror disk.

If it would appear that data errors also exist in the mirror disk, see "[\(4\) System cannot be booted. \(Boot disk data damage\)](#)."

(4) System cannot be booted. (Boot disk data damage)

Explanation

For some reason, the system disk data is invalid. You must restore data from the backup data created in advance. If restoring data from the mirror disk was unsuccessful, follow the procedures shown below.

Resolution

See "[6.1.3 Restoring \(When the System Cannot Be Booted\)](#)."

(5) System cannot be booted. (Incorrect boot disk device name)

Explanation

By mirroring system disks, the system's boot disk device name (boot-device parameter) is automatically set and updated. However, in some systems GDS cannot recognize correct device names and thus sets incorrect device names.

In that event, the following console message is output and booting the system fails.

```
ok boot
Rebooting with command: boot
Boot device: /ssm@0,0/pci@18,700000/pci@1/scsi@2/sd@0,0:a File and args:
Evaluating: boot
Can't open boot device
```

In this example the boot disk device name is "/ssm@0,0/pci@18,700000/pci@1/scsi@2/sd@0,0:a", but the correct device name is "/ssm@0,0/pci@18,700000/pci@1/scsi@2/disk@0,0:a" and thus booting the system fails.

Resolution

Create the bootdevtab file (boot disk device name file) for proper device name configuration.

1)Specify the correct boot device name and boot the system.

```
ok boot /ssm@0,0/pci@18,700000/pci@1/scsi@2/disk@0,0:a
```

For the correct boot disk device name, perform the following.

1-1) View device names of disks connected to the system.

For the following procedures, make a copy of the displayed results.

```
ok show-disks
a) /ssm@0,0/pci@18,700000/pci@1/scsi@2/disk
b) /ssm@0,0/pci@1c,700000/pci@1/scsi@2/disk
c) /ssm@0,0/pci@18,700000/pci@2/SUNW,isptwo@4/sd
```

```
d) /ssm@0,0/pci@1c,700000/pci@2/SUNW,isptwo@4/sd
q) NO SELECTION

Enter Selection, q to quit: q
ok
```

1-2) Based on a) in the displayed results, "/ssm@0,0/pci@18,700000/pci@1/scsi@2/sd@0,0:a" should be replaced with "/ssm@0,0/pci@18,700000/pci@1/scsi@2/disk@0,0:a" as that is the correct boot disk device name.

2) Create the bootdevtab file.

2-1) Create a template.

```
# sdxinfo -x BootDev > /tmp/bootdevtab
```

2-2) Edit the template.

```
# vi /tmp/bootdevtab
Before edit :
c0t0d0    /ssm@0,0/pci@18,700000/pci@1/scsi@2/sd@0,0
c1t0d0    /ssm@0,0/pci@1c,700000/pci@1/scsi@2/sd@0,0

After edit :
c0t0d0    /ssm@0,0/pci@18,700000/pci@1/scsi@2/disk@0,0
c1t0d0    /ssm@0,0/pci@1c,700000/pci@1/scsi@2/disk@0,0
```

Specify the physical disk name in the first field.

Specify the corresponding boot disk device name in the second field. Neither field requires partition assignment, such as s0 and :a.

To the template a boot disk device name assumed by GDS will be output. If the output boot disk name is incorrect, replace it with the correct device name. The correct device name can be checked as described in 1-1) through 1-2). Note that two boot disk device names may be output. In that event, delete one of the names and replace the remaining one with the correct device name.



Note

- Separate one field from another in space or tab delimited format. Lines starting with the number sign (#) are assumed as comments.
 - All disks registered with a root class are output to the template.
- Specify disks that may be used as the boot disk in the bootdevtab file. The relevant disks are:

- Original boot disk
- Mirror disk
- Spare disk connected to a boot disk's group
- Boot and spare disks in an alternative boot environment
(When using the system volume snapshot function of PRIMECLUSTER GDS Snapshot)

Do not specify the following disks in the bootdevtab file. Delete them from the bootdevtab file or comment out the relevant lines. Leaving them unedited will not affect the operation.

- Disks not used as the boot disk
- Disks for which GDS correctly assumed the boot disk device name and no assignment in the bootdevtab file is required
- When FC-AL internal disks are swapped, the boot disk device name changes.
Declare as follows so that recreation of the bootdevtab file is not required every time disks are swapped. Put the asterisk (*) after the last at mark (@). GDS will then complement the asterisk portion automatically.

```

Before edit :
c0t0d0 /ssm@0,0/pci@1d,700000/SUNW,qlc@1/fp@0,0/
ssd@w50020f230000798c,0

After edit :
c0t0d0 /ssm@0,0/pci@1d,700000/SUNW,qlc@1/fp@0,0/disk@*

```

- For MPLB disks, specify the boot disk unit names of their paths.

```

Before edit :
# mpla0

c12t16d0 /pci@1f,4000/fibre-channel@4/disk@10,0
c13t16d0 /pci@1f,4000/fibre-channel@5/disk@10,0

# mpla1

c12t17d0 /pci@1f,4000/fibre-channel@4/disk@11,0
c13t17d0 /pci@1f,4000/fibre-channel@5/disk@11,0

After edit :
# mpla0

c12t16d0 /pci@1f,4000/fibre-channel@4/sd@10,0
c13t16d0 /pci@1f,4000/fibre-channel@5/sd@10,0

# mpla1

c12t17d0 /pci@1f,4000/fibre-channel@4/sd@11,0
c13t17d0 /pci@1f,4000/fibre-channel@5/sd@11,0

```

- 2-3) Place the template in the /etc/opt/FJSVsdx/lib directory.

```

# cd /etc/opt/FJSVsdx/lib
# mv bootdevtab bootdevtab.old (when there is an existing bootdevtab
file)
# cp /tmp/bootdevtab bootdevtab
# chmod 600 bootdevtab

```

- 2-4) Confirm that the bootdevtab file was created normally.

```

# sdxinfo -x BootDev
c0t0d0 /ssm@0,0/pci@18,700000/pci@1/scsi@2/disk@0,0
c1t0d0 /ssm@0,0/pci@1c,700000/pci@1/scsi@2/disk@0,0

```

Check whether the boot disk device name is output normally.

- 3) Check whether the system's boot disk device name (boot-device parameter) is set normally.

Resolution

Restore the backed up file system.

For the restore methods, see "[6.1.2 Restoring \(When the System Can Be Booted\)](#)" or "[6.1.3 Restoring \(When the System Cannot Be Booted\)](#)."

Even if operation is performed as described in [Explanation], an error may not occur immediately.

Additionally, if the fsck(1M) command is used according to the messages shown in [Explanation], it may be possible to start the OS and it appears that the file system was restored. However, in these situations, file system data and the synchronized status of mirrors are not ensured.

When performing operation as described in [Explanation], be sure to restore the file system.

F.1.6 GDS Management View Abnormality

The following description shows resolutions to problems relevant to Web-Based Admin View and GDS Management View.

In one of the following circumstances, take action as indicated for the relevant situation.

- (1) [Web-Based Admin View abnormality](#)
- (2) [GDS Management View cannot be started.](#)
- (3) [Disk device information is not displayed correctly in GDS Management View.](#)

(1) Web-Based Admin View abnormality

Explanation

The following problems may occur in Web-Based Admin View.

- a. Web-Based Admin View top menu is not displayed. Or, even if it is, the screen is not displayed properly, or messages such as 0001 to 0050 are displayed.
- b. The screen does not change even after confirming 0002, 0003, or 0007 message. Or, the browser freezes when trying to reconnect to a previously-connected management server.
- c. The following message may be displayed on the management server or the Java console screen of the monitor node.
FJSVwvbs:OutOfMemoryError
- d. When trying to access the management server from browser, message 0005 appears and the top menu is not displayed, or the message itself may not appear.
- e. Web-Based Admin View stopped operating after changing the IP address of the public LAN.
- f. When using the PRIMECLUSTER Operation Management View, existing cluster nodes are not displayed on the each product's node list screen.
- g. The screen does not change even after confirming the dialog message.
- h. The following message is displayed on the browser status bar, and top menu is not displayed.
Applet com.fujitsu.webview.base.WebViewMain
error:java.lang.NoClassDefFoundError: com/sun/java/swing/Japplet
Or, the Operation Management View for each product cannot be started.
- i. It takes more than a few minutes before the Web-Based Admin View top menu is displayed. Or, even if it is, it freezes or message "0001 Cannot call Management server" may be displayed.
- j. When using a Java Plug-in on a client browser, and using a mouse on the Web-Based Admin View top menu or Operation Management View for each product, the screen has problems such as displaying small flecks in the area where the mouse has been used.
- k. When starting or restarting the Web-Based Admin View, or rebooting the node, the message "FJSVwvbs:webview.cnf abnormal" or "Node process abnormal end." is displayed, and Web-Based Admin View does not start.
- l. When starting or restarting the Web-Based Admin View, message "WARNING:unrelated 'http' property." is displayed.

- m. After installing the manual package, the title menu for the manual is not displayed. Or, even after removing the manual package, the title menu for the manual is still displayed.
- n. When trying to start the menu for products that are the target of management from the Web-Based Admin View top menu, the message "0023 This application must be started at the following URL" is displayed.
- o. Start button for products that are the target of management have disappeared from the Web-Based Admin View top menu.
- p. Even after disconnecting from the management server, messages 0002, 0003 or 0007 are not displayed on the client.
- q. When trying to access the management server from the browser, message 0006 appears, and the top menu is not displayed.
- r. When entering the user ID and the password for authentication, error message 0016 appears, and authentication fails.

Resolution

See "Web-Based Admin View Operation Guide" and follow the troubleshooting procedures.

(2) GDS Management View cannot be started.

Resolution

- a. When GDS is not installed.
Install GDS.
- b. When GDS is installed.
Check for the following possible causes:
 - Wrong URL specified.
 - Wrong method caused due to specifying host name.
 - Environment cannot use Swing.
 - Different browser type and level than recommended.

(3) Disk device information is not displayed correctly in GDS Management View.

Explanation

GDS does not recognize disks properly. Possible causes are as follows.

- The power of the disk unit was turned on after the system was booted.
- The disk device became unavailable for some kind of problem during system operation.
- The configuration of devices such as mphd and mplb was changed during system operation.

Resolution

Perform the following steps 1. through 2.

1. Use the format(1M) command and check whether the OS recognizes the disk unit. If it does not recognize, identify the cause based on log messages of disk drivers stored in the /var/adm/messages file and such.
2. Execute [Update Physical Disk Information] in the [Operation] menu.

F.1.7 Proxy Object Abnormality

For proxy object related errors, in one of the following circumstances, take action as indicated for the relevant situation.

- (1) [The Advanced Copy function cannot be used in master-proxy copying.](#)

(1) The Advanced Copy function cannot be used in master-proxy copying.

Explanation

The method handling copying from a master to a proxy can be checked in the CPTYPE field that is displayed with the `sdxinfo` command. When the CPTYPE field displays `soft`, instead of the Advanced Copy function of Fujitsu Storage Systems ETERNUS, the GDS soft copy function is performing the copy operation.

```
# sdxinfo -s -e long -o pvl
OBJ  NAME CLASS  GROUP  DISK  VOLUME  JRM  MODE  STATUS  COPY  CURBLKS  COPYBLKS  DLY  CPTYPE  CPSOURCE
-----
slice *   Class1 Group1 Disk1 pvl    *   *   COPY   run   93952 10027008  0  soft   *
```

The following eight possible causes can create a situation where the Advanced Copy function is not used.

(Cause a)

The disk array unit's Advanced Copy mechanism was not configured properly.

(Cause b)

After GDS Snapshot was installed, the node was not rebooted.

(Cause c)

EC or REC sessions exist between the master and another proxy and OPC or ROPC cannot be used.

(Cause d)

Multiple proxy volumes are related to a master volume, and the number of proxy volumes with EC or REC session settings has reached the upper limit (16 volumes).

(Cause e)

The mirroring configuration of proxy volumes is two-way or more multiplex mirroring.

(Cause f)

The group to which master volumes or proxy volumes belong has a lower level group connected.

(Cause g)

The disk array unit to which the master or the proxy belongs does not support the Advanced Copy function.

(Cause h)

An error occurred in the disk array unit or the Fibre Channel route.

For details, see "[A.2.18 Using the Advanced Copy Function in a Proxy Configuration.](#)"

Resolution

Check the disk array settings and the GDS object configuration and identify the cause, and take the following actions according to the cause.

- In the event of (Cause a), set up the disk array unit's Advanced Copy feature.
- In the event of (Cause b), reboot all nodes in the cluster system.
- In the event of (Cause c) or (Cause d), take the following action (Resolution 1) or (Resolution 2).

(Resolution 1)

1. Part a proxy with EC or REC session settings from the target master.
2. Cancel EC or REC sessions of step 1. using the `sdxproxy Cancel` command.

(Resolution 2)

1. Cancel the relationship between the target master and a proxy with EC or REC session settings.
- In the even of (Cause e) to (Cause g), according to need, select another object, or change the object configuration.

- In the event of (Cause h), recover the disk array unit error or the Fibre Channel route error.

F.1.8 EMC Symmetrix Abnormality

For EMC Symmetrix related errors, in one of the following circumstances, take action as indicated for the relevant situation.

- (1) SRDF is not used in master-to- proxy synchronization copying.
- (2) The joined proxy volume is in INVALID status.

(1) SRDF is not used in master-to- proxy synchronization copying.

Explanation

The method handling master-to-proxy copying can be checked in the CPTYPE field that is displayed with the `sdxinfo` command. When the CPTYPE field displays `soft` without displaying SRDF, instead of SRDF, the GDS soft copy function is performing the copy operation.

```
# sdxinfo -s -e long -o pvl
OBJ  NAME  CLASS  GROUP  DISK  VOLUME  JRM  MODE  STATUS  COPY  CURBLKS  COPYBLKS  DLY  CPTYPE  CPSOURCE
-----
slice *      Class1 Group1 Disk1 pvl      *   *   COPY   run   93952 10027008  0 soft  *
```

The following five possible causes can create a situation where SRDF is not used.

(Cause a)

The configuration of the master or proxy group is inappropriate.

The source (R1) and target (R2) devices as the SRDF pair must be connected with the master and proxy groups respectively. With the proxy group, the target (R2) device only must be connected.

(Cause b)

A pair of volumes, instead of a pair of groups, was related as the master and the proxy.

(Cause c)

When joining the master and the proxy groups, the specific was set to create proxy volumes to the proxy group with the physical slice attributes that are different from those of the master volumes.

(Cause d)

Master data was restored using the proxy with the `sdxproxy RejoinRestore` command.

When the `sdxproxy RejoinRestore` command performs master restoration, the SRDF pair is cancelled and SRDF becomes no longer available.

(Cause e)

With the `sdxproxy Join` command or the `sdxproxy Rejoin` command, the `-e` option was specified to use the soft copy function.

When the `-e softcopy` option is specified with the `sdxproxy Join` command, SRDF is not used.

When the `-e softcopy` option is specified with the `sdxproxy Rejoin` command, the SRDF pair is cancelled and SRDF becomes no longer available.

For details, see "[A.2.21 Using EMC TimeFinder or EMC SRDF in a Proxy Configuration.](#)"

Resolution

1) Check whether the possible cause is (Cause a).

1-1) Identify the physical disk names of the disks connected with the master and the proxy groups.


```
# sdxinfo -D -o mg1
OBJ      NAME      TYPE      CLASS     GROUP     DEVNAM      DEVBLKS  DEVCONNECT  STATUS
-----
disk     Disk1     mirror   Class1    mg1       emcpower10  17596416  node1:node2  ENABLE
disk     Disk2     mirror   Class1    mg1       emcpower11  17682084  node1:node2  ENABLE

# sdxinfo -D -o pg1
OBJ      NAME      TYPE      CLASS     GROUP     DEVNAM      DEVBLKS  DEVCONNECT  STATUS
-----
disk     Disk3     mirror   Class1    pg1       emcpower20  17596416  node1:node2  ENABLE
```

In this example, the physical disks emcpower10 and emcpower11 have been connected with the master group mg1, and the physical disk emcpower20 has been connected with the proxy group pg1.

If a lower level group is, or multiple disks or lower level groups are, connected with the proxy group, the possible cause is (Cause a).

1-2) Identify the Symmetrix IDs and the Symmetrix device names of the disks connected with the master and the proxy groups using the sympd command provided by SYMCLI.

```
# sympd list
Symmetrix ID: 000183600262

      Device Name          Directors          Device
-----
Physical          Sym SA :P DA :IT Config          Attribute      Sts      Cap
-----
...
/dev/rdisk/emcpower10c  000 03A:0 02B:C0 RDF1          Grp'd          RW      8592
/dev/rdisk/emcpower11c  001 03A:0 02B:D0 Unprotected  N/Grp'd        RW      8632
/dev/rdisk/emcpower12c  002 03A:0 01B:D0 RDF2          Grp'd          RW      8632
...

Symmetrix ID: 000282600920

      Device Name          Directors          Device
-----
Physical          Sym SA :P DA :IT Config          Attribute      Sts      Cap
-----
...
/dev/rdisk/emcpower20c  005 03A:0 02B:C0 RDF2          Grp'd          RW      8592
/dev/rdisk/emcpower21c  006 03A:0 01B:D0 RDF1          Grp'd          RW      8632
/dev/rdisk/emcpower22c  007 03A:0 01A:D0 RDF1          Grp'd          RW      8632
...
```

In this example, as to the physical disks emcpower10c and emcpower11c, the Symmetrix ID is 000183600262 and the Symmetrix device names are 000 and 001 respectively. As for the physical disk emcpower20c, the Symmetrix ID is 000282600920 and the Symmetrix device name is 005.



For details on the sympd command, see the Symmetrix Manager manual.

1-3) Check the source (R1) device and the target (R2) device as the SRDF pair whether they are connected with the master and proxy groups respectively using the symrdf command provided by SYMCLI.

The SymDev field and the RDev field of the symrdf list command output display the Symmetrix device names of the source (R1) device and the corresponding target (R2) device respectively.

```

# symrdf list
Symmetrix ID: 000183600262

Local Device View
-----
                STATUS   M O D E S                RDF   S T A T E S
Sym           RDF  -----
Dev  RDev  Typ:G SA RA LNK Mode Dom ACp  Tracks  Tracks Dev RDev Pair
-----
...
000  005   R1:1 RW RW NR  SYN  DIS OFF      0      0 RW  RW  Synchronized
002  006   R2:2 RW RW NR  SYN  DIS OFF      0      0 RW  RW  Split
...

Symmetrix ID: 000282600920

Local Device View
-----
                STATUS   M O D E S                RDF   S T A T E S
Sym           RDF  -----
Dev  RDev  Typ:G SA RA LNK Mode Dom ACp  Tracks  Tracks Dev RDev Pair
-----
...
005  000   R2:1 RW RW NR  SYN  DIS OFF      0      0 RW  RW  Synchronized
006  002   R1:2 RW RW NR  SYN  DIS OFF      0      0 RW  RW  Split
...

```

In this example, the source (R1) device 000 in the Symmetrix frame 000183600262 and the target (R2) device 005 in the Symmetrix frame 000282600920 are the SRDF pair. The master and the proxy have been configured properly, and the possible cause is not (Cause e).



For details on the symrdf command, see the Symmetrix Manager manual.

1-4) If the possible cause is (Cause e), release the master-proxy relationship once, and re-configure properly.

For example, according to the symrdf command results in step 1-3), you will see that the source (R1) device 006 in the Symmetrix frame 000282600920 and the target (R2) device in the Symmetrix frame 000183600262 are also the SRDF pair. The symrdf command results in step 1-2) are indicating that the physical disk names of the Symmetrix devices 006 and 002 are emcpower21 and emcpower12 respectively. The proxy configuration associated with SRDF can be created by registering emcpower21 and emcpower12 with a class, connecting them with respective groups, and relating the groups as the master and proxy.

In addition, the registration of the source (R1) device with a device group can be done after registering the disk with a class, but it must be carried out before joining the master and the proxy.

2) Check whether the possible cause is (Cause b).

```
# sdxinfo -GV -e long -o pv1
```

OBJ	NAME	CLASS	DISKS	BLKS	FREEBLKS	SPARE	MASTER	TYPE	WIDTH
group	pg1	Class1	Disk3	17596416	7503872	0	*	mirror	*

OBJ	NAME	TYPE	CLASS	GROUP	DISK	MASTER	PROXY	SKIP	JRM	MODE	CMODE	LOCK	1STBLK	LASTBLK	BLOCKS	STATUS	PSLICE	SNUM
volume	*	mirror	Class1	pg1	*	*	*	*	*	*	*	*	0	65535	65536	PRIVATE	*	*
volume	pv1	mirror	Class1	pg1	*	mv1	*	off	on	ro	*	on	65536	10092543	10027008	STOP	on	1
volume	*	mirror	Class1	pg1	*	*	*	*	*	*	*	*	10092544	17596415	7503872	FREE	*	*

In this example, the MASTER field of the group pg1 to which the proxy volume pv1 belongs displays the asterisk (*) and the MASTER field of pv1 displays the master volume name mv1, and the possible cause is (Cause b).

If the possible cause is (Cause b), the action as follows should be implemented.

2-1) Break the relationship between the master and proxy volumes.

```
# sdxproxy Break -c Class1 -p pv1
```

2-2) Join a pair of groups instead of a pair of volumes as the master and the proxy. In this example, the master group name is mg1.

```
# sdxproxy Join -c Class1 -m mg1 -p pg1 -a mv1=pv1:on
```

3) Check whether the possible cause is (Cause c).

```
# sdxinfo -V -e long -o mg1
```

OBJ	NAME	TYPE	CLASS	GROUP	DISK	MASTER	PROXY	SKIP	JRM	MODE	CMODE	LOCK	1STBLK	LASTBLK	BLOCKS	STATUS	PSLICE	SNUM
volume	*	mirror	Class1	mg1	*	*	*	*	*	*	*	*	0	65535	65536	PRIVATE	*	*
volume	mv1	mirror	Class1	mg1	*	*	*	off	on	rw	rw	off	65536	10092543	10027008	ACTIVE	on	1
volume	*	mirror	Class1	mg1	*	*	*	*	*	*	*	*	10092544	17596415	7503872	FREE	*	*


```
# sdxinfo -V -e long -o pg1
```

OBJ	NAME	TYPE	CLASS	GROUP	DISK	MASTER	PROXY	SKIP	JRM	MODE	CMODE	LOCK	1STBLK	LASTBLK	BLOCKS	STATUS	PSLICE	SNUM
volume	*	mirror	Class1	pg1	*	*	*	*	*	*	*	*	0	65535	65536	PRIVATE	*	*
volume	pv1	mirror	Class1	pg1	*	mv1	*	off	on	ro	*	on	65536	10092543	10027008	STOP	off	1
volume	*	mirror	Class1	pg1	*	*	*	*	*	*	*	*	10092544	17596415	7503872	FREE	*	*

In this example, the attributes of the physical slices of the master volume mv1 and the proxy volume pv1 are on and off respectively, and the possible cause is (Cause c).

If the possible cause is (Cause c), the action as follows should be implemented.

3-1) Break the master-proxy relationship.

```
# sdxproxy Break -c Class1 -p pg1
```

3-2) Relate the master and the proxy again. With the -a option for the sdxproxy Join command, if the physical slice attribute is not specified, the physical slice attribute of a created proxy volume conforms to that of the master volume.

```
# sdxproxy Join -c Class1 -m mg1 -p pg1 -a mv1=pv1:on
```

4) Whether the possible cause is (Cause d) or (Cause e) can be determined by checking the command action history referencing the GDS log file /etc/opt/FJSVsdx/msglog/sdxservd.log.

If the possible cause is (Cause d) or (Cause e), the master and the proxy should be disjoined once and then rejoined.

(2) The joined proxy volume is in INVALID status.

Explanation

TimeFinder and SRDF may stop operating because the exclusive lock on a Symmetrix device is not released due to an event such as a node down. In this event, copying by TimeFinder or SRDF does not operate normally, and the proxy volume status may become INVALID.

The exclusive locks for Symmetrix devices are divided into two types. Whether respective exclusive locks are left not released can be checked using the symdev and symcfg commands provided by SYMCLI.

```
# symdev -lock list
# symcfg -lock list
```

If an exclusive lock has remained, the lock number is displayed.

Resolution

Release the exclusive lock on the Symmetrix device with the following method.

As to an exclusive lock indicated by the symdev -lock list command, use the symdev command to release it. Substitute the lock number displayed with the symdev -lock list command for *number*.

```
# symdev -lock number release
```

As to an exclusive lock indicated by the symcfg -lock list command, use the symcfg command to release it. Substitute the lock number displayed with the symcfg -lock list command for *number*.

```
# symcfg -lockn number release
```



For details on the symdev command and the symcfg command, see the Symmetrix Manager manual.

F.1.9 Cluster System Related Error

For cluster system related errors, in one of the following circumstances, take action as indicated for the relevant situation.

- (1) The error message "ERROR: class: cannot operate in cluster environment, ..." is output, and the operation cannot be conducted on the class class.
- (2) The PRIMECLUSTER CF clinitreset(1M) command ends abnormally outputting an error message # 6675.
- (3) Cluster applications become "Inconsistent".
- (4) The GFS Shared File System is not mounted on node startup.

(1) The error message "ERROR: class: cannot operate in cluster environment, ..." is output, and the operation cannot be conducted on the class class.

Explanation

The local class created when cluster control facility was inactive cannot directly be used in a cluster system. When the cluster control facility is activated, the following message is output to the system log and the GDS daemon log file, and the local class becomes nonoperational.

ERROR: *class*: cannot operate in cluster environment, created when cluster control facility not ready

This error message will be output when:

- The cluster initial configuration was executed after the local class *class* had been created on a node on which that configuration was incomplete.
- The local class *class* was created in single user mode.
- The single node on which the local class *class* was created was changed over to a cluster system.

Resolution

Make the local class available in a cluster system by method a) or method b) as follows. Generally method a) should be used, but method b) should be used in order to prevent the volume data from being backed up and restored.

Method a) Re-creating the local class in the cluster system:

- 1) Activate the node in single user mode.
- 2) Back up volume data if necessary.
- 3) Delete the class.
- 4) Re-activate the cluster control facility on the node booted in multi-user mode.
- 5) Re-create the class and volumes deleted in step 3).
- 6) Restore the volume data backed up in step 2) as needed.

Method b) Converting the local class to one for a cluster system:

Convert the local class to a class for a cluster system taking the following procedures. The following illustrates the procedures when the class name is Class1.

Note

You need to delete and re-create the classes which were created while the cluster control facility was active.

- 1) Confirm the classes to be converted or re-created for each node.

```
# sdxinfo -C
OBJ   NAME   TYPE   SCOPE   SPARE
-----
class Class1 local   (local) 0
class Class2 local   Node1    0
class Class3 shared  Node1:Node2 0
```

- Classes which should be converted

Classes whose TYPE field is "local" and SCOPE field is "(local)."

In the above example, Class 1 meets the definition.

- Classes which should be re-created

Classes whose SCOPE field properly indicates a class scope (node identifier).

In the above example, Class2 and Class3 meet the definition.

2) Back up the volume data of the class to be re-created as needed.

3) Delete the class to be re-created.

4) Boot the node, on which the local class Class1 to be converted exists, in single user mode.

```
ok boot -s
...
INIT:SINGLE USER MODE
Type control-d to proceed with normal startup,
(or give root password for system maintenance): password
```

5) Stop the GDS management daemon, sdxservd, on the node where the local class Class1 exists.

```
# /etc/opt/FJSVsdx/bin/sdx_stop -s
sfdsk: received shutdown request
sfdsk: volume status log updated successfully, class=0x40000004
#
```

Confirm that the sdxservd daemon was stopped (information on sdxservd daemon processes is not displayed) in the following manner.

```
# ps -e | grep sdxservd
#
```

6) Back up the configuration database on the node where the local class Class1 exists.

```
# rm -rf /var/opt/FJSVsdx/backup/DB/Class1
# /etc/opt/FJSVsdx/bin/sdxc1trandb -B -c Class1
sdxsavedb: INFO: /dev/rdisk/c0t1d0s0: backup succeeded
sdxsavedb: INFO: /dev/rdisk/c1t1d0s0: backup succeeded
sdxsavedb: INFO: Class1: backup succeeded

# cd /var/opt/FJSVsdx/backup/DB/Class1
# ls -l
-rw-r--r--  1 root    other   14164992 May  6 09:00 c0t1d0s0
-rw-r--r--  1 root    other   14164992 May  6 09:00 c1t1d0s0
```

Note

- Verify that there is free space equal to or larger than 150 [MB] under /var/opt/FJSVsdx/backup/DB, and if it insufficient expand it.
- If an error occurs, perform the following procedure instead of proceeding to the subsequent steps:
After re-creating the local class Class1 according to method a), re-create the class deleted in step 3) and restore data.

7) On the node where the local class Class1 exists, convert the configuration database for Class1 to that for a cluster system.

```
# /etc/opt/FJSVsdx/bin/sdxcltrandb -C -c Class1
sdconvertdb: INFO: /dev/rdisk/c0t1d0s0: conversion succeeded
sdconvertdb: INFO: /dev/rdisk/c1t1d0s0: conversion succeeded
sdconvertdb: INFO: Class1: conversion succeeded
```

Note

If an error occurs, follow the steps from 10-3) to restore the configuration database, and re-create the local class Class1 according to method a). Re-create the class deleted in step 3) and restore data.

8) Re-activate the cluster control facility by rebooting the node, on which the local class Class1 exists, in multi-user mode.

```
# init 0
...
SDX:sdxshutdown: ERROR: connection timeout
...
ok boot
...
Console Login:
```

Note

The following messages are output during shutdown, but there are no problems.

```
SDX:sdxshutdown: INFO: waiting for a response from sdxservd daemon...
SDX:sdxshutdown: ERROR: connection timeout
```

9) On the node where the local class Class1 exists, verify that the configuration database for the local class Class1 was converted successfully.

```
# sdxinfo -C -c Class1
OBJ      NAME      TYPE      SCOPE      SPARE
-----
class   Class1   local     Node1      0
```

Confirm that the node identifier is displayed properly in the SCOPE field. If it is displayed properly, the process is finished.

Note

If the SCOPE field statement is improper, it means that the Class1 configuration database was not converted successfully. If that happens, you should restore the configuration database taking steps from 10-1) onward and re-create the local class Class1 according to method a). Re-create the class deleted in step 3) and restore data.

10) Restore the configuration database backed up in step 6) when an error occurred in step 7) or 9).

Perform the following steps on the node where the local class Class1 exists.

10-1) Activate the node in the single user mode.

```
ok boot -s
...
INIT:SINGLE USER MODE
Type control-d to proceed with normal startup,
(or give root password for system maintenance): password
```

10-2) Stop the GDS management daemon, sdxservd.

```
# /etc/opt/FJSVsdx/bin/sdx_stop -S
sfdsk: received shutdown request
sfdsk: volume status log updated successfully, class=0x40000004
#
```

Confirm that the sdxservd daemon was stopped (information on sdxservd daemon processes is not displayed) in the following manner.

```
# ps -e | grep sdxservd
#
```

10-3) Restore the configuration database for the local class Class1.

```
# /etc/opt/FJSVsdx/bin/sdxcltrandb -R -c Class1
sdxrestoredb: INFO: /dev/rdsk/c0t1d0s0: restore succeeded
sdxrestoredb: INFO: /dev/rdsk/clt1d0s0: restore succeeded
sdxrestoredb: INFO: Class1: restore succeeded
```

10-4) Re-activate the node in single user mode.

```
# init 0
...
SDX:sdxshutdown: ERROR: connection timeout
...
ok boot -s
...
INIT:SINGLE USER MODE
Type control-d to proceed with normal startup,
(or give root password for system maintenance): password
```



Note

The following messages are output during shutdown, but there are no problems.

```
SDX:sdxshutdown: INFO: waiting for a response from sdxservd daemon...
SDX:sdxshutdown: ERROR: connection timeout
```

10-5) Verify that the configuration database for the local class Class1 was restored normally.

```
# sdxinfo -C -c Class1
OBJ    NAME    TYPE    SCOPE    SPARE
-----
class  Class1  local   Node1    0
```

Confirm that the node identifier is displayed properly in the SCOPE field. If it is displayed properly, the restoration is finished.

11) After converting the local class, re-create the class deleted in step 3).

Then restore the volume data backed up in step 2) as needed.

(2) The PRIMECLUSTER CF `clinitreset(1M)` command ends abnormally outputting an error message # 6675.

Explanation

When a class exists in a cluster system, initializing the PRIMECLUSTER resource database with the PRIMECLUSTER CF `clinitreset` command results in that the `clinitreset` command fails outputting the following error message.

FJSVcluster: ERROR: `clinitreset: 6675: Cannot run this command because Global Disk Services has already been set up.`

When a node containing a shadow class is rebooted because of an event such as shutdown or panic, the shadow class is deleted, but the `/dev/sfdsk/Class Name` directory is not deleted. If the `clinitreset` command is executed here, the command also fails outputting the error message as above.

Resolution

1. On all nodes in the cluster system, view the configuration of objects and delete a class if any exists. If a class is deleted, volume data will be lost. If necessary, back up volume data in advance.



- For using GDS Management View, see "[5.5 Removals](#)."

- For using commands, see "[Appendix D Command Reference](#)."

2. On all nodes in the cluster system, check whether a class directory exists in the `/dev/sfdsk` directory, and delete a class directory if any exists. The following shows an example when a directory of class `Class1` exists. `_adm` and `_diag` are special files used by GDS and cannot be deleted.

```
# cd /dev/sfdsk
# ls
_adm  _diag  Class1

# rm -rf Class1
```

(3) Cluster applications become "Inconsistent".

Explanation

If a shared class is not to be used as an RMS resource, volumes included in the class are started on node startup. If a cluster application that uses those volumes are started there, the cluster application becomes "Inconsistent" because the volumes are already active. By default, classes are not to be used as RMS resources. Classes can be made available as RMS resources either by:

- Registering them in resources used by cluster applications through the Web-Based Admin View's userApplication Configuration Wizard
- Specifying them and using the `hvgdsetup -a` command

Resolution

Make the shared class available as an RMS resource with one of the following methods. After performing the procedures, restart the cluster application.

- If the class is not registered in resources used by the cluster application, register it through the userApplication Configuration Wizard.
- Execute the following command.

```
# /usr/opt/reliant/bin/hvgdsetup -a Class Name
...
Do you want to continue with these processes ? y
```

(4) The GFS Shared File System is not mounted on node startup.

Explanation

If a shared class is to be used as an RMS resource, volumes included in the class are not started on node startup. Therefore, the GFS Shared File System on those volumes is not mounted on node startup. By default, classes are not to be used as RMS resources, but they are made available as RMS resources either by:

- Registering them in resources used by cluster applications through the Web-Based Admin View's userApplication Configuration Wizard
- Specifying them and using the hvgdsetup -a command

Resolution

Take one of the following actions.

a) When using the shared class as an RMS resource, do not create the GFS Shared File System on volumes in the class, but create it on volumes in a difference class.

b) When not using the shared class as an RMS resource, make the class unavailable as an RMS resource again with one of the following methods. After performing the procedures, reboot the system.

- If the class is registered in resources used by the cluster application, remove it through the userApplication Configuration Wizard.
- If the class is not registered in resources used by the cluster application, execute the following command.

```
# /usr/opt/reliant/bin/hvgdsetup -d Class Name
...
Do you want to continue with these processes ? y
...
Do you need to start volumes in the specified disk class ? n
```

F.1.10 Shadow Object Errors

For shadow object related errors, in one of the following circumstances, take action as indicated for the relevant situation.

- (1) [A shadow volume became inaccessible.](#)

(1) A shadow volume became inaccessible.

Explanation

The configuration information of a shadow volume is not saved on the private slice, but managed in the memory. For this reason, the shadow volume configuration is cleared when the node on which the shadow volume is defined is rebooted, and thus the shadow volume becomes inaccessible.

Resolution

Delete the device special file according to the procedure described in "[Rebooting a Node](#)" in "[A.2.28 Shadow Volume](#)," and re-create a shadow volume.

F.1.11 Disk Unit Error



Note

Identifying a Failed Disk Unit

GDS Management View shows disks with I/O errors by indicating them with red disk icons. In addition, if an I/O error occurs, 1 (one) is given as a value in the E field displayed with the `sdxinfo -e` long command.

However, for hardware-related errors, identify the failed parts based on log messages of disk drivers stored in the `/var/adm/messages` file and such.

If GDS Management View and the `sdxinfo` command indicate the invalid status of slice data (invalid or INVALID), the slice data is invalid. It does not necessarily indicate that that slice has a hardware error.

For example, if an I/O error occurs on the copy source slice due to a hardware error during the synchronization copying of a mirror volume, data on the copy destination slice becomes invalid, but the copy source slice is still indicated as normal (active, stop, ACTIVE, or STOP). For details, see "F.1.1 Slice Status Abnormality."



See

For the procedures for swapping disks, see "1.2.3 Hot Swap," "5.3.4 Disk Swap," "D.8 sdxswap - Swap disk," "A.2.31 Swapping Internal Disks (FC-AL)," and "A.2.32 Swapping Internal Disks (SPARC T3-1/T3-2/T3-4)."

F.2 Collecting Investigation Material

When a problem occurs with GDS that cannot be resolved by referring to this Troubleshooting appendix, collect the investigation material and contact your local customer support.



Note

- Investigation material must be collected promptly. With the passage of time, necessary information may be lost.
- In a cluster system, collect investigation material on all the nodes.

F.2.1 Collecting with `fjsnap` Command

The `fjsnap` command, a system information collection tool, is provided in the `FJSVsnap` package of the Enhanced Support Facility. When a problem occurs with GDS, this command allows you to collect investigation material to identify the cause.

How to execute the `fjsnap` command is as follows.

1. Log in with superuser authority.
2. Execute the `fjsnap` command as follows.

```
# /opt/FJSVsnap/bin/fjsnap -h output
```

or

```
# /opt/FJSVsnap/bin/fjsnap -a output
```

output indicates either the special file name (e.g. `/dev/rmt/0`) of the output media which stores the investigation material collected with the `fjsnap` command, or the output file name.



See

For details on the `fjsnap` command, see the `README.ja` file and `README` file included in the `FJSVsnap` package, and Enhanced Support Facility user guide.

F.2.2 Collecting Initial Investigation Material with sdxsnap.sh Script

Since the amount of information collected with the `fjsnap` command is rather large, it is often required to be sent to us by postal mail, and not by e-mail, which could take a few days before our engineer can start investigating.

In order to promptly start the initial investigation, you can only collect the essential information with the `sdxsnap.sh` script. If you wish to reduce the time lapse before the actual investigation starts, collect information using both the `sdxsnap.sh` script and the `fjsnap` command, and e-mail this initial investigation material first.

Execute the `sdxsnap.sh` script by following the procedures below.

1. Login with the superuser authority.
2. Execute the `sdxsnap.sh` script.

```
# /etc/opt/FJSVsdx/bin/sdxsnap.sh -s [-d dir]
# cd dir
# tar cvf snap.tar SDXSNAP
```

Use the `-d` option to save the material in the `dir` directory. `dir` must be specified with an absolute path. If the `-d` option is not used, the material will be saved in the `/tmp` directory.

Send us the `snap.tar` file, which will be used as initial investigation material.

Appendix G Frequently Asked Questions (FAQ)

In this section, frequently asked questions (FAQ) have been compiled.

G.1 Operation Design

Questions regarding system operation design are explained.

- Q: Is operation possible without using GUI (GDS Management View)?
- Q: Is it possible to perform GDS Snapshot operations through the GUI (GDS Management View)?
- Q: Can disks with different performance specifications, such as size and revolution speed, be mirrored?
- Q: Do disks have to have equivalent attributes, such as size and revolution speed when I relate them as master and proxy?
- Q: Can I combine mirroring and striping?
- Q: Do I need to specify certain settings in order to perform a hot swap?
- Q: What is the purpose of single volume?
- Q: Can the messages output by GDS driver, daemon and command be displayed in a language other than English?
- Q: Are there any restrictions on the LUN (Logical Unit) configuration within the disk array when managing the disk array using GDS?
- Q: Are the mirrored disks differentiated as "primary" and "secondary" disks?
- Q: When reading the mirror volume, which one of the slices in the volume will the I/O request be issued to?
- Q: When mirroring a disk array that supports the hot spare function, which is preferable; using the spare disk of the disk array, or that of GDS?
- Q: Which of the GDS messages should be monitored?
- Q: I want to know GDS resident daemons to monitor resident processes.
- Q: When configuring a 24-hour operation system, would files of incremental size such as message log file have any effect?

Q: Is operation possible without using GUI (GDS Management View)?

A

Yes. Two types of operation management interfaces are available: the GUI (Graphical User Interface) and the CLI (Command Line Interface).

Q: Is it possible to perform GDS Snapshot operations through the GUI (GDS Management View)?

A

Proxy object operations are possible. To operate shadow objects, use the CLI (Command Line Interface).

Q: Can disks with different performance specifications, such as size and revolution speed, be mirrored?

A

Yes. However, when mirroring disks of different sizes, you will only be able to use the capacity of the smaller disk. For example, if a 4 GB disk and a 9 GB disk are mirrored, only 4 GB out of the 9 GB disk will be accessible.

Also, when mirroring disks with different specifications such as revolution speed, the read performance becomes unbalanced and the write performance will depend on the slower disk's performance. Therefore, it is recommended to mirror disks with the same specifications.

See "[A.3.1 Guidelines for Mirroring](#)" also.

Q: Do disks have to have equivalent attributes, such as size and revolution speed when I relate them as master and proxy?

A

You can relate disks as master and proxy even if they have different attributes, such as size and revolution speed.

However, it will result in that the read performance becomes unequal and the write performance depends on the slower disk's performance. Therefore, it is recommended to relate disks with the same attributes wherever possible.

For details see also "[A.1.8 Proxy Configuration Preconditions.](#)"

Q: Can I combine mirroring and striping?

A

You can mirror stripe groups.

However, you cannot stripe across multiple mirror groups, using them as stripe columns.

Q: Do I need to specify certain settings in order to perform a hot swap?

A

No.

As long as the disk you are using supports hot swap, no special settings are required.

For procedures, see "[5.3.4 Disk Swap](#)" and "[D.8 sdxswap - Swap disk.](#)"

Q: What is the purpose of single volume?

A

GDS is software which improves availability and manageability of the disk data. Single volume is used when users need to improve manageability, rather than availability. For details, see "[1.3.2 Centralized Disk Management.](#)"

For example, if availability is already realized by using disk arrays, the user will benefit by using the physical disk as a single volume excluded from mirroring, to improve its manageability.

Using GDS to manage the volume as a single volume not only ensures consistent management, but also allows you to include the volume in the mirroring configuration without stopping the service application, when the need arises.

Q: Can the messages output by GDS driver, daemon and command be displayed in a language other than English?

A

No. The messages are in English only.

Q: Are there any restrictions on the LUN (Logical Unit) configuration within the disk array when managing the disk array using GDS?

A

There are no restrictions on LUN RAID level or size.

Q: Are the mirrored disks differentiated as "primary" and "secondary" disks?

A

No. Disks mirrored with GDS are handled in the same way. Both the user and the application do not differentiate between the two.

However, the only exception is the system disk. The system disks may be differentiated as the primary and the secondary disks to initially be booted with OpenBoot. For details, see "[F.1.5 System Disk Abnormality.](#)"

Q: When reading the mirror volume, which one of the slices in the volume will the I/O request be issued to?

A

When a read is issued to a volume, the last access block number of each disk in the volume will be checked. The slice with the value closest to the read block number will be read. In other words, disk with the shortest seek distance will be read first.

When a write is issued to a volume, it will write to all slices in the volume (except from slice in ACTIVE status) and will return the write results after completing the write request on all slices.

Q: When mirroring a disk array that supports the hot spare function, which is preferable; using the spare disk of the disk array, or that of GDS?

A

Generally, using the spare disk of the disk array is recommended.

Q: Which of the GDS messages should be monitored?

A

Generally, the following messages given in the "GDS Messages" should be monitored.

- PANIC or WARNING level messages and Internal error messages shown in "[E.2 Driver Messages](#)"
- HALT, ERROR, or WARNING level messages and Internal error messages shown in "[E.3 Daemon Messages](#)"

For details, see "[Appendix E GDS Messages](#)."

Q: I want to know GDS resident daemons to monitor resident processes.

A

The following table contains GDS resident daemons and the description.

Name	ps -ef command's CMD	Description
sdxmond	/usr/sbin/sdxmond	Monitors GDS daemons
sdxservd	/usr/sbin/sdxservd	Required daemon
sdxlogd	/usr/sbin/sdxlogd	Required daemon
sdxexd	/usr/sbin/sdxexd	Required daemon
sdxcld	/usr/sbin/sdxcld -W	Cooperates with PRIMECLUSTER CF (Only in a cluster system)
sdxcld	/usr/sbin/sdxcld -W	Cooperates with PRIMECLUSTER CF (Only in a cluster system)
sdxcle	/usr/sbin/sdxcle -f <i>n</i>	Cooperates with PRIMECLUSTER CF (Only in a cluster system) <i>n</i> is a natural number.

If any resident daemon other than sdxmond terminates abnormally, sdxmond automatically restarts the daemon and outputs a message as follows to the GDS log file /var/opt/FJVSdx/msglog/sdxservd.log.

When succeeded in restart

```
SDX:sdxmond: WARNING: respawned daemon daemon successfully
```

When failed in restart

```
SDX:sdxmond: HALT: failed to respawn daemon, osfunc=osfunc,
errno=errno
```

If sdxmond terminates abnormally, the OS init(1M) process automatically restarts the daemon. In that event, no message is output.

As described above, the GDS resident daemons are automatically restarted in the event of abnormal termination, and it is not necessary to monitor them.

In addition, GDS Snapshot has no resident daemon.



See

.....
For the meanings of the messages and the resolutions, see "[E.3 Daemon Messages](#)."
.....

Q: When configuring a 24-hour operation system, would files of incremental size such as message log file have any effect?

A

Since GDS's message log file is managed so that it will not exceed a certain size, there are no effects that need to be considered.

G.2 Environment Configuration

Questions regarding environment configuration are explained.

- [Q: How do I make a backup of a system disk that is managed with GDS?](#)
- [Q: How do I make a backup of a disk \(other than a system disk\) that is managed with GDS?](#)
- [Q: The following message is output on the console after mirroring the system disk with GDS. Is there a problem? What is the resolution method?](#)
- [Q: I cannot execute the format\(1M\) command or prtvtoc\(1M\) on a physical disk managed by GDS \(e.g. c0t0d0s0\). How can I solve this problem?](#)
- [Q: Can I install OS upgrades while system disks are in mirroring configuration through GDS?](#)

Q: How do I make a backup of a system disk that is managed with GDS?

A

For procedures on how to make backups of file systems /(root), /usr and /var, which cannot be unmounted while the system is operating, see "[6.1 Backing Up and Restoring a System Disk](#)" and "[6.2 Backing Up and Restoring a System Disk through an Alternative Boot Environment](#)."

Q: How do I make a backup of a disk (other than a system disk) that is managed with GDS?

A

The backup of GDS's logical volume can be made in the same way as a physical slice. However, make sure you make a backup using the logical volume. For example, do not boot from the CD-ROM device or network and make a backup using the physical slice, since correct data may not be saved.

You can also efficiently create a backup with the mirroring function. For details, see "[1.3.8 Snapshots by Slice Detachment](#)" "[D.5 sdxslice - Slice operations](#)", "[5.3.2 Backup](#)" and "[6.3 Backing Up and Restoring Local Disks and Shared Disks](#)."

By making use of GDS Snapshot proxy volumes, you can also configure backup administration that minimizes operation stop time. For details, see "[1.5 Proxy Volume](#)," "[D.14 sdxproxy - Proxy object operations](#)," "[D.19 Snapshot Creation Using Command](#)," "[5.3.2.2 Backup \(by Synchronization\)](#)," "[5.3.2.3 Backup \(by OPC\)](#)" and "[6.4 Online Backup and Instant Restore through Proxy Volume](#)."

Furthermore, backing up data through a server outside the domain running the service using GDS snapshot shadow volumes can hold the effects on the service to a minimum. For details, see "[1.6 Shadow Volume](#)" and "[6.6 Backing Up and Restoring through an External Server](#)."

Q: The following message is output on the console after mirroring the system disk with GDS. Is there a problem? What is the resolution method?

```
/usr/sbin/pmconfig: /etc/power.conf line (30) failed to convert mount point /dev/sfdsk/System/dsk/root to prom name: Invalid argument
```

A

There is no problem, and no resolution is necessary.

The message says that by mirroring the system disk, the Power Management function and the suspend-resume function provided by Solaris will become unavailable.

Apart from these functions being unavailable, there is no effect.

For details on the Power Management function and the suspend-resume function, and whether the device you are using supports these functions, see the manual provided with Solaris.



The Power Management system provided by Solaris is different from the automatic power control provided by Enhanced Support Facility (ESF). The above restrictions are unrelated to ESF's automatic power control.

Q: I cannot execute the format(1M) command or prtvtoc(1M) on a physical disk managed by GDS (e.g. c0t0d0s0). How can I solve this problem?

A

Since access to physical disks is restrained by GDS, you cannot perform this operation.

For details, see "[A.2.3 Restraining Access to Physical Special File](#)."

If all you need to do with the commands format(1M) or prtvtoc(1M) is to refer to the disk contents, you may do so by booting from the CD-ROM device or network, and specifying the physical disk. However, you must exercise care since updating the data on the physical disk will damage the mirrored data. For details, see "[A.2.4 Booting from a CD-ROM Device \(or Network\)](#)."

Q: Can I install OS upgrades while system disks are in mirroring configuration through GDS?

A

No. Remove the mirrored status and then upgrade the OS.

G.3 Operation

Questions regarding system operation are explained.

- [Q: Do I need to exit all applications to perform synchronization copying on the volume? If it is not necessary, what effects should I be aware of?](#)
- [Q: How do I change the volume size?](#)
- [Q: I mistakenly updated the /etc/system file, and lost the definition for system disk mirroring. Now, I cannot boot the system. How do I restore the file?](#)
- [Q: I cannot boot the system. How do I restore it?](#)
- [Q: When parting a proxy volume in a shared class, I want to lock the proxy volume not to start up on certain nodes.](#)
- [Q: Is there a possibility that joined proxy volumes are parted automatically?](#)

Q: Do I need to exit all applications to perform synchronization copying on the volume? If it is not necessary, what effects should I be aware of?

A

You do not need to exit the applications, but there could be some affect on the performance. If the application response is seriously affected, you can either suspend the synchronization copying, or adjust the load caused by the copying.

Interrupting the copying:

```
# sdxcopy -I -c class name -v volume name
```

Resuming the copying:

```
# sdxcopy -B -c class name -v volume name
```

Adjusting the copying load:

```
# sdxcopy -P -c class name -v volume name -e delay=delay time
```

Checking the copying process:

```
# sdxinfo -S -c class name -o volume name -e long
```



See

For details, see "D.10 sdxcopy - Synchronization copying operation," and "D.6 sdxinfo - Display object configuration and status information."

Q: How do I change the volume size?

A

Volumes can be expanded with the sdxvolume -S command during the service. For details, see "1.3.7 Online Volume Expansion."

The procedure for changing the volume size while suspending access to the volume is as follows.

1. Exit the application that is accessing the volume
2. Backup the volume data as necessary.
3. Remove the volume.
4. Check if you can create a volume with the necessary size in the group in which the volume existed. You can check this as follows.

```
# sdxinfo -V -c class name -o group name
```

The volume which indicates FREE in the STATUS field will be available. The BLOCKS field indicates the volume size. When there is more than one available volume, the maximum size you can create would be the size of the largest available volume.

5. When you cannot secure sufficient size within the group by performing procedure 4., backup a different volume and remove it, or create another group.
6. Specify a new size and create the volume. When creating the volume, specifying the same name as the volume you removed in procedure 3. will reduce the amount of information you have to input when changing the device definition information in the following procedure.
7. Update the device definition information such as /etc/vfstab, as necessary.

8. Restore the backup data you collected in procedure 2.

Q: I mistakenly updated the /etc/system file, and lost the definition for system disk mirroring. Now, I cannot boot the system. How do I restore the file?

A

The file contents at the point when the system disk mirroring was configured are saved in the following file.

```
/etc/opt/FJSVsdx/bkup/etc/system
```

Try booting the system by following the procedures described below.

1. From the OpenBoot ok prompt, use the boot -a command and boot the system as instructed.
2. For the name of system file, type as below.

```
/etc/opt/FJSVsdx/bkup/etc/system
```

3. For the physical name of root device, type as below.

```
/pseudo/sfdsk@0:2,blk
```

For details on the boot(1M) command, see the Solaris manual.

Q: I cannot boot the system. How do I restore it?

A

See "[F.1.5 System Disk Abnormality](#)."

Q: When parting a proxy volume in a shared class, I want to lock the proxy volume not to start up on certain nodes.

A

Before parting a proxy volume, on nodes where you want to restrain the proxy volume from startup, turn on the proxy volume lock mode with the sdxattr -V command. For details on the lock mode, see "[Lock Volume](#)" and "[D.7 sdxattr - Set objects attributes](#)."

Q: Is there a possibility that joined proxy volumes are parted automatically?

A

In the situation where the master volume has only one ACTIVE slice, if an I/O error occurs in that slice, the proxy volume in STOP or INVALID status that is joined with the master volume through soft copy is parted automatically.

Glossary

active disk

The operating disk between disks that are connected to a switch group.

active slice

A switch slice.

Advanced Copy function

The copy function of Fujitsu Storage Systems ETERNUS. A generic term for EC (Equivalent Copy) and OPC (One Point Copy). For details, see the Fujitsu Storage Systems ETERNUS manual.

alternative root volume

The alternative volume of the root volume.

alternative volume

A volume that is directly or indirectly related to the master or the proxy. The alternative volume of a master volume is the proxy volume. The alternative volume of a proxy volume is the master volume or another proxy volume related to the master volume.

array disk

See "disk array."

attach slice

To add the detached slice to mirror volume making it accessible again as a component of the mirror volume.

automatic resource registration

The function that automatically recognizes the configuration of physical connections of hardware units then registers the hardware with the resource database in the PRIMECLUSTER system. This is performed with the clautoconfig command or the [Tool]:[Automatic configure] menu accessed from the CRM main window in the Cluster Admin window, provided by PRIMECLUSTER CF. For details see " PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide."

available size of a disk

The capacity on an SDX disk that is available for creating volumes. It equals the physical disk size rounded to a cylinder boundary, minus the private slice size.

available size of a group

The capacity that is available for a volume in a group.

background copy

The actual copying that takes place between master and proxy, after virtual completion of creating a snapshot or restoring the master when performing instant snapshot or instant restoring. When performing normal copy, the volume to which the data is being copied cannot be activated. However, when performing background copy, the volume to which the data is being copied can be activated for users to access the valid data.

backup object configuration

To output the configuration of objects within a class to a configuration file in configuration table format.

block

A unit for disk capacity or data transmission. Usually, one block is 512 bytes. Sector is used as a synonym.

break

To cancel the master-proxy relationship, and make independent again.

class

1. A disk class.
2. The generic term for disk classes and shadow classes.

close down class

To make all objects within a class unavailable for reference.

close down volume

To make the volume unavailable (invalid).

cluster application

A resource categorized as a userApplication used to group resources into a logical collection. For details, see the PRIMECLUSTER manual.

cluster control

A component contained in PRIMECLUSTER responsible for detecting various abnormalities and state transition. For details, see the PRIMECLUSTER manual.

cluster system

A highly available and reliable system with multiple nodes connected via a special network. For details, see the PRIMECLUSTER manual.

concatenate

To connect more than one physical disk. It is helpful when using multiple disks as one virtual mass disk.

concatenation

To concatenate.

concatenation group

A group whose type attribute is "concat." It is also the set of disks that will be concatenated.

configuration database

Database storing the configuration data of the objects managed by GDS. The configuration database is stored on a private slice.

configuration file

A text file that contains a configuration table.

configuration table

Text data that tabulates the object configuration of a class.

connect disk

To add an SDX disk to a disk group. User can also connect a keep disk, a single disk, or an undefined disk. Spare disk will be connected automatically by GDS when required.

connect group

To add a group to another group.

The former group is called a "lower level group," and the latter, a "higher level group."

For example, when more than one stripe group is connected to a mirror group that is the highest level group, the mirror volume that is created within the highest level group will be a logical device unit that will be both mirrored and striped.

convert configuration table

To change a configuration table according to a restoration environment in order to restore the object configuration of the class in the environment that is different from the original.

current root volume

The volume currently running as the root file system.

current volume

The master volume or the proxy volume that is described as a file system or a swap area in the /etc/vfstab file.

detach slice

To temporarily detach a mirror slice from a mirror volume, making it inaccessible as a component of the mirror volume.

Other than by user command, GDS will automatically detach slices to maintain synchronization when an I/O error of a mirrored slice is detected.

disconnect disk

To remove a disk from a disk group.

disconnect group

To remove the lower level group from a higher level group.

disk

1. An SDX disk.
2. The generic term for SDX disks and shadow disks.

disk array

A disk unit which combines multiple physical disks to be handled as one virtual disk. Spreading the access across multiple physical disks greatly improves performance and reliability. There are several access methods which are referred to as RAID (Redundant Array of Inexpensive Disks) levels. Main RAID level includes RAID 0 (striping), RAID 1 (mirroring), and RAID 5 (striping with parity).

disk class

A collection of SDX objects. Disk classes may be called "classes." Physical disks can be registered with a class and accordingly are managed by GDS. Classes are classified into root classes, local classes, and shared classes.

disk group

A collection of SDX disks and lower level groups. Disk groups may be called "groups." SDX disks and lower groups that belong to the same group become the unit of mirroring, striping, concatenating, or switching according to the type attribute of the group (mirror, stripe, concatenation, switch).

disk switch

To switch the roles of the active and inactive disks.

domain

One cluster system built by use of PRIMECLUSTER or SynfinityCluster.
A server (computer) that does not belong to a cluster system is also referred to as a domain.

EC

Copying process using the EC (Equivalent Copy) function.

EC function

An abbreviation for Equivalent Copy function. It is one of the Advanced Copy functions provided by Fujitsu Storage Systems ETERNUS. It keeps duplicating disk data within a disk unit. By performing detaching control with software, a replicate of the disk will be created.

For details, see the Fujitsu Storage Systems ETERNUS manual.

EC session

EC request from the host (server) to Fujitsu Storage Systems ETERNUS.

For details, see the Fujitsu Storage Systems ETERNUS manual.

emcpower disk

A virtual device that provides multiple access paths with the emcpower driver. Being able to use multiple paths enables spreading out of the load when accessing the disk, and increases the availability at the time of path failure.

FC-AL internal disk

An internal disk that is connected via an FC-AL (Fibre Channel Arbitrated Loop) interface.

FDU/FU

A disk file unit with multiple physical disks mounted on a single enclosure.

GDS Management View

Software which allows various operations including GDS configuration setting and operation management using the Web screen.

geometry

Physical properties of a disk such as the numbers and sizes of cylinders, heads, and sectors.

GFS

Abbreviation of Global File Services. It is software that manages two file systems, a GFS local file system and a GFS shared file system.
For details see "PRIMECLUSTER Global File Services Configuration and Administration Guide."

GFS Local File System

Local file system that features high-performance file access with contiguous block allocation of the file data, as well as a high-speed consistency recovery function after a system down, which is effective for file system switchover.

For details, see "PRIMECLUSTER Global File Services Configuration and Administration Guide."

GFS Shared File System

Shared file system that features simultaneous accessibility that maintains consistency between multiple nodes to which a shared disk unit is connected and enables the continuation of processing on another node even if a certain node fails.

For details, see "PRIMECLUSTER Global File Services Configuration and Administration Guide."

GR700 series

Disk array unit manufactured by Fujitsu.
For details see GR700 series manuals.

group

1. A disk group.
 2. The generic term for disk groups and shadow groups.
-

hard copy

To copy disk data created within the device using the disk array's hardware function so that no load is imposed on the server or the SAN.

higher level group

A group to which another group belongs.

higher level shadow group

A shadow group to which another shadow group belongs.

hot spare

A feature realizing automatic mirroring recovery while maintaining data accessibility at times of mirror disk failure.

hot swap

A feature allowing users to exchange the faulty physical disk components without stopping the application in use. Also referred to as active exchange.

inactive disk

A standby disk between disks that are connected to a switch group.

inactive slice

A standby slice that will be the operating slice when the active disk of the switch group is changed through a disk switch. The standby slice is a gathering of physically continuous blocks in the standby disk that belongs to the switch group."

instant restore

To instantly restore the master's data, by overwriting the master using the parted proxy data. Instant restore can be created by using the -e instant option with the sdxproxy RejoinRestore command or the -e instant option with the sdxproxy Restore command.

instant snapshot

To create a snapshot of a master instantly by cooperation with the OPC (One Point Copy) function provided by Fujitsu Storage Systems ETERNUS. Also it is a created snapshot. Snapshots can be created by using the -e instant option with the sdxproxy Part command.

instant snapshot by OPC

See "instant snapshot."

JBOD

An abbreviation for "Just a Bunch Of Disks." Multiple disk units are used as a single disk. See also "disk array."

join

To relate a pair of volumes or a pair of groups as a master and a proxy and synchronize the pair.

joined

The state when master and proxy are joined. Proxy volume in joined status cannot be accessed.

JRM

An abbreviation for Just Resynchronization Mechanism.

just resynchronization

A high-speed resynchronization copying process. Available when just resynchronization mode is valid.

Just Resynchronization Mechanism

A mechanism realizing just resynchronization. Abbreviated as JRM.

keep disk

A disk that retains the disk format and data when being registered with a class or being connected to a group.

local class

A "local" type class. Objects that belong to a local class can be used only on the current node. With the local class, disks connected to the current node, excluding system disks, can be registered.

local object

An object that belongs to a "local" class.

logical slice

The component of a logical volume. A logical slice consists of contiguous blocks of one physical disk or multiple blocks of multiple physical disks. A mirror volume consists of one or more logical slices that are mirrored. A single volume, a stripe volume, a volume in a concatenation group, and a switch volume consist of a single logical slice. Logical slices may be called "slices."

logical volume

A logical disk device that provides the user with a block device interface and a character (raw) device interface as well as a physical slice (disk partition). The user can access data on a physical disk through accessing the logical volume. Logical volumes are classified into single volumes, mirror volumes, stripe volumes, the volumes in concatenation groups, and switch volumes. Logical volumes may be called "volumes."

lower level group

A group that belongs to another group. A volume cannot be created within the lower level group.

lower level shadow group

A shadow group that belongs to another shadow group. A shadow volume cannot be created in the lower level shadow group.

master

See "master object."

master group

The original group from which the proxy group is copied.
The term "master" can refer to both master group and master volume.

master object

A generic term for master volume and master group.
Also simply referred to as "master."

master volume

The original volume from which the proxy volume is copied.
Volumes within master group are also called master volumes.
The term "master" can refer to both master volume and master group.

mirror group

A group whose type attribute is "mirror." It is the set of disks and lower level groups, which will be mirrored each other.

mirror slice

A logical slice that composes a mirror volume. A mirror volume is created by mirroring more than one logical slice. For each mirror volume, there will be one logical slice in each disk and each lower level group belonging to the mirror group.

mirror volume

A volume that is created within a mirror group. Its data will be made redundant by mirroring.

mirroring

To maintain redundancy by allocating the same data to multiple slices. A feature which allows you to continue access to a volume if one or more valid slices exist in the volume, even if some failures occur on the other slices.

MPHD disk

A virtual device that provides multiple access paths with the mphd driver. The user will be able to continue operation, even if an access path failure occurs.

MPLB disk

A virtual device that provides multiple access paths with the mplb driver. Being able to use multiple paths enables spreading out of the load when accessing the disk, and increases the availability at the time of path failure.

node

One computer (with one operating system).
For details, see the PRIMECLUSTER manual.

node identifier

A general name for each node within a PRIMECLUSTER system.

node name

The name returned with the `uname -n` command

NVURM

An abbreviation for Non-Volatile Update Region Map. This is information of the updated volume region stored on a private slice for just resynchronization.

object

1. An SDX object. Objects are the generic term for classes, groups, disks, volumes, and slices.
2. The generic term for SDX objects and shadow objects.

offline backup

To create backups from a volume after stopping applications that use the volume.

online backup

To create volume backup while an application is accessing the volume.

online disk migration

To change from one disk to another without stopping applications accessing the disk data.

online snapshot

To create a volume snapshot while an application is accessing the volume.

OPC

Copying process using the OPC (One Point Copy) function.

OPC function

An abbreviation for One Point Copy function. It is one of the Advanced Copy function provided by Fujitsu Storage Systems ETERNUS. Disk data within the disk unit at a certain point will be instantly copied to separate area in the disk unit. The logical copying process is completed immediately after the request to start copying is sent from the server. The actual copying process will be performed later in the background.

For details, see the Fujitsu Storage Systems ETERNUS manual.

OPC session

OPC request from the host (server) to the Fujitsu Storage Systems ETERNUS.

For details, see the Fujitsu Storage Systems ETERNUS manual.

part

To temporarily separate proxy from master so that it may be accessed independently. The parted proxy becomes a snapshot with master data of the moment it was parted. Although parted, the master-proxy relationship is maintained.

parted

The status of a proxy isolated from the master. When a proxy volume is parted, access to the proxy that is independent of access to the master is available.

physical disk

A storage media possible of handling random input and output of data. Mainly used to refer to a set of physical slices.

physical scope

A group of nodes to which physical disks are connected.

physical slice

A part of physical disk and a series of physically continued blocks. The starting sector number and the number of sectors occupied by the physical slice is registered with the VTOC.

physical special file

A device special file of a physical disk, such as `/dev/[r]dsk/c.Xt.Xd.Xs.X` and `/dev/FJSVmplb/[r]dsk/mp1b.Xs.X`.

PRIMECLUSTER

Software to configure the Cluster System. For details, see the PRIMECLUSTER manual.

PRIMECLUSTER system

A cluster system configured using PRIMECLUSTER. For details, see the PRIMECLUSTER manual.

private slice

A special slice used by GDS for control. Part of the SDX disk. Cannot be used as a logical slice.

proxy

See "proxy object."

proxy group

Replication group of master group.

It has the same volume configuration as the master group, and each volume is the proxy volume of the master volume that is within the master group.

The term "proxy" can refer to both proxy group and proxy volume.

proxy object

A generic term for proxy volume and proxy group.

Also simply referred to as "proxy."

proxy volume

A replication volume of a master volume. This may be called an alternative volume.

Proxy volumes and proxy groups may be called "proxies" collectively.

REC

Copying process by REC (Remote Equivalent Copy) function.

REC function

An abbreviation Remote Equivalent Copy function. It is one of the Advanced Copy function provided by Fujitsu Storage Systems ETERNUS. REC is EC function available between Fujitsu Storage Systems ETERNUS boxes. For details, see the Fujitsu Storage Systems ETERNUS manual.

register disk

To add a physical disk to a certain class.

The registered disk is managed by GDS subsequently.

rejoin

To rejoin the parted proxy with master.

relate

To relate a pair of volumes or a pair of groups as a master and a proxy and leave them parted.

remove disk

To remove an SDX disk from disk class.

resource

The concept describing the elements managed by PRIMECLUSTER among hardware and software that compose the PRIMECLUSTER system. In a PRIMECLUSTER system, physical disks and GDS disk classes are managed as resources.

For details, see the PRIMECLUSTER manual.

resource database

The database that manages resource information in a PRIMECLUSTER system. The resource database is managed by the PRIMECLUSTER CF cluster resource management facility. For details, see the PRIMECLUSTER manual.

restore

To restore the master's data, by overwriting the master using the parted proxy data.

restore object configuration

To create object configuration of a class according to the configuration table backed up to a configuration file.

resynchronization copying

To copy data between mirror slices to recover mirroring state when rebooting after a system failure, after a cluster failover, or after reattaching a detached slice to a volume.

Or, to copy data between master and proxy to match their data when rejoining a parted proxy to a master.

root class

A "root" type class. Objects that belong to a root class can be used only on the current node. With the root class, disks connected to the current node, including system disks, can be registered.

root disk

A disk to which the root file system will be mounted. Also referred to as "boot disk."

root object

An object that belongs to a "root" class.

root volume

The volume mounted as the root file system.

ROPC

Copying process by ROPC (Remote One Point Copy) function.

ROPC function

An abbreviation Remote One Point Copy function. It is one of the Advanced Copy functions provided by Fujitsu Storage Systems ETERNUS. ROPC is OPC function available between Fujitsu Storage Systems ETERNUS boxes. For details, see the Fujitsu Storage Systems ETERNUS manual.

SAN

An abbreviation for "Storage Area Network." It is a network connecting multiple servers and storage devices. Usually, Fibre Channel is deployed for connection.

SAN boot

The function that boots the OS from a disk array unit connected to the SAN instead of an internal disk.

scope

One of the class attributes. The scope indicates a group of nodes on which objects of the class are available.

SDX disk

A disk registered with a disk class. SDX disks may be called single disks, keep disks, spare disks, and undefined disks according to their usage. SDX disks may be called "disks."

SDX object

A virtual resource managed by GDS. SDX objects are the generic term for classes, groups, disks, volumes, and slices. SDX objects may be called "objects."

sector

See "block."

shadow class

A collection of shadow objects. Disk classes and shadow classes may be called "classes" collectively. Physical disks managed as SDX disks in another domain and those as the hard copy destinations of SDX disk data copied with disk unit copy functions can be registered with a shadow class and accordingly can be managed and accessed as shadow objects. Objects that belong to a shadow class can be used only on the current node.

shadow disk

A disk registered with a shadow class. SDX disks and shadow disks may be called "disks" collectively.

shadow group

A group created in a shadow class. The shadow group is a collection of shadow disks and lower shadow groups. Disk groups and shadow groups may be called "groups" collectively.

shadow object

A virtual resource managed by GDS. Shadow objects are the generic term for shadow classes, shadow groups, shadow disks, shadow volumes, and shadow slices. These shadow objects are available on a node where GDS Snapshot has been implemented. SDX objects and shadow objects may be called "objects" collectively.

shadow slice

The slice that composes a shadow volume. Logical slices and shadow slices may be called "slices" collectively.

shadow volume

A volume created in a shadow group. Logical volumes and shadow volumes may be called "volumes" collectively.

shared class

A "shared" type class. Objects that belong to a shared class are available on multiple nodes. With the shared class, disks connected to multiple nodes, including the current node, can be registered. Shared classes are also the units of resources used by cluster applications.

shared objects

An object that belongs to a "shared" class.

single disk

An SDX disk which does not belong to any group and is capable of creating single volumes.

single node

A node that is not Cluster System configuration.

single slice

A logical slice that consists a single volume. Single volume consists of one single slice. It is also a series of physically continued block in a single disk.

single volume

A volume that is created within a single disk that does not belong to a group. Its data is not made redundant.

size of a group

The total of available size of a group and the size of one private slice.

slice

1. A logical slice.
 2. The generic term for logical slices and shadow slices.
-

snapshot

To create a replica of a volume at a certain moment in another. It is also a created replica. GDS provides snapshot functions by slice detachment and by proxy volumes. Through use of proxy volumes, snapshot by synchronization and instant snapshot by OPC are available.

snapshot by slice detachment

To create a snapshot of a volume at a moment to a slice by detaching a slice from a mirror volume temporarily. Also it is a created snapshot.

snapshot by synchronization

To create a snapshot of a master at a moment to a proxy by parting the master and proxy which were previously joined and synchronized. Also it is a created snapshot.

soft copy

To copy disk data, using the GDS 's driver.

spare disk

A reserved disk to recover mirroring when a mirrored disk fails.

When an I/O error is detected in a disk belonging to a mirror group or its lower level group, this spare disk will automatically be connected to recover normal mirroring.

SRDF

EMC's software. An abbreviation for Symmetrix Remote Data Facility. Using Symmetrix hardware functions, disk data is copied between Symmetrix cases. For details see the SRDF manual.

start volume

To make volume accessible.

stop volume

To make volume inaccessible.

Storage Area Network

See "SAN."

stripe column

Disks and lower level groups by which striping is conducted. In other words, disks and lower level groups belonging to a stripe group.

stripe group

A disk group whose type attribute is "stripe." This is a group of disks and lower level groups by which striping is conducted.

stripe volume

A volume that is created within a stripe group. Stripe volume can distribute the I/O load among multiple disks with its striping function. Its data will not be redundant.

stripe width

The size of divided data when striping.

striping

To divide data into a certain size and then allocate the divided data alternately to multiple slices. Striping distributes the I/O load across multiple physical disks.

swap

To swap the slices comprising master and proxy that are joined and synchronized.

switch

Disk switch.

switch group

A group of the switch type. This is a gathering of an active disk and an inactive disk that are switched.

switch slice

An operating logical slice comprising a switch volume. A switch volume consists of one switch slice. The switch slice is a gathering of physically continuous blocks in the active disk that belongs to the switch group.

switch volume

A volume created in a switch group. With the disk switch function, the roles of the active and inactive disks can be switched. Data will not be redundant.

Symmetrix

Storage devices from EMC.
For details, refer to a guide to Symmetrix.

synchronization

The state when mirroring is operating normally.
Or, a state when the master and proxy data are the same.

synchronization copying

To copy data from a volume to an attached slice in order to make a mirrored state or to increase the multiplicity of the mirroring configuration, or to copy data between a master and a proxy in order to make their contents equivalent.

synchronized

The state when synchronization is maintained.

system disk

The disk that contains Solaris running on the system. It refers to the entire disk containing slice operating as one of the following file systems (or a swap area).

/, /usr, /var, or a swap area

system volume

The volume currently running as one of the following file systems (or a swap area).

/, /usr, /var, or a swap area

the highest level group

A group that does not belong to any other group. Volumes can be created within the highest level group.

the highest level shadow group

A shadow group that does not belong to another shadow group. A shadow volume can be created in the highest level shadow group.

TimeFinder

EMC's software. Using Symmetrix hardware functions, disk data is copied within a Symmetrix case. For details see the TimeFinder manual.

undefined disk

An SDX disk that is neither a single disk, keep disk or a spare disk. Its usage is not specified yet.

update

To copy (overwrite) data from a master to a parted proxy. The updated proxy is a snapshot that retains an image of the master at a certain moment.

volume

1. A logical volume.
2. The generic term for logical volumes and shadow volumes.

VTOC

Abbreviation for Volume Table Of Contents. The list of starting sector numbers of a physical slice, the numbers of sectors, etc. which exist on a physical disk.

It is stored in the first sector of the physical disk.

Web-Based Admin View

Software for system operation and management using the Web screen.

For details, see "Web-Based Admin View Operation Guide."