

FUJITSU Software PRIMECLUSTER Global Disk Services

Configuration and Administration Guide 4.6

Linux

J2UL-2506-03ENZ0(00)
May 2022

Preface

This manual describes the function, setting up, and managing of the following software:

Software (Feature)	Abbreviation	Description
Global Disk Services	GDS	Volume management software provided by PRIMECLUSTER.
Global Disk Services Snapshot	GDS Snapshot	Optional software of GDS provided by PRIMECLUSTER GD Snapshot.
Global Disk Services I/O Monitor Option	GDS I/O Monitor Option	Optional software of GDS provided by PRIMECLUSTER GD I/O Monitor Option.

Purpose

This manual aims to help users understand GDS, GDS Snapshot, and GDS I/O Monitor Option, including environment configuration, operation, and maintenance.

Target Reader

This manual is devoted to all users who will operate and manage GDS, GDS Snapshot, and GDS I/O Monitor Option, and programmers who will create applications running in GDS environments.

Organization

This manual consists as follows:

Headers	Contents
Chapter 1 Functions	Explains the features and function overview of GDS, GDS Snapshot, and GDS I/O Monitor Option.
Chapter 2 Structure	Explains the concepts and rules of objects used by GDS, GDS Snapshot, and GDS I/O Monitor Option.
Chapter 3 Notes on Design	Explains notes when designing the system where GDS, GDS Snapshot, and GDS I/O Monitor Option are used.
Chapter 4 Installation	Explains how to install GDS, GDS Snapshot, and GDS I/O Monitor Option.
Chapter 5 GDS Management View	Explains how to start and exit GDS Management View and its screen configuration and functions.
Chapter 6 Settings	Explains how to set up GDS, GDS Snapshot, and GDS I/O Monitor Option.
Chapter 7 Operation and Maintenance	Explains how to operate and maintain GDS, GDS Snapshot, and GDS I/O Monitor Option.
Chapter 8 Configuration Change	Explains how to change the configuration of GDS, GDS Snapshot, and GDS I/O Monitor Option.
Chapter 9 Removing Configuration	Explains how to remove the configuration of GDS, GDS Snapshot, and GDS I/O Monitor Option.
Appendix A General Notes	Explains general notes for using GDS, GDS Snapshot, and GDS I/O Monitor Option.

Headers	Contents
Appendix B Command Reference	Explains the commands available in GDS, GDS Snapshot, and GDS I/O Monitor Option, and the PRIMECLUSTER commands used for registering and deleting the shared disk resources.
Appendix C GDS Messages	Explains the contents, possible causes, and resolutions for GDS messages that appear when setting or operating GDS, GDS Snapshot, and GDS I/O Monitor Option.
Appendix D Troubleshooting	Explains resolutions for abnormality of objects and physical disks managed by GDS.
Appendix E Frequently Asked Questions (FAQ)	Lists frequently asked questions regarding GDS, GDS Snapshot, and GDS I/O Monitor Option.
Appendix F Systemd Services and Startup Daemons in GDS	Explains the GDS systemd service and startup daemon.
Appendix G Release Information	Explains new features, manual update details, and changes in specifications depending on the versions of GDS, GDS Snapshot, and GDS I/O Monitor Option.
Glossary	Explains terminologies of GDS, GDS Snapshot, and GDS I/O Monitor Option. Please refer to it as necessary.

GDS: Global Disk Services

GDS Snapshot: Global Disk Services Snapshot

GDS I/O Monitor Option: Global Disk Services I/O Monitor Option

Related documentation

Please refer to the following documents according to need.

- PRIMECLUSTER Concepts Guide
- PRIMECLUSTER Installation and Administration Guide
- PRIMECLUSTER Installation and Administration Guide Cloud Services
- 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
- Software Release Guide PRIMECLUSTER GD
- PRIMECLUSTER GD Installation Guide
- Software Release Guide PRIMECLUSTER GD Snapshot
- PRIMECLUSTER GD Snapshot Installation Guide
- Software Release Guide PRIMECLUSTER GD I/O Monitor Option
- PRIMECLUSTER GD I/O Monitor Option Installation Guide
- FJQSS (Information Collection Tool) User's Guide
- PRIMEQUEST Virtual Machine Function User's Manual

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:

- [EFI] indicates that the description is for UEFI boot environment.
- [IOmonitor] indicates that the description is for GDS I/O Monitor Option.
- [RHEL7] indicates that the description is for RHEL7.
- [RHEL8] indicates that the description is for RHEL8.

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

- Device Mapper Multipath is abbreviated as DM-MP.
- ETERNUS SF AdvancedCopy Manager is abbreviated as ACM.
- FUJITSU Hybrid IT Service FJcloud-O is abbreviated as FJcloud-O.
- FUJITSU Hybrid IT Service FJcloud-Baremetal is abbreviated as FJcloud-Baremetal.
- FUJITSU Hybrid IT Service for Microsoft Azure is abbreviated as "for Azure."
- FUJITSU Hybrid IT Service for AWS is abbreviated as "for AWS."
- FJcloud-V sold by FUJITSU LIMITED and NIFCLOUD sold by FUJITSU CLOUD TECHNOLOGIES LIMITED are abbreviated as "NIFCLOUD" in this manual.
- "for Azure" sold by FUJITSU LIMITED and Microsoft Azure sold by Microsoft Corporation in the United States are abbreviated as "Azure" in this manual.
- "for AWS" sold by FUJITSU LIMITED and AWS (Amazon Web Services) sold by Amazon.com, Inc. are abbreviated as "AWS" in this manual.
- PRIMEQUEST 3000/2000 Series are abbreviated as PRIMEQUEST.
- Red Hat Enterprise Linux is abbreviated as RHEL.
- Red Hat OpenStack Platform is abbreviated as RHOSP.

- The following Symfoware products are abbreviated as Symfoware Server(Native).
 - Symfoware Server V11 or earlier
 - Symfoware Server Standard Edition(Native interface)
 - Symfoware Server Enterprise Edition(Native interface)

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

Amazon Web Services is a registered trademark of Amazon.com, Inc. or its affiliates in the United States and/or other countries.

Microsoft, Microsoft Edge, Azure, Internet Explorer, and Windows are trademarks of the Microsoft group of companies.

Linux(R) is the registered trademark of Linus Torvalds in the U.S. and other countries.

Red Hat is a registered trademark of Red Hat, Inc. in the U.S. and other countries.

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

Dell EMC, EMC, PowerPath, SAN Manager, SRDF, Symmetrix, and TimeFinder are trademarks of Dell Inc. or its subsidiaries.

VMware and other VMware product names are trademarks or registered trademarks of VMware in the United States and/or other jurisdictions.

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

December 2019, First edition May 2022, Third edition

All Rights Reserved, Copyright (C) FUJITSU LIMITED 2019-2022.

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 [EFI].....	3
1.2.1.2 Mirroring between Disk Array Unit.....	4
1.2.1.3 Shared Disk Mirroring.....	5
1.2.1.4 Mirroring among Servers.....	6
1.2.2 Hot Spare.....	7
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.....	11
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
1.3.5.1 Class Scope.....	14
1.3.5.2 Starting and Stopping Volume.....	14
1.3.5.3 Access Mode.....	15
1.3.5.4 Lock Volume.....	15
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.....	17
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 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 Alternative Boot Environment [EFI].....	33
1.6 Shadow Volume.....	34
1.6.1 Accessing Volumes from External Server.....	34
1.6.2 Using Copy Function of Disk Unit.....	35
1.7 GDS I/O Monitor Option Features.....	36
1.8 I/O Response Time Assurance [IOmonitor].....	38
1.8.1 I/O response delay for one slice on the mirror volume.....	38
1.8.2 I/O response delay on all slices of the mirror volume.....	39
1.8.3 I/O response delay for the slice of a single volume.....	40
Chapter 2 Structure.....	42
2.1 SDX Object.....	42
2.1.1 Disk Class.....	42
2.1.2 SDX Disk.....	46
2.1.3 Disk Group.....	49
2.1.4 Logical Volume.....	52
2.1.5 Logical Slice.....	57
2.2 GDS Snapshot Objects.....	61
2.2.1 Proxy Object.....	61
2.2.2 Shadow Object.....	65

2.2.2.1 Shadow Class.....	66
2.2.2.2 Shadow Disk.....	68
2.2.2.3 Shadow Group.....	71
2.2.2.4 Shadow Volume.....	72
2.2.2.5 Shadow Slice.....	77
2.3 Rules.....	78
2.3.1 Object Name.....	78
2.3.2 Number of Classes.....	79
2.3.3 Number of Disks.....	79
2.3.4 Number of Groups.....	80
2.3.5 Number of Volumes.....	80
2.3.6 Number of Keep Disks [EFI].....	80
2.3.7 Creating Group Hierarchy.....	80
2.3.8 Disk Label.....	82
2.3.9 Disk Size.....	83
2.3.10 Group Size.....	84
2.3.11 Volume Size.....	84
2.3.12 Proxy Configuration Preconditions.....	85
2.3.13 Number of Proxy Volumes.....	85
2.3.14 Proxy Volume Size.....	85
2.3.15 Proxy Group Size.....	86
Chapter 3 Notes on Design.....	87
3.1 Guidelines.....	87
3.1.1 Guidelines for Mirroring.....	87
3.1.2 Guidelines for Striping.....	87
3.1.3 Guidelines for Concatenation.....	88
3.1.4 Guidelines for Combining Striping with Mirroring.....	88
3.2 DM-MP (Device Mapper Multipath).....	88
3.3 Dell EMC Storage Units.....	88
3.4 Managing System Disks.....	89
3.5 System Disk Mirroring Preconditions.....	89
3.6 Initializing Disk.....	90
3.7 Keep Disk [EFI].....	90
3.8 Hot Spare.....	91
3.8.1 Number of Disks within Class.....	91
3.8.2 Hot Spare for Hierarchized Mirror Groups.....	93
3.8.3 Number of Spare Disks.....	94
3.8.4 Spare Disk Size.....	94
3.8.5 Hot Spare for the Root Class.....	94
3.8.6 Hot Spare for Proxy Volumes.....	94
3.8.7 Shadow Classes.....	94
3.8.8 Disk Array Unit's Hot Spare Function.....	94
3.8.9 Spare Disk Failure.....	94
3.8.10 Synchronization Copying Invoked by Hot Spare.....	94
3.8.11 Required Time for Synchronization Copying Invoked by Hot Spare.....	94
3.8.12 Hot Spare Mode (Spare Disk Selection Mode).....	95
3.8.13 Spare Disk Manual Connection.....	96
3.9 Mirroring among Servers.....	97
3.9.1 Network Configuration.....	97
3.9.2 Disk.....	98
3.9.3 Object Configuration.....	99
3.9.4 Configuration and Settings of the Cluster System.....	99
3.9.5 Data Amount and Data Update Amount.....	100
3.9.6 Required Time for Copying Processing.....	101
3.9.7 Note.....	101
3.10 Volume Access Mode.....	101

3.11 Data Consistency at the time of Simultaneous Access.....	102
3.12 Online Volume Expansion.....	102
3.13 Disk Switch.....	103
3.14 Creating Snapshot by Slice Detachment.....	103
3.15 Online Backup and Restoration Using Proxy Volume.....	103
3.16 Difference between Mirror Slice and Proxy Volume.....	105
3.17 Ensuring Consistency of Snapshot Data.....	105
3.18 Instant Snapshot by OPC.....	105
3.19 Using Dell EMC TimeFinder or Dell EMC SRDF in Proxy Configuration.....	106
3.20 Backing Up and Restoring through External Server.....	108
3.21 Virtual Environment.....	110
3.21.1 KVM Environment and VMware Environment.....	110
3.21.2 Cloud Environment.....	113
3.21.3 RHOSP Environment.....	113
3.22 Disaster Recovery.....	114
3.22.1 Building Disaster Recovery System.....	115
3.22.2 Data Synchronization of Operation System and Disaster Recovery System.....	115
3.22.3 Switching to Disaster Recovery System.....	115
3.22.4 Restoration to Operation System from Disaster Recovery System.....	116
3.23 I/O Response Time Assurance [IOmonitor].....	116
Chapter 4 Installation.....	118
4.1 Notes on Installing OS.....	118
4.2 Installing GDS.....	118
4.3 Installing GDS Snapshot.....	118
4.4 Installing GDS I/O Monitor Option.....	118
4.5 Setting Up DM-MP.....	118
4.6 Configuring KVM Environment.....	119
4.6.1 Adding Virtual Disk.....	119
4.6.2 Major Number of the virtio Block Device.....	120
4.6.3 Setting up Shared Disk.....	120
4.6.4 Setting up virtio Block Device Used for Mirroring among Servers.....	121
4.7 Configuring VMware Environment.....	121
4.8 Disk Setting for Performing Mirroring among Servers.....	122
4.8.1 Setting Tuning Parameters.....	122
4.8.2 Checking and Setting Required Packages.....	123
4.8.3 Creating iSCSI Target.....	124
4.8.4 Establishing iSCSI Session.....	127
4.8.5 Confirming the Correspondence of iSCSI Device Name of Each Node.....	131
4.9 Preparation for Starting GDS Management View.....	132
4.9.1 Deciding the User Group.....	132
4.9.1.1 User Group Types.....	132
4.9.1.2 Creating User Groups.....	132
4.9.1.3 Registering to a User Group.....	132
4.9.2 Setting up the Client Environment.....	133
4.10 Settings to Use Dell EMC PowerPath.....	133
Chapter 5 GDS Management View.....	134
5.1 Starting GDS Management View.....	134
5.1.1 Starting Web-Based Admin View Operation Menu.....	134
5.1.2 Web-Based Admin View Operation Menu Functions.....	134
5.1.3 Starting GDS Management View.....	135
5.2 Screen Configuration.....	136
5.3 Menu Configuration and Functions.....	139
5.3.1 [General] Menu.....	139
5.3.2 [Settings] Menu.....	140
5.3.3 [Operation] Menu.....	145
5.3.4 [View] Menu.....	148

5.3.5 [Help] Menu.....	148
5.4 Icon Types and Object Status.....	149
5.5 Object Information.....	152
5.6 Operation.....	153
5.7 Exiting GDS Management View.....	156
5.8 Uninstalling Java Application.....	157
Chapter 6 Settings.....	158
6.1 Creating Excluded Device List.....	158
6.1.1 Excluded Device List.....	158
6.1.2 To Use DM-MP (Device Mapper Multipath).....	158
6.1.3 To Use Dell EMC Storage Units.....	158
6.1.4 To Use Mirroring among Servers.....	160
6.2 Using iSCSI Disk.....	160
6.3 Shared Disk Resource Registration.....	160
6.3.1 What Is Shared Disk Resource Registration?.....	160
6.3.2 Flow of Shared Disk Resource Registration.....	161
6.3.3 Preconditions and Important Points.....	161
6.3.4 Shared Disk Unit Resource Registration Procedure.....	162
6.4 Updating Physical Disk Information.....	165
6.5 System Disk Mirroring Settings [EFI].....	165
6.5.1 Checking Partition Configuration of System Disk.....	166
6.5.2 Setting System Disk Mirroring.....	166
6.5.2.1 Setting with GDS Management View.....	166
6.5.2.1.1 Operation Flow.....	166
6.5.2.1.2 Operation Procedure.....	167
6.5.2.2 Setting with Commands.....	173
6.5.3 Checking System Disk Settings.....	173
6.5.4 Checking Physical Disk Information and Slice Numbers.....	175
6.5.5 Collecting Backups of System Disk.....	178
6.5.6 Backing up EFI Configuration Information.....	178
6.6 Class, Group, Volume Settings.....	179
6.6.1 Setting with GDS Management View.....	179
6.6.1.1 Single Volume Configuration Settings.....	179
6.6.1.2 Other Volume Configuration Settings.....	179
6.6.1.3 Class Configuration.....	180
6.6.1.4 Cluster System Class Configuration.....	184
6.6.1.5 Group Configuration.....	186
6.6.1.6 Volume Configuration.....	188
6.6.2 Setting with Commands.....	192
6.6.2.1 Creating Single Volume.....	192
6.6.2.2 Creating Mirror Volume.....	192
6.6.2.3 Creating Netmirror Volume.....	193
6.6.2.4 Creating Large-Capacity Volume Using Concatenation.....	194
6.6.2.5 Creating Stripe Volume.....	194
6.6.2.6 Creating Mirror Volume by Combining Striping and Mirroring.....	195
6.7 File System Configuration.....	196
6.7.1 Creating the File System.....	196
6.7.2 File System Auto Mount.....	197
6.7.3 In Cluster Systems.....	198
6.8 raw Device Binding.....	198
6.9 Volume's Block Special File Access Permission.....	198
6.10 Proxy Configuration.....	200
6.10.1 Setting with GDS Management View.....	201
6.10.1.1 Join.....	201
6.10.1.2 Relate.....	205
6.10.2 Setting with Commands.....	207

6.10.2.1 Join.....	207
6.10.2.2 Relate.....	207
6.11 Cloning.....	208
6.11.1 Summary.....	208
6.11.2 Preconditions.....	209
6.11.3 Preparation.....	209
6.11.4 Copying Disk Data.....	209
6.11.5 Setting up Copy Destination System.....	209
6.12 I/O Response Time Assurance [IOmonitor].....	211
6.12.1 Setting I/O Response Time Assurance.....	211
6.12.2 Hang-up Monitoring Function.....	212
Chapter 7 Operation and Maintenance.....	214
7.1 Viewing Configurations/Statuses and Monitoring Statuses.....	214
7.1.1 Confirming SDX Object Configuration.....	214
7.1.2 Viewing Proxy Object Configurations.....	219
7.1.3 Monitoring Object Status.....	223
7.1.4 Viewing Object Statuses.....	225
7.2 Synchronization Copy Operation of Mirror Volume.....	225
7.2.1 Using GDS Management View.....	225
7.2.2 Using Commands.....	229
7.3 Disk Swap.....	229
7.3.1 Setting with GDS Management View.....	229
7.3.1.1 Operation Flow.....	229
7.3.1.2 Operation Procedure.....	230
7.3.2 Using Commands.....	234
7.3.3 Swapping Disks of Netmirror Group.....	234
7.3.3.1 Hot Swap.....	234
7.3.3.2 Cold Swap.....	240
7.3.4 Swapping Physical Disks Registered with Shadow Classes.....	244
7.3.5 Notes.....	244
7.3.5.1 Physical Disk Size.....	244
7.3.5.2 Physical Disks That Cannot Be Swapped.....	244
7.3.5.3 LinuxSwapping Physical Disks when the Proxy Volume Is in a Mirroring Configuration.....	245
7.3.5.4 Swapping Internal Disks Registered with Root Classes or Local.....	245
7.4 Backing Up and Restoring System Disk.....	246
7.4.1 Settings Before Backing Up.....	247
7.4.2 Backing Up [EFI].....	247
7.4.3 Settings After Backing Up.....	250
7.4.4 Settings Before Restoring.....	251
7.4.5 Restoring (When System Can Be Booted) [EFI].....	251
7.4.6 Restoring (When System Cannot Be Booted) [EFI].....	254
7.4.7 Settings After Restoring.....	256
7.5 Backing Up and Restoring System Disk through Alternative Boot Environment [EFI].....	257
7.5.1 System Configuration.....	257
7.5.2 Summary of Backup.....	259
7.5.3 Summary of Restore.....	260
7.5.4 Summary of Procedure.....	261
7.5.5 Configuring Environment.....	261
7.5.6 Backing Up.....	264
7.5.7 Restoring.....	267
7.6 Backup (Offline).....	272
7.7 Backup (by Slice Detachment).....	273
7.7.1 Using GDS Management View.....	273
7.7.1.1 Operation Flow.....	273
7.7.1.2 Operation Procedure.....	274
7.7.2 Using Commands.....	279

7.8 Backup (by Synchronization)	280
7.8.1 Using GDS Management View	282
7.8.1.1 Operation Flow	282
7.8.1.2 Operation Procedure	283
7.8.2 Using Commands	288
7.9 Backup (by OPC)	291
7.9.1 Using GDS Management View	292
7.9.1.1 Operation Flow	292
7.9.1.2 Operation Procedure	292
7.9.2 Using Commands	295
7.10 Restoring	298
7.10.1 Using GDS Management View	298
7.10.1.1 Operation Flow	298
7.10.1.2 Operation Procedure	299
7.10.2 Using Commands	303
7.10.2.1 Without Using GDS Snapshot	303
7.10.2.2 Using GDS Snapshot	304
7.11 Backing Up and Restoring through Disk Unit's Copy Functions	307
7.11.1 Configuring Environment	307
7.11.2 Backing Up	308
7.11.3 Restoring from Backup Disks	312
7.12 Backing Up and Restoring through External Server	313
7.12.1 Backing Up and Restoring Logical Volume with No Replication	313
7.12.1.1 System Configuration	314
7.12.1.2 Summary of Backup	314
7.12.1.3 Summary of Restore	316
7.12.1.4 Summary of Procedure	317
7.12.1.5 Configuring Environment	318
7.12.1.6 Backing Up	319
7.12.1.7 Restoring	323
7.12.2 Backing Up and Restoring through Snapshot by Slice Detachment	327
7.12.2.1 System Configuration	328
7.12.2.2 Summary of Backup	328
7.12.2.3 Summary of Restore	330
7.12.2.4 Summary of Procedure	331
7.12.2.5 Configuring Environment	332
7.12.2.6 Backing Up	333
7.12.2.7 Restoring	337
7.12.3 Backing Up and Restoring Using Snapshots from Proxy Volume	343
7.12.3.1 System Configuration	343
7.12.3.2 Summary of Backup	344
7.12.3.3 Summary of Restore from Proxy Volume	345
7.12.3.4 Summary of Restore from Tape	346
7.12.3.5 Summary of Procedure	348
7.12.3.6 Configuring Environment	349
7.12.3.7 Backing Up	350
7.12.3.8 Restoring from Proxy Volume	353
7.12.3.9 Restoring from Tape	354
7.12.4 Backing Up and Restoring by Disk Unit's Copy Function	358
7.12.4.1 System Configuration	359
7.12.4.2 Summary of Backup	360
7.12.4.3 Summary of Restore from BCV	362
7.12.4.4 Summary of Restore from Tape	363
7.12.4.5 Summary of Procedure	365
7.12.4.6 Configuring Environment	366
7.12.4.7 Backing Up	367
7.12.4.8 Restoring form BCV	372

7.12.4.9 Restoring from Tape.....	374
7.13 Backing Up and Restoring Object Configurations.....	379
7.13.1 Backing Up.....	379
7.13.2 Restoring.....	380
7.14 Volume Expansion.....	381
7.14.1 Volume Expansion of Local or Shared Class.....	381
7.14.2 Volume Expansion Using Commands [EFI].....	381
7.15 I/O Response Time Assurance [IOmonitor].....	388
7.16 Mirroring Among Servers.....	388
7.16.1 I/O Error Due to Node Stop or Network Error.....	388
7.16.2 Stopping or Restarting the Node.....	389
7.16.3 Operation when the LEFTCLUSTER Node Exists.....	389
7.16.4 Restoration from the LEFTCLUSTER State.....	390
7.16.5 Lock Volume of the Netmirror Volume.....	390
7.16.6 Operation on Only One Node.....	391
7.16.7 Restoration to 2-node Operation after Operation on Only One Node.....	391
7.16.8 Restoration when the Latest Disk Cannot Be Selected Automatically.....	391
7.16.9 Stopping both Nodes.....	394
7.16.10 Restarting the Node during Synchronization Copying.....	395
7.16.11 Switchover of the Cluster Application when the Network is Disconnected.....	396
7.16.12 Configuration Change when the Node is Stopped.....	396
7.17 Maintenance in Cloud Environment.....	396
7.18 Removing Errata Kernel.....	396
Chapter 8 Configuration Change.....	398
8.1 Changing Web-Based Admin View Settings.....	398
8.2 Changing Object Configuration.....	398
8.2.1 Using GDS Management View.....	398
8.2.1.1 Operation Flow.....	398
8.2.1.2 Class Configuration.....	399
8.2.1.3 Group Configuration.....	404
8.2.1.4 Volume Configuration.....	407
8.2.2 Using Commands.....	408
8.3 Changing the Settings of I/O Response Time Assurance [IOmonitor].....	409
8.3.1 Changing the I/O response time.....	409
8.3.2 Changing the System Panic Mode.....	410
8.4 Deleting the Settings of the I/O Response Time Assurance [IOmonitor].....	410
8.5 Migration from Single Nodes to Cluster System.....	411
8.6 Server Migration.....	411
8.7 Storage Migration.....	413
8.7.1 Using GDS Snapshot Online Disk Migration.....	413
8.7.1.1 Using GDS Management View.....	414
8.7.1.1.1 Operation Flow.....	414
8.7.1.1.2 Operation Procedure.....	415
8.7.1.2 Using Commands.....	417
8.7.2 Using Copy Function of Storage Units and sdxconfig Command.....	418
8.8 Changing Driver Major Number.....	420
8.8.1 Changing sfdsk Driver Major Number.....	420
8.8.2 Changing sfdsklog Driver Major Number.....	422
8.9 Upgrading Procedure for Dell EMC PowerPath.....	422
8.10 Changing IP Addresses Used for Mirroring among Servers.....	425
8.11 Adding Disks Used for Mirroring among Servers.....	428
Chapter 9 Removing Configuration.....	431
9.1 Breaking Proxy.....	431
9.1.1 Using GDS Management View.....	431
9.1.2 Using Commands.....	432
9.2 Removing File System.....	432

9.3 Removing Volumes, Groups, and Classes.....	433
9.3.1 Using GDS Management View.....	433
9.3.1.1 Operation Flow.....	433
9.3.1.2 Removing Volume.....	433
9.3.1.3 Removing Group.....	435
9.3.1.4 Removing Class.....	436
9.3.2 Using Commands.....	437
9.4 System Disk Mirroring Cancellation [EFI].....	438
9.4.1 Unmirroring System Disk with GDS Management View.....	438
9.4.1.1 Unmirroring the System Disk [EFI].....	438
9.4.1.2 Operation Flow.....	438
9.4.2 Cancellation with Commands.....	441
9.5 OS Re-installation.....	443
9.6 Deleting iSCSI Settings Used for Mirroring among Servers.....	444
9.6.1 Delete iSCSI devices.....	444
9.6.2 Deleting iSCSI Initiator Settings.....	447
9.6.3 Deleting iSCSI Target Settings.....	448
Appendix A General Notes.....	450
A.1 Restraining Access to Physical Special File.....	450
A.2 Booting from a CD-ROM Device.....	450
A.3 Operation in Cluster System.....	451
A.4 Just Resynchronization Mechanism (JRM).....	451
A.5 Backing Up and Restoring Object Configuration (sdxcfg).....	453
A.6 Physical Disk Recognition.....	453
A.7 Object Operation When Using Proxy.....	454
A.8 Using Advanced Copy Function in Proxy Configuration.....	454
A.9 Shadow Volume.....	455
A.10 Changing system run level.....	458
A.11 Root Class Operation [EFI].....	459
A.12 Operating When There Is Disk in DISABLE Status or There Is Class not Displayed with sdxcfg Command.....	459
A.13 Notes on Mirroring among Servers.....	459
Appendix B Command Reference.....	461
B.1 GDS Commands.....	462
B.1.1 sdxcfg - Class Operations.....	462
B.1.2 sdxdisk - Disk Operations.....	462
B.1.3 sdxcfg - Group Operations.....	472
B.1.4 sdxcfg - Volume Operations.....	476
B.1.5 sdxcfg - Slice Operations.....	481
B.1.6 sdxcfg - Display Object Configuration and Status Information.....	484
B.1.7 sdxcfg - Change Objects Attributes.....	501
B.1.8 sdxcfg - Swap Disk.....	509
B.1.9 sdxcfg - Restore Failed Object.....	511
B.1.10 sdxcfg - Synchronization Copying Operation.....	514
B.1.11 sdxcfg - Root File System Mirroring Definition and Cancellation [EFI].....	515
B.1.12 sdxcfg - Configuration Parameter Operations.....	519
B.1.13 sdxcfg - Object Configuration Operations.....	520
B.1.14 sdxcfg - Display Device Information of Internal Disk.....	525
B.2 GDS Snapshot Commands.....	526
B.2.1 sdxcfg - Proxy Object Operations.....	526
B.2.2 sdxcfg - Shadow Disk Operations.....	537
B.2.3 sdxcfg - Shadow Group Operations.....	542
B.2.4 sdxcfg - Shadow Volume Operations.....	545
B.3 Command Reference.....	548
B.3.1 clautoconfig(8)-Execute resource registration.....	548
B.3.2 cldelrsc(8)-Delete resource.....	549

Appendix C GDS Messages.....	552
C.1 Web-Based Admin View Messages (0001-0099).....	552
C.2 Driver Messages.....	552
C.2.1 Panic Messages (20000-20099).....	553
C.2.2 Error Messages (21000-21099).....	554
C.2.3 Warning Messages (22000-22099).....	555
C.2.4 Information Messages (24000-24099).....	562
C.2.5 Internal Error Messages (26000-26001).....	566
C.3 Daemon Messages.....	567
C.3.1 Halt Messages (40000-40099).....	569
C.3.2 Error Messages (42000-42099).....	570
C.3.3 Warning Messages (44000-44099).....	593
C.3.4 Information Messages (46000-46199).....	604
C.3.5 Internal Error Messages (48000).....	625
C.4 Command Messages.....	626
C.4.1 Error Messages (60000-60499).....	628
C.4.2 Warning Messages (62000-62099).....	724
C.4.3 Information Messages (64000-64099).....	730
C.4.4 Fix Messages (66000).....	744
C.4.5 Internal Error Messages (68000).....	744
C.5 Operation Management View Messages.....	745
C.5.1 Error Messages (5000-5099).....	745
C.5.2 Warning Messages (5000,5100-5199).....	749
C.5.3 Information Messages (5200-5299).....	759
Appendix D Troubleshooting.....	769
D.1 Resolving Problems.....	769
D.1.1 Slice Status Abnormality.....	769
D.1.2 Disk Status Abnormality.....	776
D.1.3 Volume Status Abnormality.....	778
D.1.4 Class Status Abnormality.....	794
D.1.5 System Disk Abnormality [EFI].....	798
D.1.6 Errors in the Mirroring Among Servers.....	811
D.1.7 GDS Management View Abnormality.....	822
D.1.8 Proxy Object Abnormality.....	825
D.1.9 Errors in Dell EMC Storage Units.....	828
D.1.10 Cluster System Related Error.....	832
D.1.11 File System Error.....	837
D.1.12 Shadow Object Errors.....	841
D.1.13 Disk Unit Error.....	842
D.1.14 OS Messages.....	842
D.1.15 Other Errors.....	848
D.2 Collecting Investigation Material.....	849
D.2.1 Collecting Investigation Material with pclsnap Command.....	849
D.2.2 Collecting Investigation Material with sdxsnap.sh Script.....	849
D.2.3 Collecting Investigation Material with FJQSS.....	850
Appendix E Frequently Asked Questions (FAQ).....	851
E.1 Operation Design.....	851
E.2 Environment Configuration.....	854
E.3 Operation.....	855
Appendix F Systemd Services and Startup Daemons in GDS.....	858
F.1 Explanation Formats.....	858
F.2 systemd Service Lists.....	858
Appendix G Release Information.....	861
G.1 New Features.....	861

G.1.1 New Features in 4.3A00.....	861
G.1.2 New Features in 4.3A10.....	861
G.1.3 New Features in 4.3A20.....	861
G.1.4 New Features in 4.3A30.....	862
G.1.5 New Features in 4.3A40.....	863
G.1.6 New Features in 4.4A00.....	864
G.1.7 New Features in 4.5A00.....	864
G.1.8 New Features in 4.5A10.....	865
G.1.9 New Features in 4.6A00.....	865
G.1.10 New Features in 4.6A10.....	865
G.1.11 New Features in 4.6A20.....	865
G.2 Manual Changes.....	865
G.2.1 4.3A00 Manual Changes.....	865
G.2.2 4.3A10 Manual Changes.....	866
G.2.3 4.3A20 Manual Changes.....	879
G.2.4 4.3A30 Manual Changes.....	882
G.2.5 4.3A40 Manual Changes.....	886
G.2.6 4.4A00 Manual Changes.....	895
G.2.7 4.5A00 Manual Changes.....	903
G.2.8 4.5A10 Manual Changes.....	910
G.2.9 4.6A00 Manual Changes.....	911
G.2.10 4.6A10 Manual Changes.....	912
G.2.11 4.6A20 Manual Changes.....	914
G.3 Specification Changes Other Than New Features.....	915
G.3.1 Specification Changes in 4.3A00.....	915
G.3.1.1 Support for Export to NFS Clients.....	915
G.3.1.2 Change of I/O Amount Supported for Simultaneous Processing.....	915
G.3.2 Specification Changes in 4.3A10.....	915
G.3.2.1 Disk Label Type of a Disk with the Size smaller than 1 TB.....	915
G.3.3 Specification Changes in 4.3A20.....	916
G.3.3.1 Disk Label Type of a Disk with the Size smaller than 1 TB.....	916
G.3.4 Specification Changes in 4.3A30.....	916
G.3.4.1 Restoring Object Configuration (sdxconfig Restore Command).....	916
G.3.4.2 File System Configuration screen of GDS Management View.....	917
G.3.4.3 Changes of the warning message "5119" in the Volume Configuration screen.....	917
G.3.5 Specification Changes in 4.3A40.....	917
G.3.6 Specification Changes in 4.4A00.....	917
G.3.7 Specification Changes in 4.5A00.....	917
G.3.8 Specification Changes in 4.5A10.....	917
G.3.9 Specification Changes in 4.6A00.....	917
G.3.9.1 Operation after Restoration from LEFTCLUSTER in Mirroring among Servers Configuration.....	917
G.3.9.2 Unsupported Functions in RHEL8.....	917
G.3.10 Specification Changes in 4.6A10.....	917
G.3.11 Specification Changes in 4.6A20.....	918
Glossary.....	919

Chapter 1 Functions

This chapter describes the features and functions of Global Disk Services (hereinafter GDS), Global Disk Services Snapshot (hereinafter GDS Snapshot), and Global Disk Services I/O Monitor Option (hereinafter GDS I/O Monitor Option).

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 closely related functions below:

- 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 Linux server. In a PRIMECLUSTER system, GDS also allows users to manage all disk units connected to all servers in an integrated manner for shared in a SAN (Storage Area Network) storage environment as well as for local disk units connected to specific servers.

Figure 1.2 SAN (Storage Area Network)




In general, 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. The problem is that 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.

This section explains functions which protect data from unexpected hardware problems and improve system availability.


1.2.1 Disk Mirroring

Disk mirroring utilizes multiple disks to maintain data synchronization.

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

With GDS disk mirroring, applications can continue disk accesses even if one mirror disk (a disk that is mirrored) fails. Therefore, all the data in the disk is not damaged and applications can continue normal operations.

Figure 1.4 Disk Mirroring



GDS supports mirroring of single disk units as well as mirroring between 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.

Examples of special mirroring implementation are explained below.

1.2.1.1 System Disk Mirroring [EFI]

System disk mirroring can mirror system disks with Linux operating system 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 reboot in a situation where disk failure has occurred.



See

.....
For the system conditions where the system disk can be mirrored, see "[3.5 System Disk Mirroring Preconditions](#)."
.....

Figure 1.5 System Disk Mirroring




1.2.1.2 Mirroring between Disk Array Unit

GDS can provide mirroring between high-performance and high-reliability disk arrays.

Mirroring two disk arrays connected with Fibre Channel can protect data from unexpected accidents and power shortage.

Moreover, disk units with redundant access paths can also be mirrored with specific software.

Figure 1.6 Mirroring between Disk Array Unit



1.2.1.3 Shared Disk Mirroring


GDS can mirror shared disk units connected with a cluster system composed of multiple servers (also referred to as nodes).

Such mirroring 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 Global File Services (hereinafter GFS) Shared File System, that provide simultaneous access from multiple servers to shared disks, not to mention switch-over or standby type applications.

In the virtual environment (such as KVM, VMware, and RHOSP), shared disks can be mirrored using GDS which is operated on the guest OS.

Figure 1.7 Shared Disk Mirroring



1.2.1.4 Mirroring among Servers

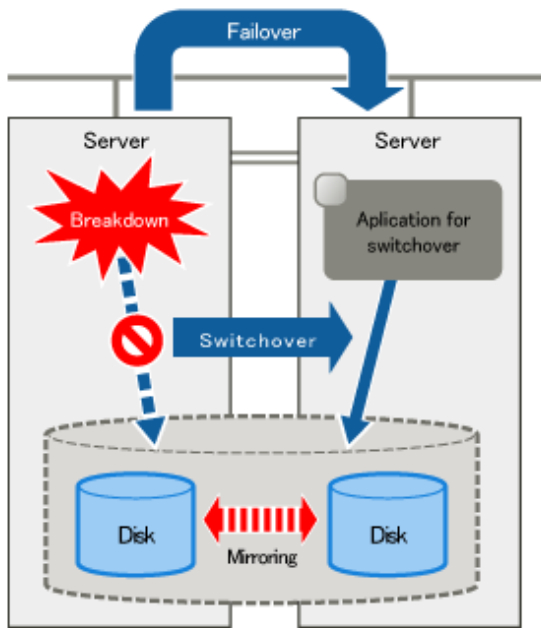
Mirroring among servers is a function that mirrors the local disks (such as an internal disk) of each server via the network in a cluster system configured with two servers.

Mirroring among servers can be used for taking over the data between nodes in a switch-over type cluster application that performs the standby operation. Compared to using the shared disk device, its features are as follows:

- The cluster system can be configured at lower cost.
- It does not require the setup and operation skills for the disk storage array.
- It is suitable for work with a small amount of data and a small amount of data update.

In a public cloud where shared disk drives are not available, Mirroring among Servers is used to pass data between nodes in a cluster system. Since data can be transferred between different zones, when a network failure occurs in the entire Zone or when Zones become abnormal due to a large-scale disaster, business can be quickly restored by failing over to the standby system with fewer operations.

Figure 1.8 Mirroring among Servers



Note


- If the mirroring among servers is used in a VMware environment, VMware vCenter Server is required. For details, refer to "Forced stop method in a VMware environment" in "3.9.4 Configuration and Settings of the Cluster System."
- The mirroring among servers is suitable for the operation if the amount of data and the amount of data update is small. If the amount of data and the amount of data update of the operation is large, and data accessibility performance or data availability is especially important, use the shared disk.
For details, see "3.9.5 Data Amount and Data Update Amount."
- In the configuration of mirroring among servers, much time to switch the cluster application is required than in the configuration of shared disk. An additional few seconds are required for one netmirror volume to switch the cluster application. To switch the cluster application promptly, use the shared disk.
- When the disk on another node cannot be accessed due to a network error, and when the resynchronization copying is performed after the network is restored, I/O to the volume is suspended for a few seconds for one volume. Lower number of volumes is recommended.
The time to suspend I/O is about 4 to 5 seconds though it changes depending on the system.
- In this version, the configuration of the mirroring among servers using Symfoware Server(Native) is supported only in a cloud environment (except an FJcloud-Baremetal environment). When using Symfoware Server(Native) in a physical environment or in a virtual environment, use the shared disk.
- When using the mirroring among servers, see also "A.13 Notes on Mirroring among Servers."

1.2.2 Hot Spare

Overview

"hot spare" can automate mirroring recovery using spare disks in the event of mirrored disk failure.

Figure 1.9 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 to recover normal mirroring.

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, the spare disk will automatically be disconnected after the failed disk is swapped with another disk and synchronization copying to the new disk is complete.



Note

The hot spare function cannot be used for system disks.



Note

Points of Concern for Hot Spare


See "[3.8 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 1.10 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 can complete copying, which usually takes a few minutes for each 1 GB, by quickly copying only portions where data mirrored by synchronization is not retained.

Figure 1.11 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.4 Just Resynchronization Mechanism \(JRM\)](#)."

Note

Root File System Volume Resynchronization [EFI]

Even when the OS is shut down normally, access from the OS to the root file system (/) volume is not canceled. For this reason, at server startup, resynchronization of the root (/) volume is always performed.

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.

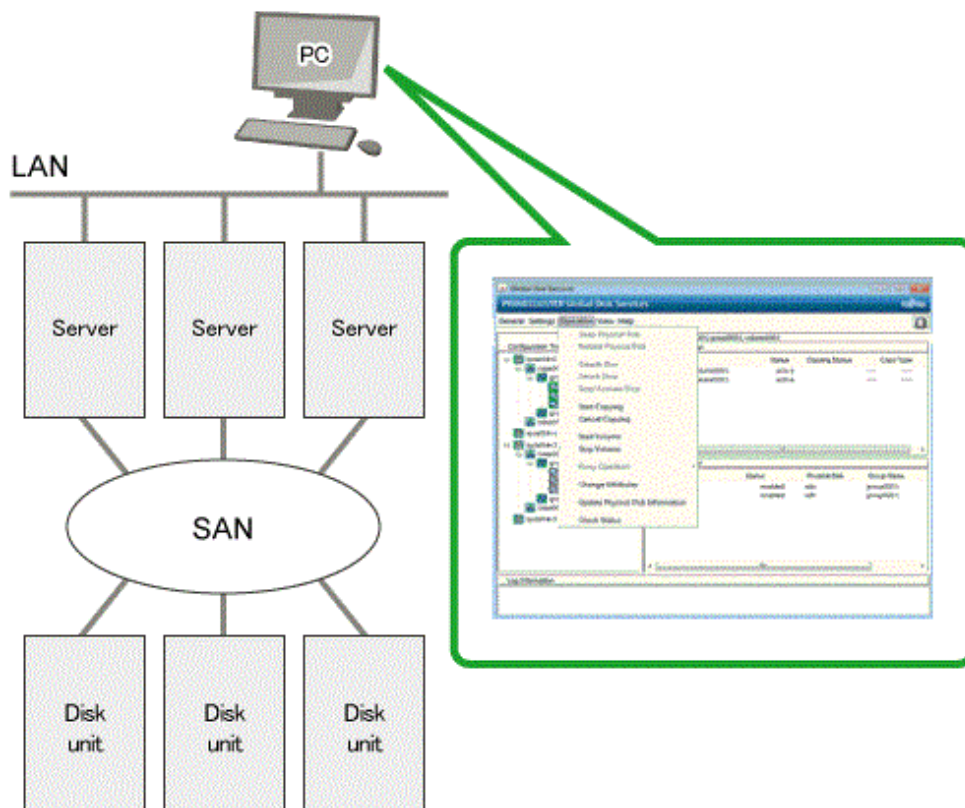
Functions improving disk manageability are explained below.

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 interfaces provide the GUI (Graphical User Interface) that allows easy and intuitive operation even for inexperienced Linux system users, and the CLI (Command Line Interface) that is useful for automatic processing, operation logging, and liaising with other applications.

Figure 1.12 Operation Management Interface



See

For the operation methods of the GUI, see "[Chapter 5 GDS Management View](#)."

For the usage methods of CLI, see "[Appendix B Command Reference](#)."

For the operations available on GUI and on CLI, see "[5.6 Operation](#)."

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 1.13 Centralized Disk Management



 Note

System Disk Management [EFI]

For the environment conditions where GDS can manage the system disk, see "2.1.1 Disk Class."

1.3.3 Name Management


In Linux System, disks are named in "sdX" format, where "X" is an alphabetical character assigned by the OS, and the administrator differentiates the disks by the assigned consecutive characters.

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 1.14 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 1.15 Single System Image Environment





Note

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.


1.3.5.1 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 the operation which has a certain disk unit group connected to three nodes; node A, node B, and node C. Some disk units are accessed only from node A and node B, and other disk units are accessed only from node B and node C.

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

Figure 1.16 Class Scope



1.3.5.2 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 1.17 Starting and Stopping Volume

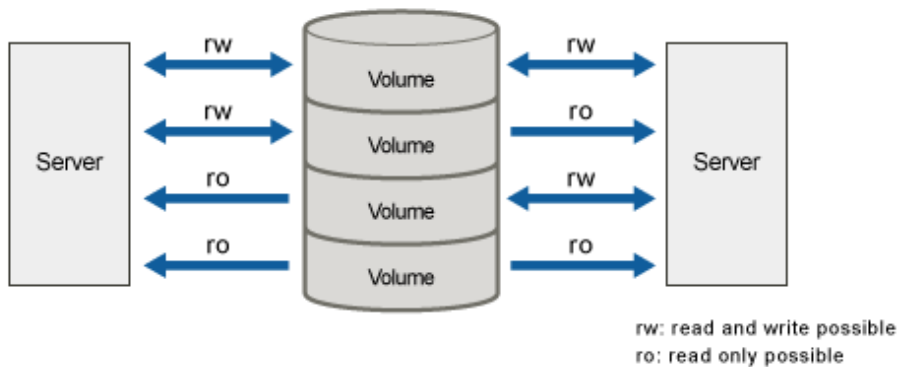


1.3.5.3 Access Mode

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 1.18 Access Mode



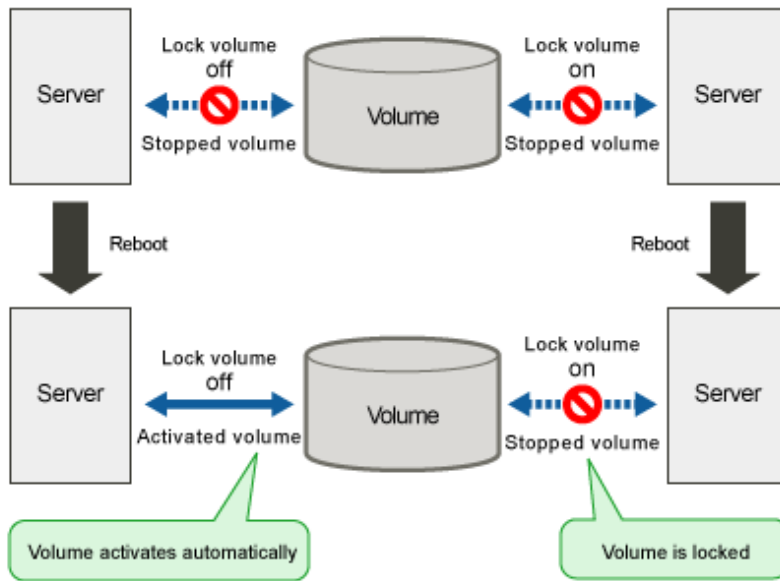
1.3.5.4 Lock Volume

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, a volume may be started unexpectedly by rebooting a node.

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 1.19 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 disk label (partition table).

For physical disks in a Linux system, in general, a maximum of 128 partitions can be used.

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


For details, see "2.1.4 Logical Volume" and "B.1.4 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.

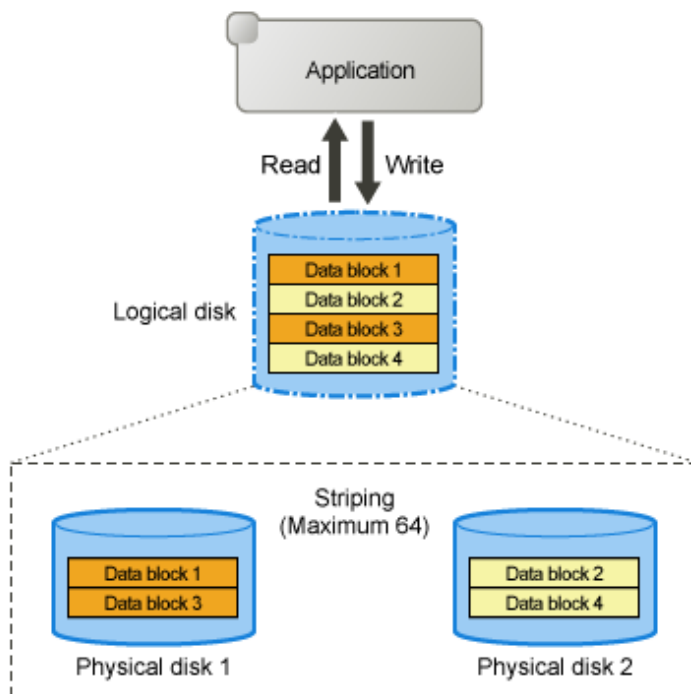
Figure 1.20 Disk Concatenation



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 1.21 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 1.22 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 1.23 Online Volume Expansion



Concatenation and Online Volume Expansion

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



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 "[3.12 Online Volume Expansion](#)."

Figure 1.24 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, online volume expansion will be available by migration of the volume to a disk with sufficient free space. For the volume migration, use the GDS Snapshot online disk migration function.



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

Figure 1.25 Online Disk Migration and Online Volume Expansion



Online System Volume Expansion [EFI]

By using a combination of this function and the GDS Snapshot's alternative boot environment creation function, online volume expansion is available for system volumes.



See

For the GDS Snapshot's alternative boot environment creation function, see "[1.5.6 Creating Alternative Boot Environment \[EFI\]](#)." For the procedures for expanding system volumes, see "[7.14.2 Volume Expansion Using Commands \[EFI\]](#)."

Operating Instructions

Use the `sdxvolume -S` command. For details, see "[B.1.4 sdxvolume - Volume Operations](#)."



Note

Points of Concern for Online Volume Expansion

See "[3.12 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, an operation mode where process loads do not influence on each other can also be established.

After applications that use detached slices end, it is necessary to reattach the detached slices to volumes in order to create snapshots again. 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 1.26 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 may 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 ext4, unmount the volume with the `umount(8)` command before creating a snapshot, and mount it with the `mount(8)` command afterwards. Then, 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.4 Just Resynchronization Mechanism \(JRM\)](#)."

1.4 GDS Snapshot Features

GDS Snapshot is an optional product of GDS 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 data erasure by the user or data crash due to an application malfunction. Therefore, data back up is mandatory. In a conventional 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 conventional 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 installed, 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 1.27 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 1.28 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 1.29 Liaising with Volume Management Function



The snapshot functions using a proxy volume are explained here.

1.5.1 Snapshot by Synchronization

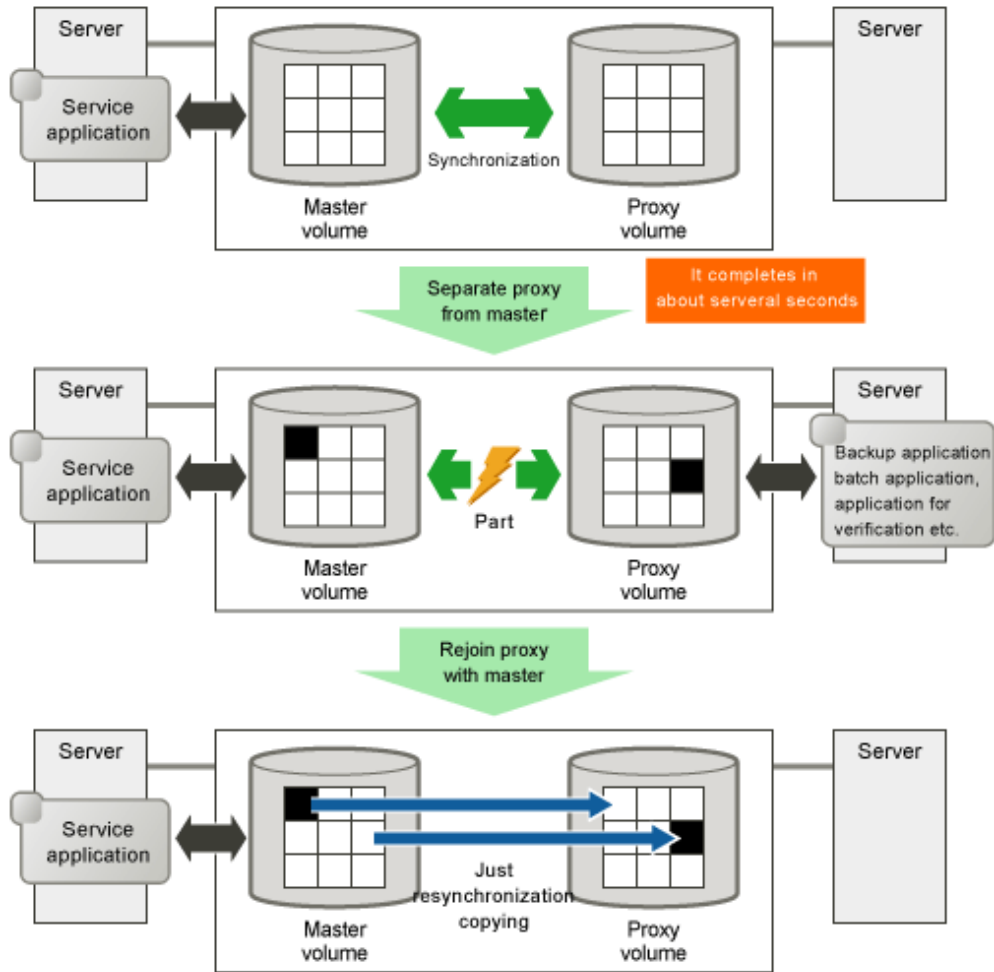
A master volume snapshot can be created by synchronizing data of the master volume and the proxy volume in advance and then separating the synchronized proxy volume from the master volume at any moment.

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 1.30 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 1.31 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.4 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 ext4, unmount the volume with the `umount(8)` command before creating a snapshot, and mount it with the `mount(8)` command afterwards. Then, 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 "[3.15 Online Backup and Restoration Using Proxy Volume](#)."



The Difference between a Mirrored Slice and Synchronized Proxy Volume

Although data matches on mirrored slices or synchronized master volumes and proxy volumes, the purposes of use 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 1.32 Difference between a Mirrored Slice and Synchronized Proxy Volume



1.5.2 Snapshot Function with No Load to Server/SAN



Note

The cooperation with the following copy functions are not supported in this version:

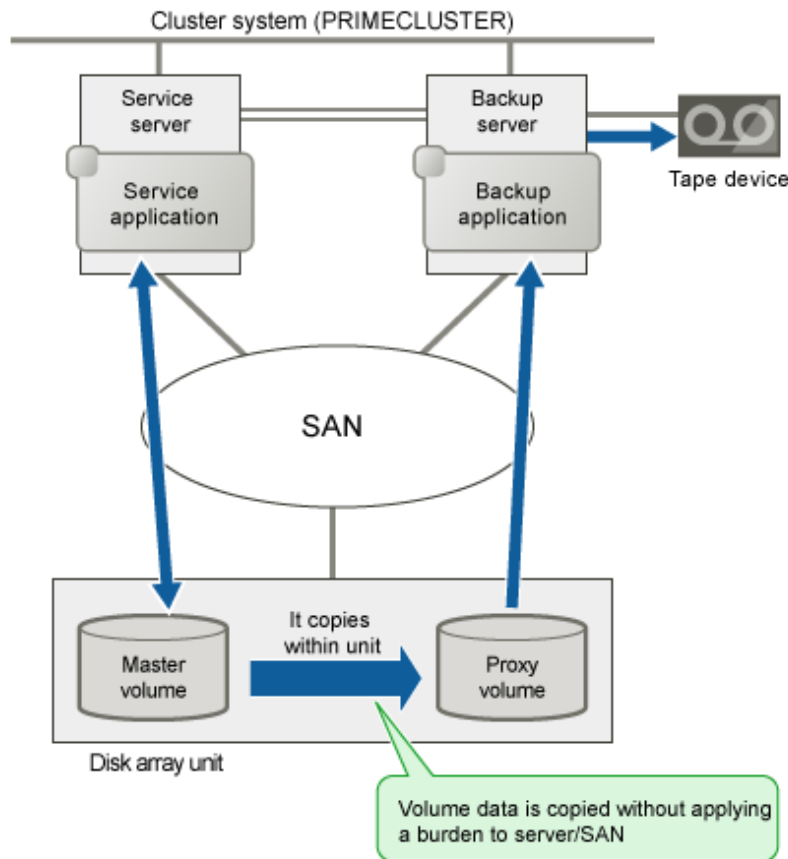
- Dell EMC TimeFinder
- Dell EMC SRDF

By cooperation with the Advanced Copy function of ETERNUS Disk storage system or TimeFinder and SRDF of Dell EMC storage units, 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 1.33 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 ETERNUS Disk storage system, snapshots can be created instantly at any given moment.

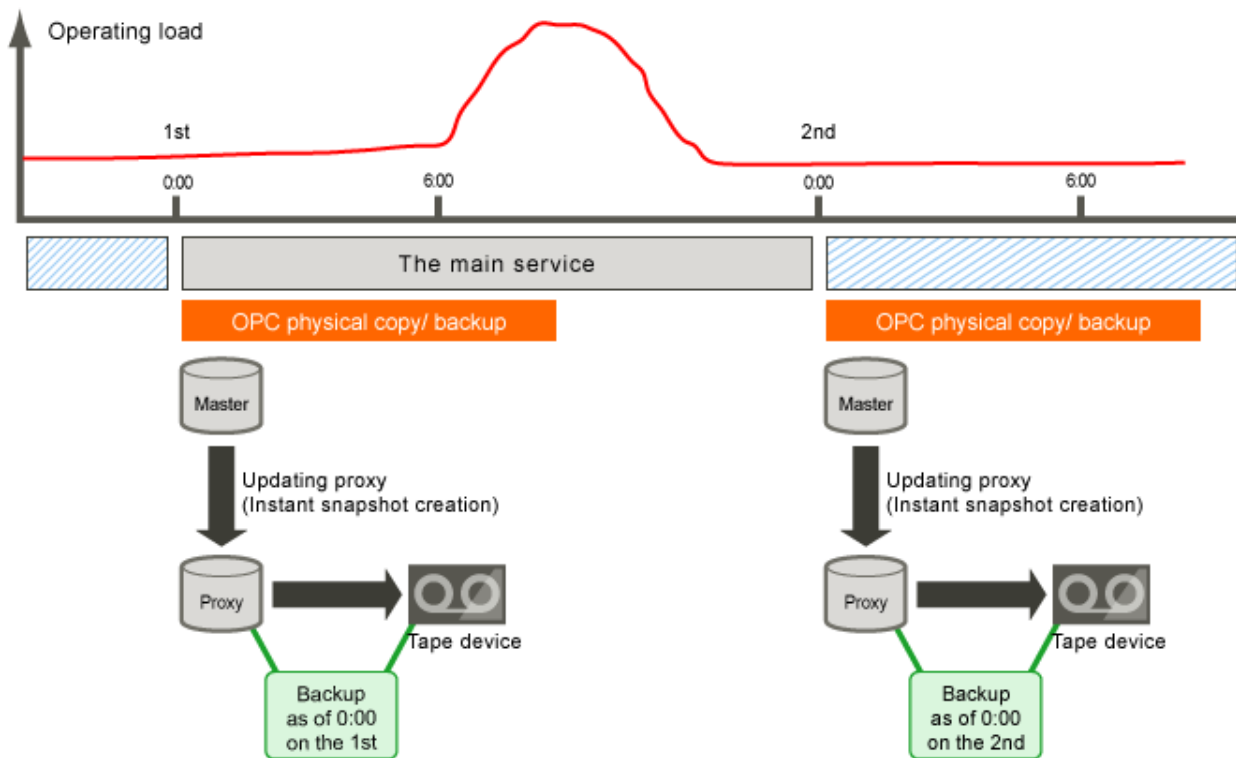
Unlike snapshot by synchronization, instant snapshot by OPC does not require synchronization with volumes in advance, so scheduling is unnecessary. However, as the copying process will take place within the disk array after the snapshot has been created, I/O performance of the disk will be reduced until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

QuickOPC function is also supported. When you use QuickOPC function in creating snapshot, only the differences (updated parts) are copied from the second time. This will reduce copy amount. Therefore, QuickOPC function is suitable for backup of large scale database and so on which require shortening of time.

Note

Instant snapshot by OPC is only available on ETERNUS Disk storage system with the OPC function.

Figure 1.34 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.

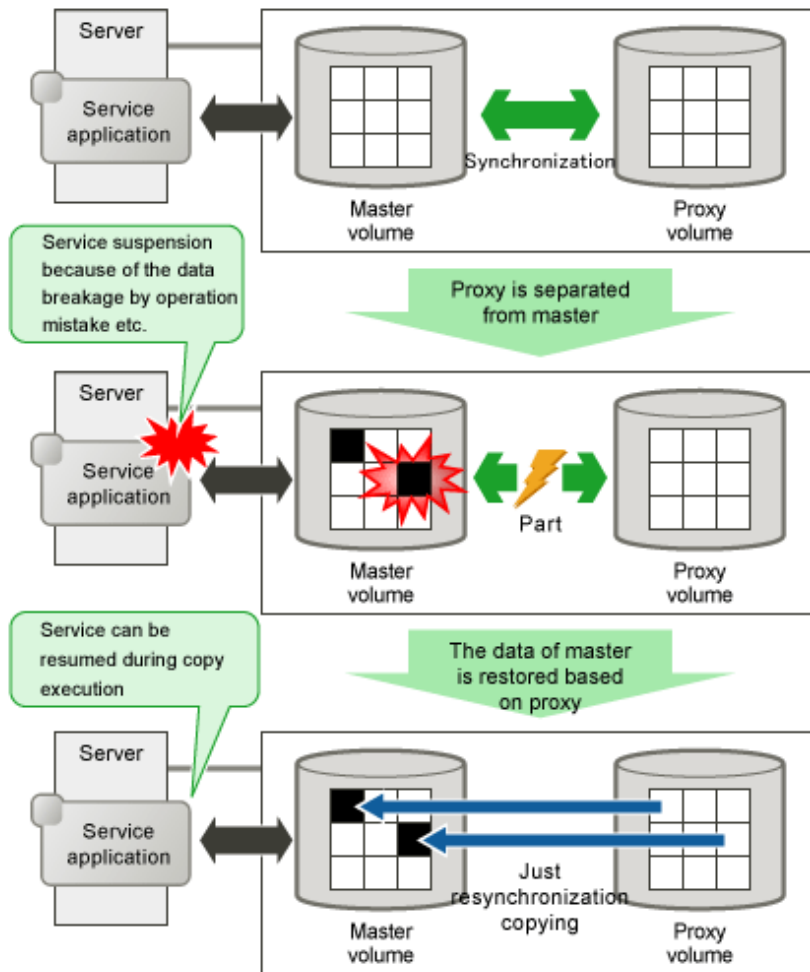


Note

System Volume Restoration [EFI]

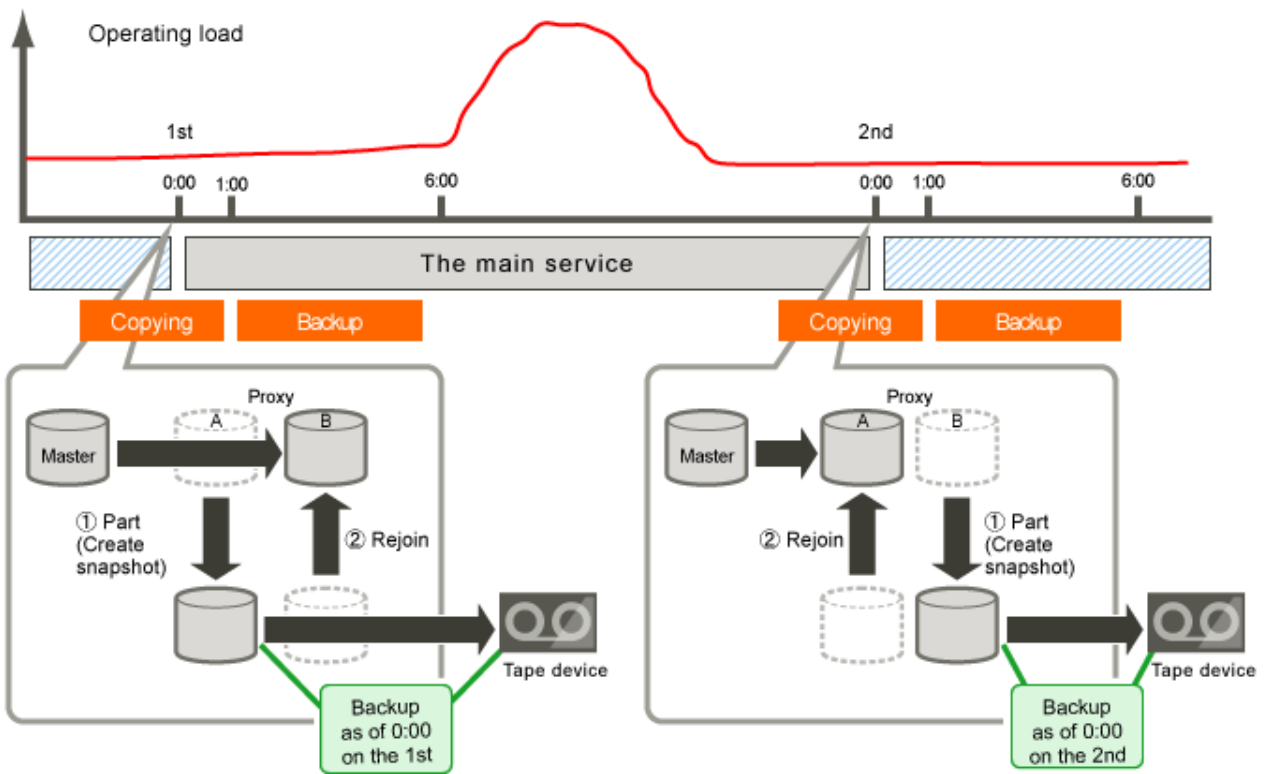
To execute instant restore, the master volume and the proxy volume must be stopped temporarily. An active system volume running as a file system such as /, /usr, and /var cannot be the restore destination or source of instant restore because it cannot be stopped. For the method of restoring system volumes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

Figure 1.35 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 1.36 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 1.37 Online Disk Migration



1.5.6 Creating Alternative Boot Environment [EFI]

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 restoration work for system volumes (`/`, `/usr`, `/var`, `/boot`, `/boot/efi`, 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.

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

Enable the system to boot from snapshot volumes (proxy volumes) instead of system volumes in the current boot environment. This configuration changes system files (`fstab`, `grub.cfg`, and `dracut.conf`) 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.



Note

SAN Boot Environment

For system disk snapshot in the SAN boot environment using ETERNUS Disk storage system, the Advanced Copy function is not available.



See

For examples of using an alternative boot environment, see "[7.5 Backing Up and Restoring System Disk through Alternative Boot Environment \[EFI\]](#)" and "[7.14.2 Volume Expansion Using Commands \[EFI\]](#)."

1.6 Shadow Volume



Note

Shadow volumes are not supported in this version.

Volumes in other domains (cluster systems or single servers) connected to a SAN can be recognized as volumes in the current node (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 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 mode and the procedure for backup and restore by shadow volumes through an external server, see "[3.20 Backing Up and Restoring through External Server.](#)"

Figure 1.38 Backup Server



1.6.2 Using Copy Function of 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 mode that uses the copy destination disk areas created with the copy function of a disk unit is available as well as the operation mode that uses proxy volumes, as replications of the volumes used in the service.

 See


For the operation mode and the procedure for backup and restore using the copy function of a disk unit and using shadow volumes, see "[7.11 Backing Up and Restoring through Disk Unit's Copy Functions.](#)"

Figure 1.39 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 1.40 Disaster Recovery System



1.7 GDS I/O Monitor Option Features

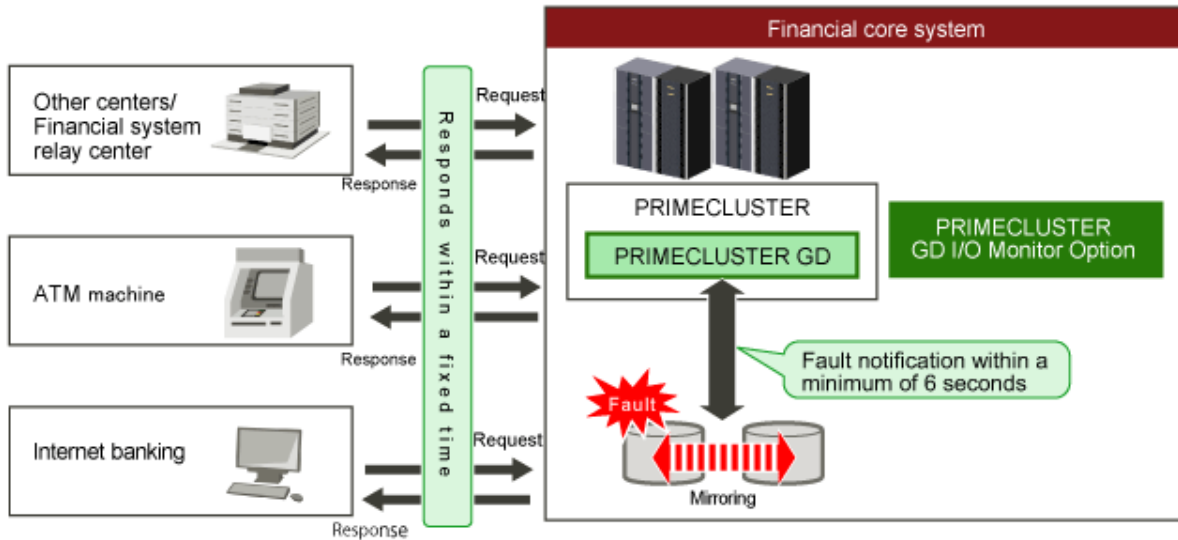
GDS I/O Monitor Option is an optional software of GDS that provides additional function to GDS.

The transaction process must be finished within a given time in the database of the core system.

However when hardware or transfer route error occurs in the I/O device, I/O response may be delayed in order for the retry process or recovery and degradation to operate.

GDS I/O Monitor Option provides I/O response time assurance to ensure I/O response within a given time (default time is within 77 seconds) in order to avoid effects caused by I/O delay of the disk.

Figure 1.41 I/O Response Time Assurance



- Features

The I/O responds within the specified time (default is within 77 seconds, shortest is within 6 seconds), even if the disk unit fails.

- Effectiveness

- Response can be realized in a short time for processing requests.
- System slowdown and stopping of services caused by I/O response delay can be prevented.

- Function overview

- Mirror volume (mirroring the disk)

The error generated disk due to the I/O response delay is detached from mirroring and the operation can be continued by using the other normal disk.

Even when an error occurs with the other disk, the operation continues by switching to standby node (system restart for non-clustered system).

- Single volume (not mirroring the disk)

When I/O response delay occurs, the operation continues by switching to standby node (system restart for non-clustered system).

- The I/O response time can be set for each class. The following settings are possible:

- Set different I/O response time on the disk that stores the system disk and user data.
- Divide the disks to be used for each application to set different I/O response time.

- System configuration

- Server: Physical environment of PRIMEQUEST 3000/2000 Series
- Disk: Server built-in disk, ETERNUS DX series (FC connection)
- Volume configuration of GDS: Mirror volume, single volume

Note

GDS I/O Monitor Option cannot be used in a system that uses a virtual environment or a cloud environment.

See

- For operation details, see "[1.8 I/O Response Time Assurance \[IOMonitor\]](#)."
- For general notes on system configuration and design, see "[3.23 I/O Response Time Assurance \[IOMonitor\]](#)."
- For the settings method, see "[6.12 I/O Response Time Assurance \[IOMonitor\]](#)," "[8.3 Changing the Settings of I/O Response Time Assurance \[IOMonitor\]](#)," and "[8.4 Deleting the Settings of the I/O Response Time Assurance \[IOMonitor\]](#)."
- For general notes on operation and maintenance, see "[7.15 I/O Response Time Assurance \[IOMonitor\]](#)."

1.8 I/O Response Time Assurance [IOMonitor]

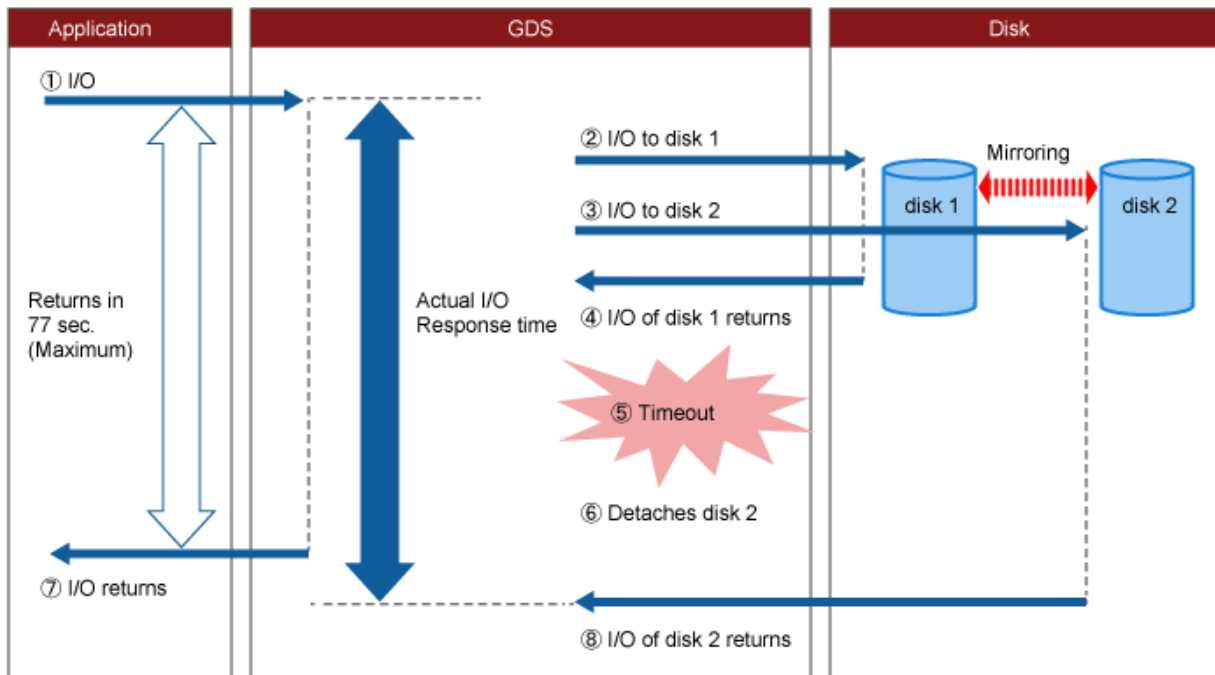
Here the I/O operations when setting up I/O Response Time Assurance Function will be explained.

1.8.1 I/O response delay for one slice on the mirror volume

The graph below shows an example of the I/O operations when an I/O response delay occurs on a disk with the I/O response time set to 77 seconds.

The I/O returns in a maximum time of 77 seconds to the application even if the I/O response of the other disk (disk 2 in this example) is delayed. (After this, the I/O of disk 2 returns.)

Figure 1.42 I/O Response Time Assurance (When I/O response delay occurs on one slice of the mirror volume)



[Process]

- ① Application issues I/O request to GDS volume
- ② GDS driver issues I/O request to disk 1
- ③ GDS driver issues I/O request to disk 2
- ④ I/O of disk 1 returns normally
- ⑤ I/O of disk 2 has timed out
- ⑥ GDS driver detaches disk 2
- ⑦ I/O issued from the application returns in a maximum time of 77 seconds
- ⑧ I/O of disk 2 returns

GDS: Global Disk Services


1.8.2 I/O response delay on all slices of the mirror volume

The figure below shows the I/O operations when an I/O delay has occurred on all slices that make up the mirror volume if system panic mode is enabled.

When the specified I/O response time is exceeded, the I/O of disk1 and disk2 are timed out and the GDS I/O Monitoring Option panics the system. Cluster failover occurs because of this. The time necessary for failover is shorter than if system panic mode is disabled.

For cluster systems, it is recommended to set the system panic mode to enable.

Figure 1.43 I/O Response Time Assurance (When an I/O response delay occurs on all slices of the mirror volume)



[Process]

- ① Application issues I/O request to GDS volume
- ② GDS driver issues I/O request to disk 1
- ③ GDS driver issues I/O request to disk 2
- ④ I/O of disk 1 and disk 2 has timed out
- ⑤ I/O aborted
- ⑥ System has panicked
- ⑦ Cluster failover

GDS: Global Disk Services

 Information

If system panic mode is set to disabled, if an I/O response delay or I/O error occurs on all the slices making up the mirror system, or if the class was closed, an error occurs in the volume stop processing on the previous operating node while cluster failover processing is being done and because of the double fault RMS panics the previous operating node.


1.8.3 I/O response delay for the slice of a single volume

The image below is an example of the I/O operations where an I/O delay has occurred on the slice of a single volume with system panic mode is enabled.

If the specified I/O response time is exceeded, the I/O of disk1 will time out and the GDS I/O Monitor Option will panic the system. Because of this the cluster will fail over. The time necessary for failover is less than if the system panic mode is set to disabled.

For cluster systems, it is recommended to set the system panic mode to enable.

Figure 1.44 I/O Response Time Assurance (When an I/O response delay occurs on the slice of a single volume)



[Process]

- ① Application issues I/O request to GDS volume
- ② GDS driver issues I/O request to disk 1
- ③ I/O of disk 1 has timed out
- ④ I/O aborted
- ⑤ System has panicked
- ⑥ Cluster failover

GDS: Global Disk Services

 Information

.....
 If system panic mode is set to disabled, if an I/O response delay or I/O error occurs on all the slices making up the mirror system, or if the class was closed, an error occurs in the volume stop processing on the previous operating node while cluster failover processing is being done and because of the double fault RMS panics the previous operating node.

Chapter 2 Structure

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.

By installing the GDS I/O Monitor Option which is optional software for GDS, the I/O response time guarantee function for the GDS class can be set.

This chapter systematically describes the virtual object structure that is the underlying basis of GDS, GDS Snapshot, and GDS I/O Monitor Option 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.

In GDS, these objects are called disk classes, SDX disks, disk groups, logical volumes, and logical slices.

Figure 2.1 Figure: Interrelation of SDX Objects



2.1.1 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 "[8.5 Migration from Single Nodes to 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 [EFI]

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

The following disks can be managed; the system disk, mirror destination disks, spare disks, and disks on which snapshots of the system disks are created.

For the environment conditions where the root class can be used, see "[3.4 Managing System Disks](#)."

Local

Objects managed in this type of class are available only to 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.

I/O Response Time Assurance [IOmonitor]

Attribute indicating the operation mode of the I/O Response Time Assurance. Either of the following values can be set.

on

I/O Response Time Assurance is enabled.

off

I/O Response Time Assurance is disabled.

I/O Response Time [IOmonitor]

When the I/O Response Time Assurance is enabled, this attribute indicates the maximum time it takes until the I/O responds.

System Panic Mode [IOmonitor]

This attribute is specified to panic or not to panic the system when the I/O Response Time Assurance is enabled if one of the following incidents occurs:

- I/O errors occur on all the slices that make up the volume.
- The class was closed.

Either of the values below can be set.

on

System panic mode is enabled.

off

System panic mode is disabled.

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 "[6.6.1.3 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 "[9.3.1.4 Removing Class](#)."

Status Display

Class status can be displayed with the `sdxinfo` command. For details on GDS Management View, see "[7.1 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 "[8.2.1.2 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 "[D.1.4 Class Status Abnormality](#)."

Reference

There are the following points of concern for classes.

Rules

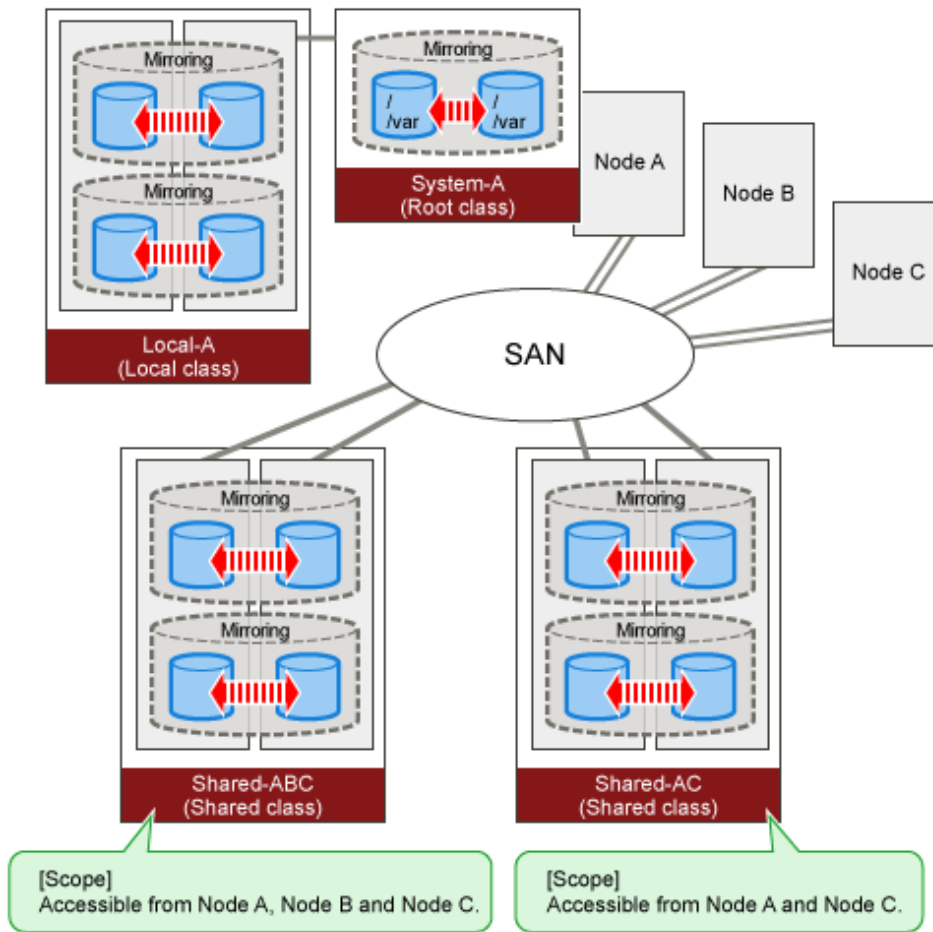
[2.3.1 Object Name](#)

[2.3.2 Number of Classes](#)

Important Points

[3.4 Managing System Disks](#)

Figure 2.2 Disk Class



2.1.2 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.

Netmirror

The disk has been connected to a netmirror group.

Note

If there is no description, the explanation of a disk for a netmirror type is the same as that for a mirror type.

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 [EFI]

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 "[6.6.1.3 Class Configuration](#)."

Delete

A disk can be deleted from a class with the `sdxdisk -R` command.
For details on GDS Management View, see "[9.3.1.4 Removing Class](#)."

Connect

A disk can be added to a certain group with the `sdxdisk -C` command.
For details on GDS Management View, see "[8.2.1.3 Group Configuration](#)."

Disconnect

A disk can be disconnected from the group with the `sdxdisk -D` command.
For details on GDS Management View, see "[8.2.1.3 Group Configuration](#)."

Swap

A disk is made ready for disk swap with the `sdxswap -O` command.
For details on GDS Management View, see "[7.3.1.2 Operation Procedure](#)."

Recover

Swapped disks can be recovered with the `sdxswap -I` command.

For details on GDS Management View, see "[7.3.1.2 Operation Procedure.](#)"

Status Display

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

For details on GDS Management View, see "[7.1 Viewing Configurations/Statuses and Monitoring Statuses.](#)"

Change Attributes

Disk attribute values can be changed with the `sdxattr -D` command.

For details on GDS Management View, see "[8.2.1.2 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

Since disk identification information (class and disk names) is not set, data on the disk is inaccessible.

When disk identification information stored in the private slice of a disk is not set at the time of booting the system, the disk will be in DISABLE status. For example, if a user reboots the system after formatting a disk mistakenly, it will be in this status. If this happens, all slices related to the disk are made the NOUSE status.

For details when the disk is in DISABLE status, see "[D.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

[2.3.3 Number of Disks](#)

[2.3.6 Number of Keep Disks \[EFI\]](#)

[2.3.9 Disk Size](#)


Important Points

[3.6 Initializing Disk](#)

[3.7 Keep Disk \[EFI\]](#)

[A.1 Restraining Access to Physical Special File](#)

Figure 2.3 SDX Disk



2.1.3 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. The type attribute takes either of the following values:

- Mirror
- Netmirror
- Stripe
- Concat
- Switch

By connecting multiple disks to a group, one of the following functions can be performed for each group depending on the type attribute value of the group:

- Mirroring
- Mirroring among servers
- Striping
- Concatenation
- Switching

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.



Note

If there is no description, the explanation of a group for a netmirror type is the same as that for a mirror type.

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 following.

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 among servers can be supported.

netmirror

Disks belonging to the same group will be mirrored to each other among servers. A maximum of 2 disks can be connected to a group collectively. In other words, a maximum of two-way multiplex mirroring among servers 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.

concatenation

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. The size you can set is (a value of two raised to the power) x 512 bytes and it conforms the following conditions:

- Minimum value: 512 bytes
- Maximum value: the minimum value among the following
 - (A value of two raised to the 30th power) x 512 bytes (= 512 GB)
 - Available size of the smallest disk in a group
 - Available size of the smallest lower group in a group

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 "[6.6.1.5 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 "[9.3.1.3 Removing 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 "[8.2.1.3 Group Configuration](#)."

Disconnect

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

For details on GDS Management View, see "[8.2.1.3 Group Configuration](#)."

Status Display

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

For details on GDS Management View, see "[7.1 Viewing Configurations/Statuses and Monitoring Statuses](#)."

Change Attributes

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

For details on GDS Management View, see "[8.2.1.3 Group Configuration](#)."

Reference

There are the following points of concern for groups.

Rules

[2.3.4 Number of Groups](#)

[2.3.7 Creating Group Hierarchy](#)

Important point

[3.13 Disk Switch](#)

Guide


[3.1.1 Guidelines for Mirroring](#)

[3.1.2 Guidelines for Striping](#)

[3.1.3 Guidelines for Concatenation](#)

[3.1.4 Guidelines for Combining Striping with Mirroring](#)

Figure 2.4 Disk Group



2.1.4 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 six 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.

- Netmirror Volume Created in a Netmirror Group

When multiple disks are connected to the netmirror group, data will be redundant through the mirroring function among servers.

- 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.

For physical disks in a Linux system, in general, a maximum of 128 partitions can be used. 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 and be used.

Note

If there is no description, the explanation of a volume for a netmirror type is the same as that for a mirror type.

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 following.

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 following.

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 consisting the volume is registered with the disk label. The value can be set to one of the following. 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 disk label. 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 disk label. When the volume is a switch volume, slices on all of the disks connected to the switch group will be registered with the disk label. 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 disk label.

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 "[6.6.1.6 Volume Configuration](#)."

Delete

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

For details on GDS Management View, see "[9.3.1.2 Removing Volume](#)."

Start

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

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

Stop

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

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

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 "[7.1 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 "[8.2.1.4 Volume Configuration](#)."

Copy control

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

For details on GDS Management View, see "[7.2.1 Using GDS Management View](#)."

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 "[D.1.3 Volume Status Abnormality](#)."

Reference

There are the following points of concern for volumes.

Rules

[2.3.5 Number of Volumes](#)

[2.3.11 Volume Size](#)

Important Points

[3.10 Volume Access Mode](#)

[3.11 Data Consistency at the time of Simultaneous Access](#)

[3.12 Online Volume Expansion](#)

[A.4 Just Resynchronization Mechanism \(JRM\)](#)

Figure 2.5 Logical Volume

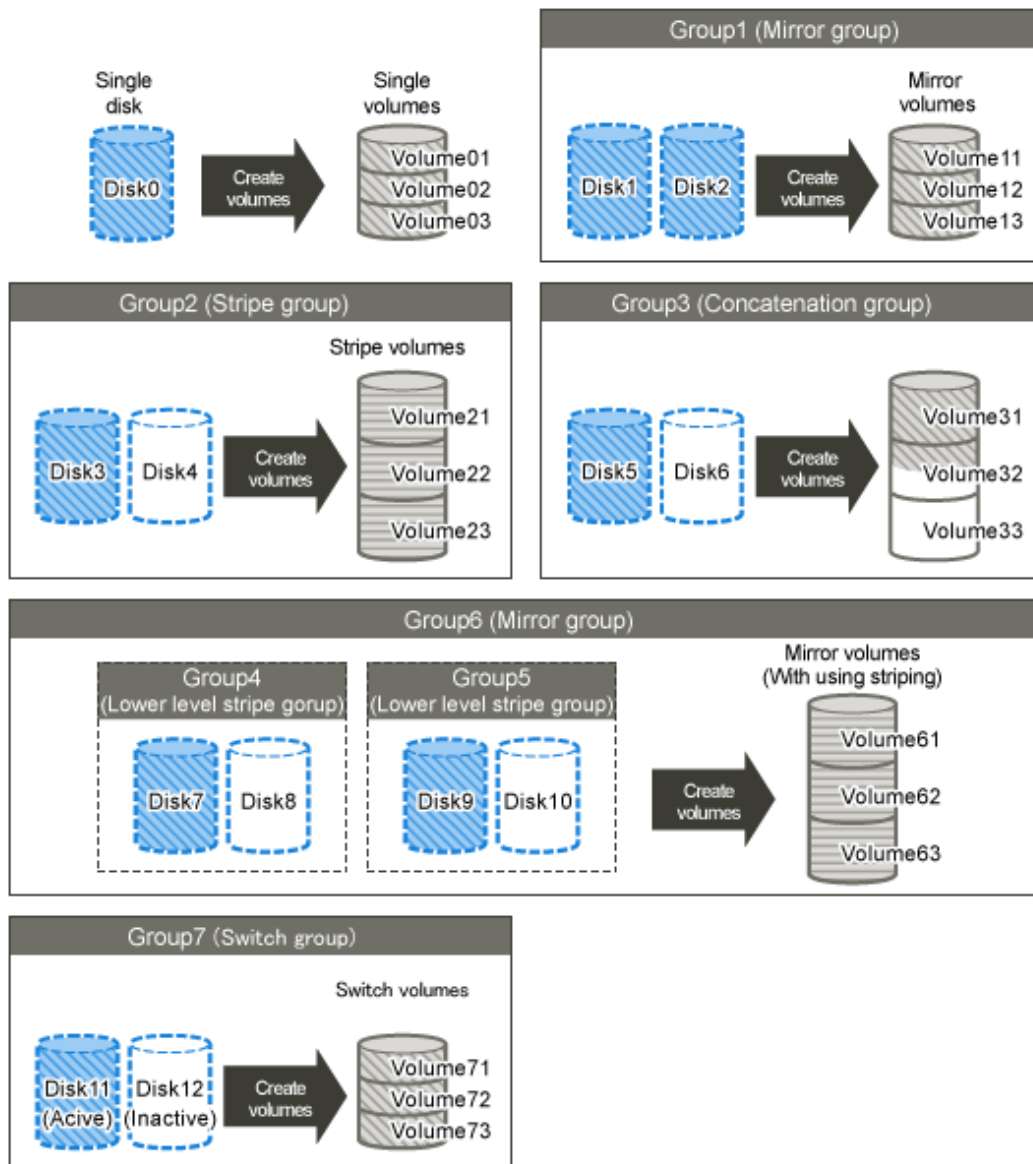



Figure 2.6 Logical Partitioning (Physical Slice Attribute)



2.1.5 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.



Note

If there is no description, the explanation of a slice belonging to a mirror volume is the same as a slice of a netmirror volume.

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 following.

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 "[7.7.1.2 Operation Procedure.](#)"

Attach

The temporarily detached slice can be attached to the volume again with the `sdxslice -R` command. For details on GDS Management View, see "[7.7.1.2 Operation Procedure.](#)"

Activate

The detached slice can be activated with the `sdxslice -N` command. For details on GDS Management View, see "[7.7.1.2 Operation Procedure.](#)"

Stop

The detached slice can be stopped with the `sdxslice -F` command. For details on GDS Management View, see "[7.7.1.2 Operation Procedure.](#)"

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 "[7.1 Viewing Configurations/Statuses and Monitoring Statuses.](#)"

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).

However, in the case of slices of netmirror type, the data may be in an invalid and inaccessible state. For details, see "[7.16.3 Operation when the LEFTCLUSTER Node Exists.](#)"

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).

However, in the case of slices of netmirror type, the data may be in an invalid state. For details, see "[7.16.3 Operation when the LEFTCLUSTER Node Exists.](#)"

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

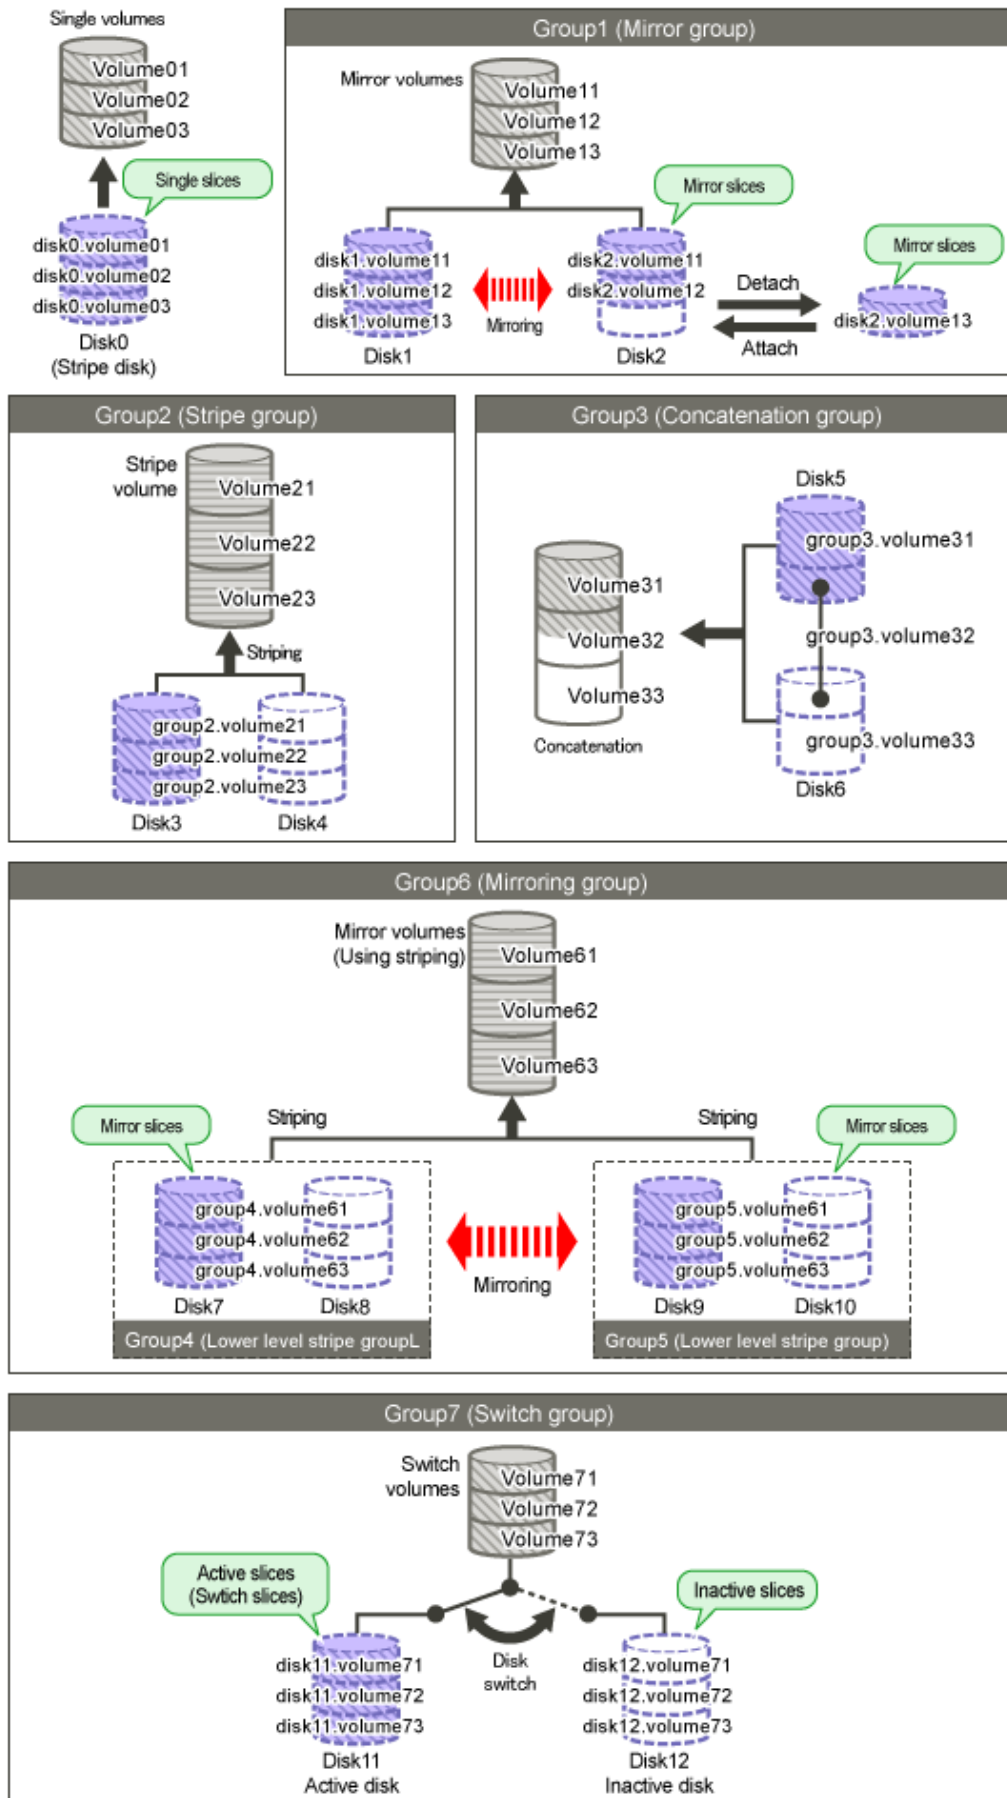
There are the following points of concern for slices.

Important Points

[3.14 Creating Snapshot by Slice Detachment](#)

[A.4 Just Resynchronization Mechanism \(JRM\)](#)

Figure 2.7 Logical Slice



2.2 GDS Snapshot Objects

2.2.1 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 "[6.10.1.1 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 "[7.8.1.2 Operation Procedure](#)."

Rejoin

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

For details on GDS Management View, see "[7.8.1.2 Operation Procedure](#)."

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 "[7.10.1.2 Operation Procedure](#)."

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 "[8.7.1.1.2 Operation Procedure](#)."

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 "[6.10.1.2 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 "[7.9.1.2 Operation Procedure](#)."

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 "[7.10.1.2 Operation Procedure](#)."

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 [EFI]

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 canceled and they will return to individual objects with the `sdxproxy Break` command.

For details on GDS Management View, see "[9.1 Breaking 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

[2.3.12 Proxy Configuration Preconditions](#)

[2.3.13 Number of Proxy Volumes](#)

[2.3.14 Proxy Volume Size](#)

[2.3.15 Proxy Group Size](#)

Important Points

[3.16 Difference between Mirror Slice and Proxy Volume](#)

[3.17 Ensuring Consistency of Snapshot Data](#)

[3.18 Instant Snapshot by OPC](#)

[3.19 Using Dell EMC TimeFinder or Dell EMC SRDF in Proxy Configuration](#)

[A.4 Just Resynchronization Mechanism \(JRM\)](#)

[A.7 Object Operation When Using Proxy](#)

[A.8 Using Advanced Copy Function in Proxy Configuration](#)

Figure 2.8 Operating Proxy Objects




Figure 2.9 Swapping Master Volume and Proxy Volume





Figure 2.10 Swapping Master Group and Proxy Group



2.2.2 Shadow Object



Note

Shadow objects are not supported in this version.

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 one of the following conditions can be managed as shadow objects.

- Disks registered with a class (local class or shared class) of GDS in another domain and managed as SDX disks
- Disks to which the private slice was copied from an SDX disk registered with a class (local class or shared class) of GDS 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.9 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 "[2.3.1 Object Name](#)."

2.2.2.1 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 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 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 "7.1 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.

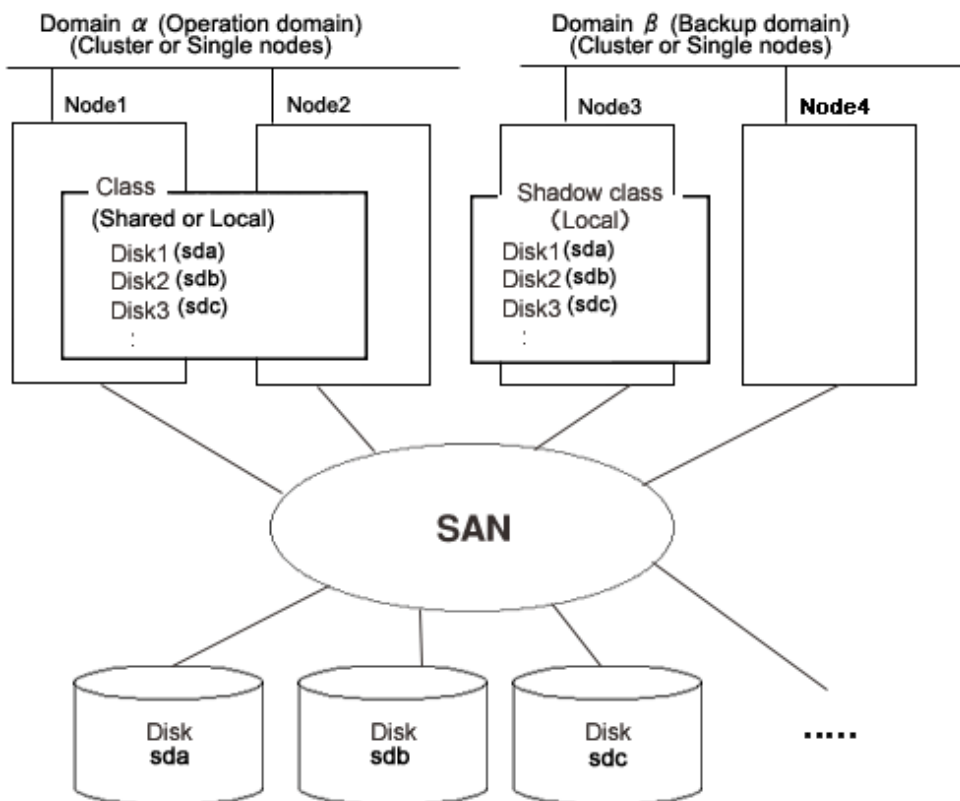
Rules

[2.3.1 Object Name](#)

[2.3.2 Number of Classes](#)

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.

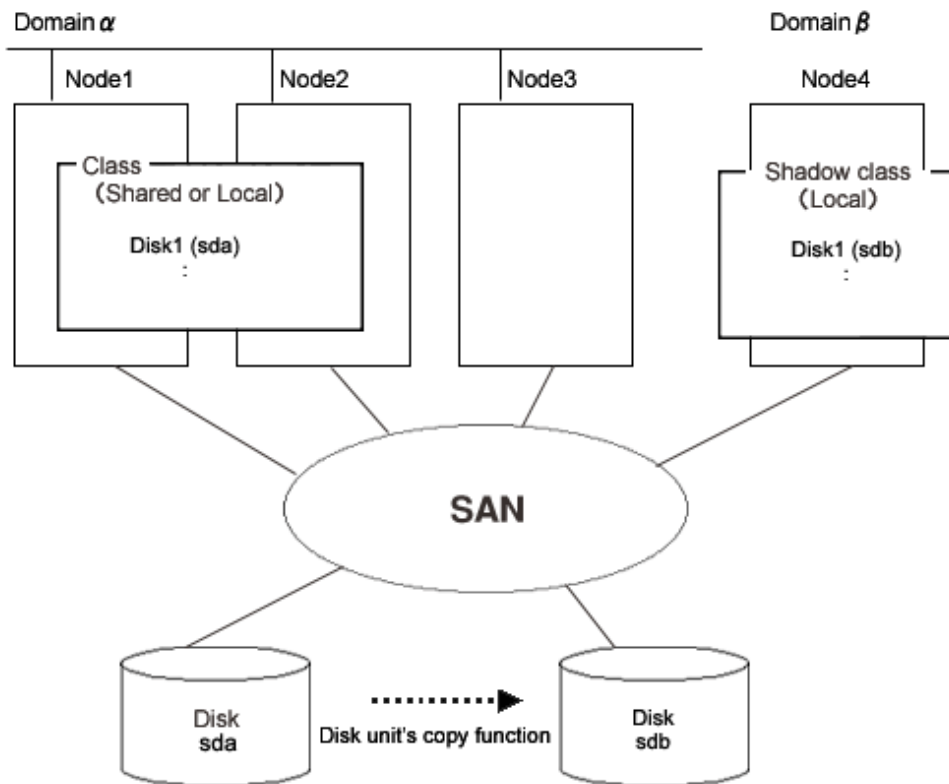
Figure 2.11 Example of Common Configuration with a Shadow Class



The disk (sdb in the figure below) to which the private slice has been copied from an SDX disk (sda 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 2.12 Example of Configuration Using Disk Unit's Copy Function and a Shadow Class



Note

Physical Device Name

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

2.2.2.2 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 "[7.1 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

[2.3.3 Number of Disks](#)

[2.3.9 Disk Size](#)


Important Points

[3.6 Initializing Disk](#)

A.1 Restraining Access to Physical Special File

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 2.13 An SDX Disk and a Shadow Disk



When the following conditions are met, a physical disk (sdb in the figure below) can be managed as a shadow disk by being registered with a shadow class in certain domain (domain alpha) or another domain (domain beta). A shadow disk must be named as same as the SDX disk (sda in the figure below).

- A physical disk (sda in the figure below) is registered with a class in domain alpha and managed as a SDX disk.
- An sda private slice is copied to sdb by the copy function of a disk device.

Figure 2.14 Disk Unit's Copy Function and a Shadow Disk



Note

Physical Device Name

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

2.2.2.3 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. The size you can set is (a value of two raised to the power) x 512 bytes and it conforms the following conditions:

- Minimum value: 512 bytes
- Maximum value: the minimum value among the following
 - (A value of two raised to the 30th power) x 512 bytes (= 512 GB)
 - Available size of the smallest disk in a group

- Available size of the smallest lower group in a group

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 "[7.1 Viewing Configurations/Statuses and Monitoring Statuses.](#)"

Reference

There are the following points of concern for shadow groups.

Rules

[2.3.4 Number of Groups](#)

[2.3.7 Creating Group Hierarchy](#)

2.2.2.4 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 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.9 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 the disk label.

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 "[7.1 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

[2.3.5 Number of Volumes](#)

[2.3.11 Volume Size](#)

Important Point


[3.10 Volume Access Mode](#)

[A.9 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 2.15 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.

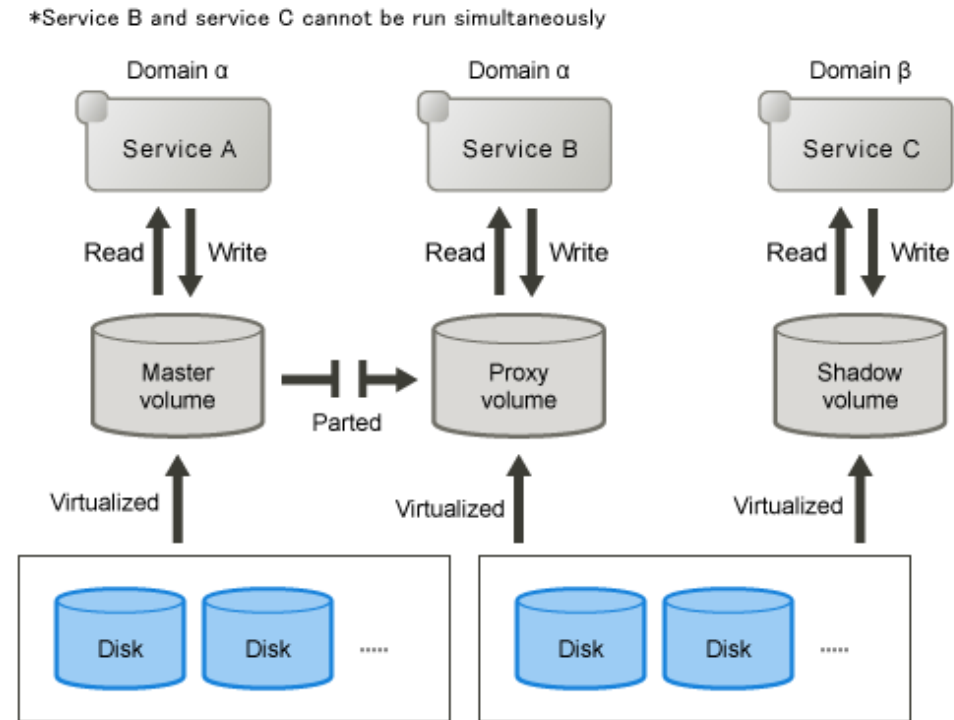
Figure 2.16 Mirror Slice Temporary Detachment and a Shadow Volume

* 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 2.17 A Proxy Volume and a Shadow Volume



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 2.18 Disk Unit's Copy Function and a Shadow Volume



2.2.2.5 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 "[7.1 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 "[D.1.1 Slice Status Abnormality](#)."

2.3 Rules

2.3.1 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 "[8.5 Migration from Single Nodes to 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.

2.3.2 Number of Classes

The number of root classes [EFI] 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. [EFI]
- 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 "[2.3.3 Number of Disks](#)" and "[2.3.5 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.

2.3.3 Number of Disks

The number of disks you can register with one class has the following limitations:

- To root class [EFI], you can register up to 100 disks.
- To local class 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 netmirror group, a maximum of 2 disks can be connected. In other words, a maximum of two-way multiplex mirroring among servers is supported.
- 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.
- To a proxy group of a root class [EFI], only one disk can be connected.

2.3.4 Number of Groups

The number of groups you can create within one class has the following limitations:

- Within root class [EFI], you can create up to 100 groups.
- Within local class or shared class, you can create up to 1024 groups.

2.3.5 Number of Volumes

There are the following limits to the number of volumes that can be created within a group in the root class [EFI]:

- The number of volumes with the physical slice attributes set to "on" is limited to a maximum of 14.
- 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 4 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 the disk label.
- When you perform the proxy operation per group, the number of volumes within the master group must be 400 or less.

In addition, there are the following limits to the number of volumes that can be created within a class:

- For the root class [EFI], the number is limited to a maximum of 256.
- For local and shared classes, 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.

2.3.6 Number of Keep Disks [EFI]

With root class, you can register up to 100 keep disks.

2.3.7 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)

A netmirror group cannot be nested.

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.

2.3.8 Disk Label

Disk label is an area at the front of each disk on which geometry and slice information are stored.

In RHEL, two types of disk labels below can be used:

- MSDOS label (Also called MBR label)
- GPT label



Information

Creation and change of the disk labels can be executed with the parted(8) command.

Disk label types

Disk label types of disks registered with a class are indicated as below:

For a root class [EFI]

It will be GPT type.

For local and shared classes

It will be the same as the disk label type of a class. Disk label type of a class is determined by the size of the disk first registered with the class and the disk label type as below.

Disk first registered with the class		Disk in the class
Size	Disk label type before registered	Disk label type after registered (Disk label type of a class)
Smaller than 2 TB	No label MSDOS type	MSDOS type
	GPT type	GPT type
2 TB or larger	No label MSDOS type GPT type	

Also, the size of a disk that can be registered with the class is determined by the disk label type of a class.

- For MSDOS type: Smaller than 2 TB

- For GPT type: Smaller than 2 TB or 2 TB or larger (both available)

The disk label type of a class can be viewed in the LABEL field of Class information that is displayed with the `sdxinfo -C -e` label command.

2.3.9 Disk Size

Available size of a disk

Within 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.

Size of a private slice

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. In the shared class and the local class, the size of a private slice in the class 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. Below estimate values represent the maximum size of the private slice that is reserved in the physical disk. The private slice size never exceeds the estimated values.

- In the root class [EFI]

10.5 MB

- In the shared class and the local class

- When the size of the disk first registered with the class is 10 GB and below:

32 MB

- When the size of the disk first registered with the class is over 10 GB:

32 MB + (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 or netmirror volume is created after registering a disk larger in size than the first registered disk with the class. Therefore, the following points should be considered when registering the disk to the class.

- When creating a mirror volume

When the JRM log area for the volume cannot be reserved, JRM cannot be used for the volume. Therefore you are recommended to first register a disk of the maximum size with the class.

- When creating a netmirror volume

When the JRM log area cannot be reserved, a netmirror volume cannot be created.

In a class where a netmirror volume is created, 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 the values stored in their private slices. 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.

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.

1 TB or larger disks

1 TB or larger disks can be managed in local and shared classes.

For 1 TB or larger disks, register disks with the class where the disk label type is GPT and manage them. For details, see "[2.3.8 Disk Label](#)."

2.3.10 Group Size

Available size of a group

Within the group area, the available capacity for creating volume area is called available size of a group.

- Available size of a mirror group

The same as the available size of the smallest disk or lower level group that is connected to the mirror group.

- Available size of the stripe group

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.

- Available size of the concatenation group

Is the total size of available size of disks connected to the concatenation group.

Group Size

Added size of one private slice to the available size of a group.



.....
For details about available size of a disk, cylinder disk and size of a private slice, see "[2.3.9 Disk Size](#)."
.....

2.3.11 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 [EFI]

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.

2.3.12 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.
- The proxy of the root class is performed per group.

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 QuickOPC function is available.
- To the master and proxy groups, no lower level group is connected.
- To create snapshots in group unit, the volume layouts (offsets and sizes) of the master group and the proxy group are consistent.
- See also "[A.8 Using Advanced Copy Function in Proxy Configuration](#)" and "[3.18 Instant Snapshot by OPC.](#)"

2.3.13 Number of Proxy Volumes

To one master volume, multiple proxy volumes can be related.

For the root class, the number of proxy volumes that can be related to one master volume is up to two.

For the local and shared classes, 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.

2.3.14 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 [EFI]

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.

2.3.15 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 [EFI], 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.

Chapter 3 Notes on Design

This chapter describes notes when designing the system where GDS, GDS Snapshot, and GDS I/O Monitor Option are used.

3.1 Guidelines

3.1.1 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 are 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.

3.1.2 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. This 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.

- Striping configurations cannot be used in RHEL8.

3.1.3 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.

- Concatenation configurations cannot be used in RHEL8.

3.1.4 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 "[3.1.1 Guidelines for Mirroring](#)" and "[3.1.2 Guidelines for Striping](#)" also.

3.2 DM-MP (Device Mapper Multipath)

DM-MP is the standard OS multipath function of RHEL.

mpath devices of DM-MP can be managed in local and shared classes.

When using DM-MP, note the following points.

- A device to which DM-MP is applied for the virtio block device (/dev/vdX) on the KVM guest, cloud environment (except an FJcloud-Baremetal environment), and RHOSP environment cannot be registered to GDS.
- Do not use DM-MP for devices (local disks or iSCSI devices) used for the mirroring among servers. If you configure DM-MP, register these devices in the black list.
- If there are multiple mpath devices in the system, registering a disk to a shared class and a local class may fail.
The number of mpath devices that registering a disk fails depends on the system processing performance. Only as a guide, it is 12,000/the number of multipaths. If the number of multipaths is 2, it is 6,000.

3.3 Dell EMC Storage Units

In local and shared classes, disks of Dell EMC storage units can be managed.

For available Dell EMC storage units, contact field engineers.

Multipath Software

The following multipath software can be used:

- In a physical environment
Dell EMC PowerPath

- In a VMware environment
Dell EMC PowerPath/VE or VMware NMP (Native Multipathing Plugin)

Devices that Cannot be Managed

The following devices cannot be managed by GDS.

- native devices configuring emcpower devices
- BCV (Business Continuance Volume) devices
- SRDF target (R2) devices
- GateKeeper devices
- CKD (Count Key Data) devices
- VC MDB (Volume Configuration Management Data Base) used by Dell EMC's SAN management software (Volume Logix, ESN Manager, SAN Manager and so on).

3.4 Managing System Disks

- For managing system disks, register them as keep disks with the root class. [EFI]
- Volumes in local classes and shared classes cannot be used as:
/ (root), /usr, /var, /opt/, /boot, /boot/efi, swap areas, dump saving area (such as /var/crash)
To manage disks for these uses with GDS, register the disks with the root class.
- With respect to dump devices and disks used as kdump dump saving area, GDS does not support them.



See

.....
For details on dump devices and dump saving area, see the dump function manual.
.....

3.5 System Disk Mirroring Preconditions

- Conditions for OS
 - RHEL8(Intel64) (UEFI boot environment)
 - RHEL7(Intel64) (UEFI boot environment)
- Conditions for server
 - PRIMEQUEST 3000 Series
 - PRIMEQUEST 2000 Series
 - PRIMERGY (VMware environment)
- Conditions for system disk configuration and settings
 - The disk label is the GPT type.
 - The number of slices is 14 or less.
 - The disk size is smaller than 1 TB.
 - Sufficient free space.
For details on necessary free space, refer to "2.3.9 Disk Size."
 - The areas of /, /usr, /var, /boot, and /boot/efi are located in the same disk.
When mirroring the swap area, the swap area must be located in the same disk as well.

- The device of the system volumes (/ , /var, /usr, /boot, /boot/efi, or swap area) is defined in the /etc/fstab file by the following formats:
device name, LABEL, or UUID
- Without LVM (Logical Volume Manager) configuration

Information

System disk mirroring of the following environments is also available.

- SAN boot environment
- KVM host OS
- VMware guest

Note

- When mirroring the system disk in SAN boot environment, contact field engineers to check if the environment is supported.
- In an iSCSI boot environment, management and mirroring of system disk cannot be done.
- When mirroring the system disk, do not write GRUB_DISABLE_LINUX_UUID=true or GRUB_DISABLE_UUID=true in the /etc/default/grub file.

3.6 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/fstab and so on, it is required to change it to the device special file name for the following volume.

Block device special file:

```
/dev/sfdsk/class_name/dsk/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.

Note

Keep disks can be registered only with the root class.

3.7 Keep Disk [EFI]

By specifying a disk as "keep" type and registering it with root class, you can retain the disk data while configuring mirroring. The disk to be registered with root class as a keep disk must conform to the following conditions:

- The disk label is the GPT type.
- The number of slices is 14 or less.

- The system disk has sufficient free space.

For details on free space area, see "2.3.9 Disk Size."

Additionally, when registering physical disks other than system disks (with running /, /usr, /var, /boot, /boot/efi, 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.


3.8 Hot Spare

3.8.1 Number of Disks within Class

If using the hot spare function with a local class or shared class, the number of disks including the spare disks within the class is to be set to four or more. If the number of disks within the class is three, the spare disk will not be effective.

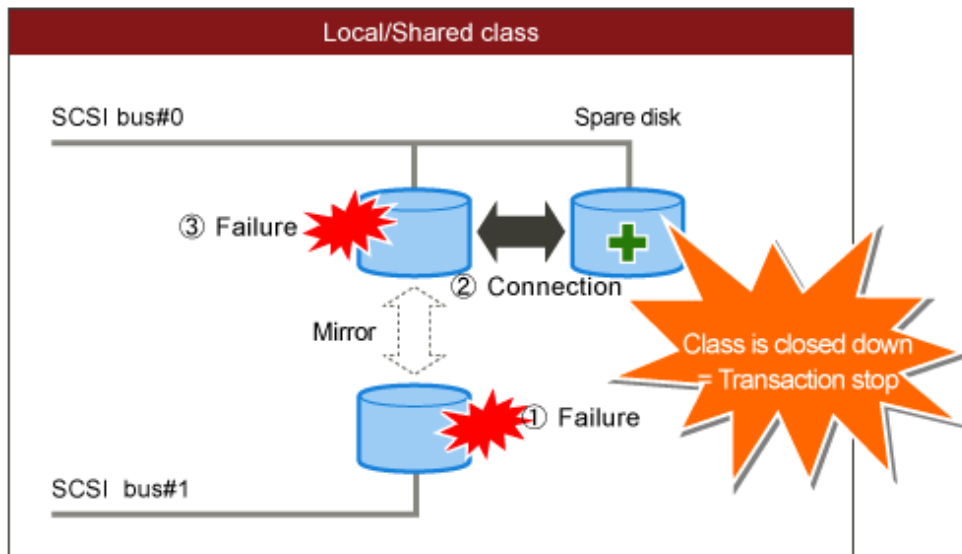
For the configuration in example 1, with only the SCSI Bus#0 malfunctioning, the class gets closed down and the transactions stop.

Figure 3.1 Example 1 of an unsuitable configuration




For the configuration in Example 2, after one of the disks has malfunctioned and the spare disk has been connected to the mirror group, the class gets closed down and transactions stop when one more disk malfunctions.

Figure 3.2 Example 2 of an unsuitable configuration



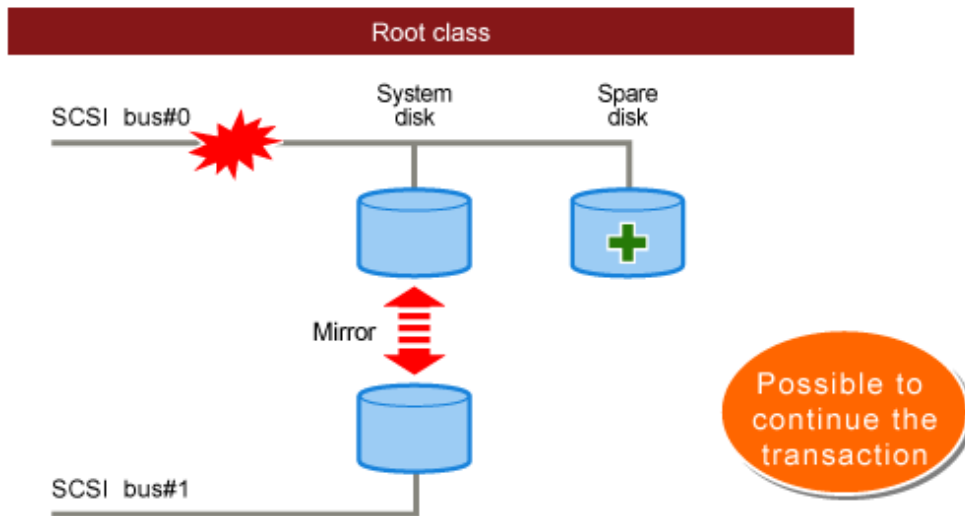
By increasing the number of disks within the class, it is possible to avoid the transaction stop caused by a single hardware malfunction.

Figure 3.3 Example 1 of a suitable configuration



If using the root class, insofar as there is a normal disk, the class will not close down, so setting the number of disks within the root class including the spare disk to three is not a problem.

Figure 3.4 Example 2 of a suitable configuration




 Note

The hot spare function cannot be used in the root class in this version.

3.8.2 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. Therefore, the capacity of the spare disk must be larger than the available size of the lower group.

Figure 3.5 Hot Spare for Hierarchized Mirror Groups



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.

3.8.3 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.

3.8.4 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.

3.8.5 Hot Spare for the Root Class

The hot spare function cannot be used in the root class.

3.8.6 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.

3.8.7 Shadow Classes



Shadow objects are not supported in this version.

Spare disks cannot be registered with shadow classes.

3.8.8 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.

3.8.9 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.

3.8.10 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.



For details, see "[B.1.12 sdxparam - Configuration Parameter Operations](#)."

3.8.11 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 0.5 (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`.

3.8.12 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 for 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.

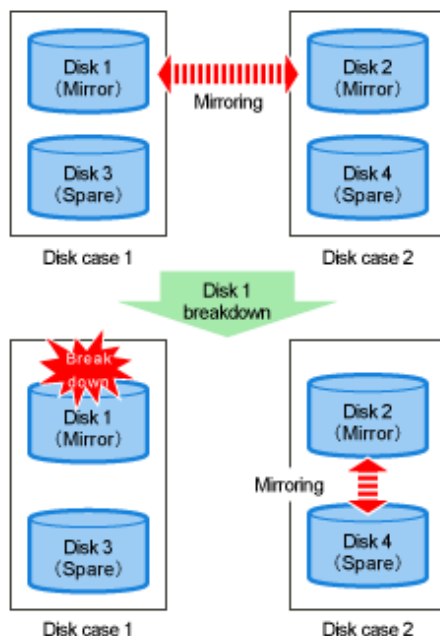
If an I/O error occurs in a disk irrelevant to a disk array unit (such as an internal disk), this method selects a spare disk that is connected to a controller different from that the failed disk is connected to first.

If no applicable spare disk connected to a group, this method selects 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, there are possibilities that the I/O cable of the disk case has an error, the entire disk case is down, or the controller has 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 3.6 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 3.7 Hot Spare in Internal Mode



Note

When an I/O error occurs in a disk of the disk array unit, in order for the Hot Spare functions to work as specified by the Hot Spare mode, the configuration must fulfill the following three conditions:

- The multiplicity of mirroring is two.
- The two disks to be mirrored belong to different disk array cases.
- The spare disks, which are registered in the classes to which the mirrored disks belong, are in the same disk array cases to which the mirrored disks belong.

3.8.13 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 3.7 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.

See

For disk type changing methods, see "Changing the Class Attributes" in "8.2.1.2 Class Configuration" when using the GDS Management View, or "B.1.7 sdxattr - Change 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 "8.2.1.3 Group Configuration" when using the GDS Management View, or the description about the -C option in "B.1.2 sdxdisk - Disk Operations" when using the command.

3.9 Mirroring among Servers

3.9.1 Network Configuration

The mirroring function among servers utilizes iSCSI to read/write the data of a local disk on another node via the network.

- The network which is used in the mirroring among servers need to be a private network. This network is usable also as other networks of the cluster system, such as the cluster interconnect, the administrative LAN, and the public LAN.

However, if you are concerned that the I/O processing to the network mirror volumes may affect on other communications or vice versa, you must have the network dedicated to the mirroring among servers.

- The IPv4 address or the IPv6 address can be used.
- For IPv6 address, a global unicast address and a unique local unicast address can be used. A link-local unicast address cannot be used.
- When configuring the network used for the mirroring among servers to be redundant:
 - Use Virtual NIC mode or NIC switching mode of Global Link Services (hereinafter GLS).
 - To use NIC switching mode of GLS, specify the physical IP address for the IP address used for the mirroring among servers.
 - To use NIC switching mode of GLS (Physical IP takeover), do not register the takeover IP address to the cluster application.
 - Do not change the monitoring time of GLS.

Redundancy mode	Shared with interconnect		Not shared with interconnect	
	IPv4	IPv6	IPv4	IPv6
No redundancy	Y	Y	Y	Y
Virtual NIC mode of GLS	N	N	Y	Y
NIC switching mode of GLS	N	N	Y	N

GLS: Global Link Services

- When using IP aliasing, use IP address that is always available on the node.
- To improve the system availability, avoid sharing the network used for the mirroring among servers with the cluster interconnect, the administrative LAN, and the public LAN. Configuring the network to be redundant is also recommended.
- Do not allow any operations that disable the communication done by using the IP address set by the process described in "4.8 Disk Setting for Performing Mirroring among Servers."

(Operation examples)

- Stopping NIC that is not redundant or the virtual interface of GLS by using the ifconfig command or the ip command of OS.
- Stopping the virtual interface by using any GLS commands.

In the configuration of mirroring among servers, besides the disk of any other nodes, GDS also accesses the disk of own node by using iSCSI. If the operations described above are allowed on the operating node, the disk of both nodes cannot be accessed from the operating node. In this case, the disk of the operating node cannot be accessed by the standby node. This will cause the following errors.

- The cluster application is switched over.

- First one of the two netmirror slices that consist the netmirror volume where the I/O error is detected is detached from the netmirror volume, and the slice becomes INVALID.
- If the slice on the node where the cluster application is switched becomes INVALID, the cluster application cannot be started on the switched node, and the service will stop.

Note

Sharing the network with the cluster interconnect or the administrative LAN

- When the I/O speed (data traffic per second) from the application to the netmirror volume exceeds 50% of the bandwidth of the cluster interconnect, I/O to the netmirror volume may affect the heartbeat processing of the cluster system. In this case, do not share the network for the mirroring among servers with the network for the cluster interconnect.
- Do not share the network for the mirroring among servers with the network for the administrative LAN as well when the I/O speed to the netmirror volume exceeds 50% of the bandwidth of the administrative LAN.
- When sharing the network for the cluster interconnect with the network for the mirroring among servers, the CIP IP address cannot be used for the mirroring among servers. Assign another IP address to the network device specified in the cluster interconnect.
- When sharing the network for the cluster interconnect with the network for the mirroring among servers, the network for the mirroring among servers cannot be configured to be redundant. In this case, the risk of operation being stopped will be increased as follows:
 - Due to a network error of the mirroring among servers, the disk on the standby node cannot be accessed from the operating node. As a result, the disk on the standby node is detached from the mirroring. If the cluster application is switched at this time, the operation stops. For details, see "[7.16.11 Switchover of the Cluster Application when the Network is Disconnected.](#)"

Note

Sharing the network with the public LAN

When sharing the network between the public LAN and the mirroring among servers, if an error occurs in the network, an error occurs both in the communication of service application and in writing to the netmirror volume. In this case, the service will stop due to the following behaviors.

- The slice on the disk of the standby node is detached because the disk on the standby node cannot be accessed from the operating node.
- Triggered by the communication error of the service application, failover of the cluster application occurs.
- On the new operating node, the slice on the disk of the own node is detached and cannot be used. The disk of the remote node cannot be accessed due to a network error. In this case, the service will stop.

To continue the service during the single failure of the network, configure the network to be redundant.

If the total data traffic per second may exceed the network bandwidth by sharing the network, configure the unshared network.

3.9.2 Disk

The following devices can be used for the mirroring among servers.

- Internal disks (hard disks, SSD)
 - RAID disks by RAID card also can be used.
- Physical disks (RMDK), virtual disks (VMDK) in a VMware environment
- virtio-SCSI devices and virtio block devices of the KVM guest
- Block devices in a cloud environment (except an FJcloud-Baremetal environment)
- Block storage provided by the OpenStack Block Storage service (cinder) in an RHOSP environment

- Partitions in the above disks

For example, when the server has only one internal disk, use it for installing the OS, and make the partition created in a free space as a target of the mirroring among servers.

The following devices can not be used for the mirroring among servers.

- External disk units (such as ETERNUS) in a physical environment
- The disk already registered to the GDS class
- The device whose compatible device name is not `sdX` or `vdX`, such as GDS volumes, GDS slices, and LVM logical volumes

3.9.3 Object Configuration

Set a disk of each node to be accessed as the iSCSI device from another node. As well as the shared disk device, register the iSCSI device as a resource, and then register it in the shared class. Connect the iSCSI device with the netmirror group.

The object configuration must satisfy the following conditions.

- The number of iSCSI devices for mirroring among servers : 16 or less
- The shared class where the netmirror group is created
 - The number of classes: 2 or less
 - The number of nodes in the scope: 2-node
 - The number of disks in the class: 8 or less
 - Any groups except netmirror type, a single disk, and a spare disk cannot be created
 - The shared class must be registered to the cluster application as the Gds resource
- The number of disks in the netmirror group: 2
- The JRM attribute of the netmirror volume: on



See

For how to set a disk of each node to be accessed as the iSCSI device from another node, see "[4.8 Disk Setting for Performing Mirroring among Servers.](#)"



Note

- When the cluster application is switched, an additional few seconds are required for one netmirror volume. For this reason, lower number of netmirror volumes is recommended.
- When the disk on another node cannot be accessed due to a network error, and when the resynchronization copying is performed after the network is restored, I/O to the volume is suspended for a few seconds for one volume. Lower number of volumes is recommended. The time to suspend I/O is about 4 to 5 seconds though it changes depending on the system.
- If both nodes are restarted before the shared class where the netmirror volume is created is registered to the cluster application as the Gds resource, the start lock is set in the netmirror volume to protect the data. At this time, the netmirror volume cannot be started. In this case, register the shared class to the cluster application to start the cluster application.

3.9.4 Configuration and Settings of the Cluster System

- Configuration of the Cluster System

The mirroring among servers is available for the following conditions.

- Operation mode
 - Standby operation (cold standby) ,1:1Active standby or mutual standby

- Scalable operation (Only when using Symfoware Server (Native) hot standby in a cloud environment)
- The number of cluster applications: 2 or less
- The number of nodes: 2 nodes
- Setting CF remote service
 - Enable the function to execute the command on other nodes (cfsh).
 - For details, see "PRIMECLUSTER Installation and Administration Guide" and "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide."
- Forced stop method in a VMware environment
 - Set the forced stop method using the VMware vCenter Server functional cooperation.
 - The forced stop method using the I/O fencing function cannot be used.

Note

If the configuration satisfies any of the following conditions, the synchronization copying and volume startup after restoration from the LEFTCLUSTER state may be suppressed due to data protection.

- The number of cluster applications is 2.
- The survival priority of the node is not set according to the survival scenarios for specific application survival.

See

- For details, see "[7.16.4 Restoration from the LEFTCLUSTER State.](#)"
- For details on the survival priority and the survival scenarios, see "PRIMECLUSTER Installation and Administration Guide."

3.9.5 Data Amount and Data Update Amount

The mirroring among servers is suitable for the operation if the amount of data and the amount of data update is small for the following reasons:

- Why the mirroring among servers is suitable for the operation if the amount of data is small
 - When a disk is replaced, the data of the disk is copied to the replaced disk so that the disk is mirrored. The operation can be continued during copying, however, availability of the operation becomes low. When the data amount is large, copying processing takes time and it leads to the longer period of the low operational availability. For the required time for copying processing, see "[3.9.6 Required Time for Copying Processing.](#)"
- Why the mirroring among servers is suitable for the operation if the amount of data update is small
 - In the mirroring among servers, disk data is written via the network. In this case, the write performance is lowered compared to the environment without mirroring.
 - When an I/O error occurs due to the network error, the data difference occurs among the target disks of the mirroring among servers. If the network is restored, the resynchronization copying is automatically performed, and the data in the area where the data difference occurs is copied. If the amount of data update is small, the equivalence can be restored for a short time.

Note

The resynchronization copying which is executed after the following events copies the entire slice.

- When an I/O error occurs in the memory area of difference information inside the private slice
- When the resynchronization copying in process is canceled due to shutting the node, switching the cluster application, the `sdxcopy` command, etc.

If the amount of data and the amount of data update of the operation is large, use the shared disk.

3.9.6 Required Time for Copying Processing

Estimated time of copying processing of the mirroring among servers is 1 minute per 1 GB.

Size of the netmirror volume	Estimated time of copying processing
100 GB	1hour 40 minutes
300 GB	5 hours
1 TB	17 hours
4 TB	2.8 days
6 TB	4.3 days

Actual time of copying processing changes according to the system due to the performance of server, network, or disk.

3.9.7 Note

The following functions cannot be used for disks, groups, and volumes in which the mirroring among servers performs:

- Hot spare
- Create a hierarchy of groups (in combination with the striping or the concatenation)
- Creating snapshot by slice detachment
- GDS Snapshot (proxy volumes and shadow volumes)
- Volume expansion

3.10 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 "[B.1.7 sdxattr - Change Objects Attributes.](#)"
- For the method for setting the "current access modes" of logical volumes, see "[B.1.4 sdxvolume - Volume Operations.](#)"
- For the method for setting the "current access mode" for shadow volumes, see "[B.2.4 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 "[B.1.6 sdxinfo - Display Object Configuration and Status Information.](#)"

3.11 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.

3.12 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 "[7.4 Backing Up and Restoring System Disk.](#)"
.....

- 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.

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 Shared File System cannot be expanded.

- **Capacity Expansion of Disk Registered in the Class**

Do not expand the capacity of disk (LUN) registered in the class. If do so, GDS does not work properly.

3.13 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.

3.14 Creating 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, it is impossible to perform snapshot creation by slice detachment. In addition, slices cannot be detached from shadow volumes and from mirror volumes in the root class.

3.15 Online Backup and Restoration Using Proxy Volume

This subsection describes the modes 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 following functions, backup and restore processing that do not impose the load on the server or SAN become possible.

- Advanced Copy function of ETERNUS Disk storage system
EC (including REC) or OPC (including QuickOPC)
- Copy function of Dell EMC storage units (not supported in this version)
Dell EMC TimeFinder or Dell EMC SRDF

The following table shows backup and restore methods and their features.

Operation	Disk Unit's Copy Function	Master Configuration (*2)	Proxy Configuration (*2)	Feature
Backup(Sync)	-	single mirror mirror + stripe mirror + concat	single mirror mirror + stripe mirror + 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.

Operation	Disk Unit's Copy Function	Master Configuration (*2)	Proxy Configuration (*2)	Feature
	EC REC TimeFinder (*1) SRDF (*1)	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 QuickOPC	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 QuickOPC	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 mirror + stripe mirror + concat	single mirror mirror + stripe mirror + 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 QuickOPC	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) Not supported in this version

(*2) In the table above, configurations indicate objects as follows.

Configuration	Description
single	<p>This configuration means any one of the following objects.</p> <ul style="list-style-type: none"> - A single volume created in a single disk - A mirror group not hierarchized, with the multiplicity of one - A volume created in a mirror group not hierarchized, with the multiplicity of one
mirror	<p>This configuration means either the following objects.</p>

Configuration	Description
	<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
mirror + stripe	<p>This configuration means either the following objects.</p> <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
mirror + concat	<p>This configuration means either the following objects.</p> <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, SRDF), see "[1.5.2 Snapshot Function with No Load to Server/SAN](#)," "[A.8 Using Advanced Copy Function in Proxy Configuration](#)," "[3.18 Instant Snapshot by OPC](#)" and "[3.19 Using Dell EMC TimeFinder or Dell EMC SRDF in Proxy Configuration](#)."

3.16 Difference between Mirror Slice and 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

.....
 See "[Figure 1.32 Difference between a Mirrored Slice and Synchronized Proxy Volume](#)" in "[1.5.1 Snapshot by Synchronization](#)."

3.17 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, if the volume (master) has been used as a file system such as GFS and ext4, before and after creating snapshots, unmount the system and remount it in order to ensure the snapshot data integrity.

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 "[3.15 Online Backup and Restoration Using Proxy Volume](#)."

3.18 Instant Snapshot by OPC

The following functions are based on OPC functions (including QuickOPC function) of ETERNUS Disk storage system.

- 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 only available for using a disk array unit with the OPC function in ETERNUS Disk storage system.

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.8 Using Advanced Copy Function in 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 "[D.1.3 Volume Status Abnormality.](#)"

3.19 Using Dell EMC TimeFinder or Dell EMC SRDF in Proxy Configuration



Note

The cooperation with Dell EMC TimeFinder and Dell EMC SRDF are not supported in this version. When creating snapshots using the Dell EMC storage units that support TimeFinder or SRDF, configure the disk array units as TimeFinder and SRDF are disabled in advance.

In a proxy configuration, by using Dell EMC's TimeFinder and SRDF, the copy functions of Dell EMC 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 "D.1.4 Class Status Abnormality."
- The BCV and target (R2) devices connected to the proxy group may not be described in the Excluded Device List. However, native devices configuring such devices must be described in the list. For details on the Excluded Device List, see "3.3 Dell EMC Storage Units."
- 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. [EFI]
- 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 canceled, the TimeFinder function or SRDF function is not used. BCV pairs and SRDF pairs are canceled 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.

3.20 Backing Up and Restoring through 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 "[7.12.1 Backing Up and Restoring Logical Volume with No Replication.](#)"

2. Backing up and restoring through snapshot by slice detachment



See

For details, see "[7.12.2 Backing Up and Restoring through Snapshot by Slice Detachment.](#)"

3. Backing up and restoring using snapshots from a proxy volume



See

For details, see "[7.12.3 Backing Up and Restoring Using Snapshots from Proxy Volume.](#)"

4. Backing up and restoring by the disk unit's copy function



See

For details, see "[7.12.4 Backing Up and Restoring by 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
2	Available	Slices detached temporarily	-	One of (*3): mirror (concat + mirror) (stripe + mirror)	GDS
3	Available	Proxy volumes	-	One of (*4): single mirror concat + mirror stripe + mirror	GDS GDS Snapshot

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	
			Dell EMC TimeFinder or Dell EMC SRDF (*6)	One of (*6): mirror	
4	Available	Non-SDX disks as the disk unit's copy function destinations	Dell EMC TimeFinder or Dell EMC SRDF (*7)	Any	GDS

GDS: Global Disk Services

GDS Snapshot: Global Disk Services Snapshot

(*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.

Type	Description
	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 ["3.14 Creating Snapshot by Slice Detachment."](#)

(*4):

See ["2.3.12 Proxy Configuration Preconditions."](#)

(*5):

See ["2.3.12 Proxy Configuration Preconditions"](#) and ["A.8 Using Advanced Copy Function in Proxy Configuration."](#)

(*6):

See ["2.3.12 Proxy Configuration Preconditions"](#) and ["3.19 Using Dell EMC TimeFinder or Dell EMC SRDF in Proxy Configuration."](#)

(*7):

There are two operation modes that use Dell EMC TimeFinder or Dell 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.

3.21 Virtual Environment

This section describes GDS functions which are available in the virtual environments (KVM, VMware, cloud, and RHOSP) and the notes as follows:



See

To configure and operate GDS in the virtual environments, you must have the knowledge of virtual system design, installation, and operation.

See the manual of the virtual machine function in use beforehand.

3.21.1 KVM Environment and VMware Environment

GDS functions which are available in the virtual environments (KVM and VMware) are as follows:

Purpose	Operating environment of GDS	Class	Description
Mirroring a system disk on the host	Host (KVM)	Root class	Available in UEFI boot environment. The settings are the same with the conventional way where the virtual machine function is not used.
Mirroring a system disk on the guest	Host (KVM)	Local class Shared class	Create a volume of a local class or a shared class on the host and install the guest OS there.

Purpose	Operating environment of GDS	Class	Description
	Guest (VMware)	Root class	Available in UEFI boot environment. Mirroring of the following system disks, which were connected by the virtual SCSI controller, is available: <ul style="list-style-type: none"> - Physical disk (RDM) on a VMware guest - Virtual disk (VMDK) on a VMware guest
Mirroring a virtual disk on the guest	Guests (KVM and VMware)	Local class Shared class	Supporting the mirroring in order to improve the availability. Mirroring of the following devices are available: <ul style="list-style-type: none"> - virtio block devices on a KVM guest - virtio-SCSI device on a KVM guest - Physical disk (RDM) on a VMware guest connected by the virtual SCSI controller - Virtual disk (VMDK) on a VMware guest connected by the virtual SCSI controller The settings are the same as the conventional way where the virtual machine function is not used.
Management of the shared disk for the cluster system Mirroring among servers of the cluster system		Shared class	GDS manages the shared disk for the cluster system built between guests and enables the mirroring among servers. The following devices can be managed: <ul style="list-style-type: none"> - virtio block devices on a KVM guest - virtio-SCSI device on a KVM guest - Physical disk (RDM) on a VMware guest connected by the virtual SCSI controller - Virtual disk (VMDK) on a VMware guest connected by the virtual SCSI controller The settings are the same as the conventional way where the virtual machine function is not used.

GDS: Global Disk Services

GDS Snapshot functions which are available in the virtual environments (KVM and VMware) are as follows:

Purpose	Operating environment of GDS Snapshot	Class	Description
Snapshot of a system disk on the host Configuration of an Alternative Boot Environment on the host	Host (KVM)	Root class	Available in UEFI boot environment. The settings are the same with the conventional way where the virtual machine function is not used.
Snapshot of data volume on the host	Host (KVM)	Local class Shared class	The settings are the same with the conventional way where the virtual machine function is not used.
Snapshot of a system disk on the guest	Host (KVM)	Local class Shared class	Create a volume of a local class or a shared class on the host and install the guest OS there. Create a snapshot on the host.

Purpose	Operating environment of GDS Snapshot	Class	Description
			To ensure the consistency of snapshot data, you need to stop the guest OS during snapshot creation.
	Guest (VMware)	Root class	Available in UEFI boot environment. Available for the system disk of the following formats connected by the virtual SCSI controller: <ul style="list-style-type: none"> - Physical disk (RDM) on a VMware guest - Virtual disk (VMDK) on a VMware guest
Configuration of an Alternative Boot Environment on the guest	Guest (VMware)	Root class	Available in the following environment: <ul style="list-style-type: none"> - RHEL6 (Intel64) in a UEFI boot environment - RHEL7(Intel64) in a UEFI boot environment Available for the system disk of the following formats connected by the virtual SCSI controller: <ul style="list-style-type: none"> - Physical disk (RDM) on a VMware guest - Virtual disk (VMDK) on a VMware guest
Snapshot of local volume on the guest	Host (KVM)	Local class	Create a local volume on the host and register it as a virtio block device on the guest OS. Create a snapshot on the host.
	Guests (KVM and VMware)	Local class	<ul style="list-style-type: none"> - The settings are the same with the conventional way where the virtual machine function is not used. - Snapshot cannot be created by using the Advanced Copy function of ETERNUS in the following environments: <ul style="list-style-type: none"> - KVM guest - Virtual disk (VMDK) on a VMware guest
Snapshot of shared volume on the cluster system.	Guests (KVM and VMware)	Shared class	<ul style="list-style-type: none"> - The settings are the same with the conventional way where the virtual machine function is not used. - Snapshot cannot be created by using the Advanced Copy function of ETERNUS in the following environments: <ul style="list-style-type: none"> - KVM guest - Virtual disk (VMDK) on a VMware guest

GDS Snapshot: Global Disk Services Snapshot



Note

VMware vSAN

- Virtual disk (VMDK) of VMware vSAN can be registered to the class. However, it cannot be used for the mirroring among servers.
- For the virtual disk redundancy of vSAN, perform either the mirroring function of GDS or redundant function of vSAN, and decide which function to use according to the system requirements.
For example, when constructing a system that does not make the administrator of the virtual machine aware of the redundant disk, the server administrator can construct such system by using the vSAN function to configure the redundant disk.
On the other hand, if you want to make the disk redundant to be used in the virtual machine in a cloud environment (except an FJcloud-

Baremetal environment) where the disks are not made redundant by vSAN function, you can use the mirroring function of GDS on the virtual machine.

Note

When using GDS on both host and guest of KVM

- For the disk used only on KVM guest, configure the following on the host.
 - Register the disk to the Excluded Device List of GDS.
For the Excluded Device List, see "[6.1 Creating Excluded Device List](#)."
 - Do not register the disk to the PRIMECLUSTER resource database.
For the resource registration, see "[6.3 Shared Disk Resource Registration](#)."
- It is recommended to create different class names for the host and guest.

Information

How to copy the volumes created in a KVM guest by using the Advanced Copy function of ETERNUS

When using GDS Snapshot on the KVM guest, the Advanced Copy function of ETERNUS cannot be used. To use the Advanced Copy function of ETERNUS, use ACM. For details, see the manual of ACM.

3.21.2 Cloud Environment

In the cloud environment (except an FJcloud-Baremetal environment), the following function can be used in a cluster system composed of two virtual servers in the same availability zone.

Function	Description
Mirroring among servers	Data between cluster nodes can be taken over with the mirroring among servers function using the block storage attached to a virtual server as additional storage.

In the FJcloud-Baremetal environment, the following functions can be used on the Bare Metal server.

Function	Description
Management of shared disk	The shared disk device connected to multiple Bare Metal servers with iSCSI can be registered to the shared class.
Management of local disk	The disk device connected to the Bare Metal server with iSCSI can be registered to the local class.

See

For details about the usage in the cloud environment, see "[PRIMECLUSTER Installation and Administration Guide Cloud Services](#)."

3.21.3 RHOSP Environment

In the RHOSP environment, the following functions can be used on the virtual machine instance.

Function	Description
Management of shared disk	The shared disk device connected to multiple virtual machine instances with iSCSI can be registered to the shared class. GDS Snapshot can be also used.
Mirroring among servers	Data between cluster nodes can be taken over with the mirroring among servers function using the block storage provided by the OpenStack Block Storage service (cinder).

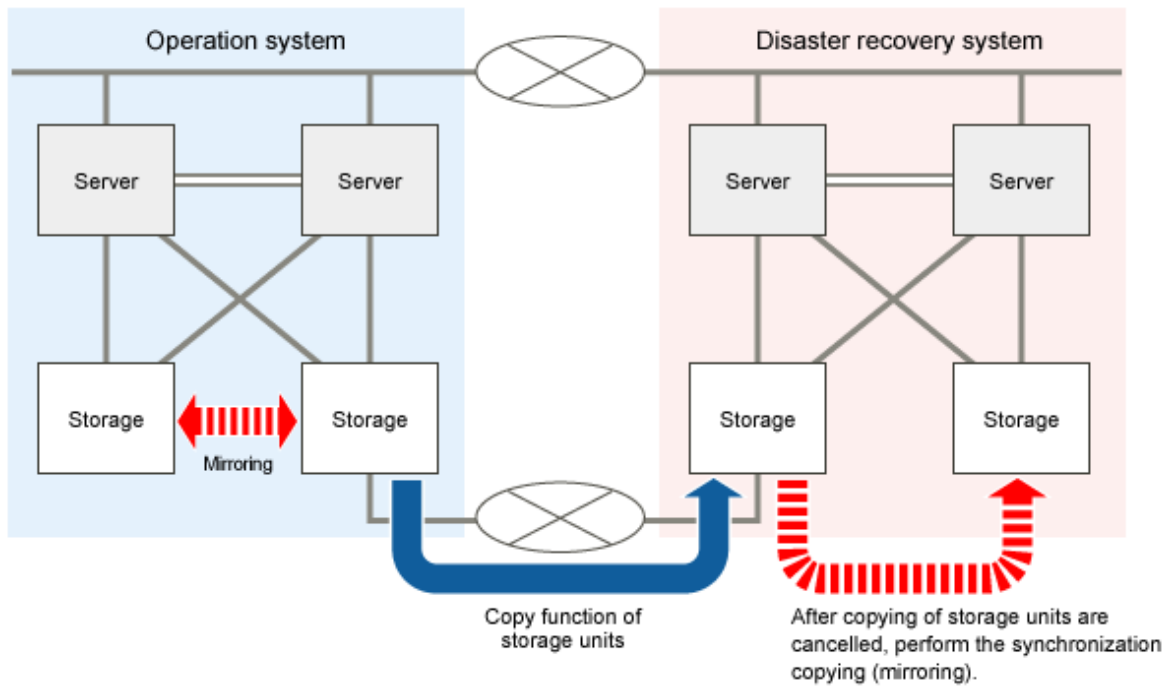
Function	Description
Management of local disk	The disk device connected to the virtual machine instance with iSCSI can be registered to the local class. GDS Snapshot can be also used.

GDS Snapshot: Global Disk Services Snapshot

3.22 Disaster Recovery

In preparation for large-scale disasters such as earthquakes, this section describes the outline method to configure a replication (hereafter a disaster recovery system) of the operation system where GDS is used in a remote location and so on to run the service.

Figure 3.8 Disaster recovery system



For copying data to a disaster recovery system from the operation system, use the copy function of storage units. Either of the following copy functions are available:

- Copy function per slice
Example: REC of ETERNUS Disk storage system
- Copy function per LUN
Example: ACM CCM and Dell EMC SRDF

 See

In construction and operation of a disaster recovery system, use the `sdxconfig` command. For the method of using the `sdxconfig` command, see the following:

- "[B.1.13 sdxconfig - Object Configuration Operations](#)"
- "[7.13 Backing Up and Restoring Object Configurations](#)"
- "[7.11 Backing Up and Restoring through Disk Unit's Copy Functions](#)"
- "[A.5 Backing Up and Restoring Object Configuration \(sdxconfig\)](#)"

3.22.1 Building Disaster Recovery System

This sub-section describes the procedure to build a replication of the operation system where GDS is used in a remote location and so on as a disaster recovery system.

The outline of building procedure is as follows:

1. Back up object configuration information of local and shared classes on the operation system. (sdxcfg Backup command)
2. Convert the configuration files backed up in step 1 for a disaster recovery system as below to create configuration files for a disaster recovery system. (sdxcfg Convert command)
 - Remove disks other than the copy destination of the storage copy process from each mirror group. (sdxcfg Convert -e remove command).
 - Change the physical disk names to physical disk names of a disaster recovery system. (sdxcfg Convert -e replace command)
3. Set the cluster system (CF, SF, shared disk resource) on the disaster recovery system.
4. When using the copy function per slice of storage (such as REC of ETERNUS Disk storage system), use the configuration files created in step 2 for a disaster recovery system to restore object configuration information on a disaster recovery system. (sdxcfg Restore command)
After executing the sdxcfg Restore command, reboot the server where the command was executed.
5. Copy data of the local class and shared class to one case on the disaster recovery system from one case on the operation system by using the copy function of storage.
6. When using the copy function per LUN of storage (such as ACM CCM), use the configuration files created in step 2 for a disaster recovery system to restore object configuration information. (sdxcfg Restore command)
In the sdxcfg Restore command, -e chkps option is set.
After executing the sdxcfg Restore command, reboot the server where the command was executed.
7. Set the shared class (change type from local to shared) on the disaster recovery system. (sdxattr -C command)
8. Set the cluster application on the disaster recovery system.

3.22.2 Data Synchronization of Operation System and Disaster Recovery System

This sub-section describes the procedure of synchronizing disk data of the local class and the shared class of operation system to the disaster recovery system.

The outline of procedure is as follows:

1. Stop the application and RMS on the disaster recovery system.
If there is a volume started up in the disaster recovery system, stop the volume. (sdxvolume -F command)
2. When using the copy function per LUN of storage (such as ACM CCM), delete the local class and shared class in the disaster recovery system. (sdxcfg Remove command)
Before deleting the shared class, change the shared class to the local class by the sdxattr -C command.
3. Start the copy function of data to one case on the disaster recovery system from one case on the operation system by using the copy function of storage.

3.22.3 Switching to Disaster Recovery System

This sub-section describes the procedure when switching the service from the operation system to a disaster recovery system in the event of disasters.

The outline of switching procedure is as follows:

1. Cancel copying of storage.
2. When using the copy function per LUN (such as ACM CCM) of storage, restore object configuration information on a disaster recovery system by using the configuration files of the disaster recovery system created in step 2 in "[3.22.1 Building Disaster](#)

[Recovery System.](#)" (sdxconfig Restore command)

In the sdxconfig Restore command, -e chkps option is set.

After executing the sdxconfig Restore command, reboot the server where the sdxconfig Restore command was executed.

After the server is restarted, set the shared class (change type from local to shared) using (sdxattr -C command) and set using (hvgdsetup -a command) to use the shared class in the cluster application.

3. When mirroring storage on a disaster recovery system, add a disk to a mirror group on a disaster recovery system. (sdxdisk -M and sdxdisk -C commands)

3.22.4 Restoration to Operation System from Disaster Recovery System

This sub-section describes the procedure when the operation system is restored by using data on a disaster recovery system after restoring hardware of the operation system.

The outline of restoration procedure is as follows:

1. Cancel copying of storage.
2. Restore the data from the system backup on the operation system.
3. When GDS objects of local class or shared class exist on the operation system, delete object configuration information. (sdxconfig Remove command)
Before deleting the shared class, change the shared class to the local class by the sdxattr -C command.
If the object configuration information cannot be deleted due to storage device replacement etc, the configuration information must be deleted forcefully. For the forceful deletion method, contact field engineers.
4. Convert the configuration files backed up in step 1 in "[3.22.1 Building Disaster Recovery System.](#)" Remove disks other than the copy destination of the storage copy process from each mirror group. (sdxconfig Convert -e remove command)
5. When using the copy function per slice of storage (such as REC of ETERNUS Disk storage system), restore object configuration information on the operation system. (sdxconfig Restore command)
After executing the sdxconfig Restore command, reboot the server where the sdxconfig Restore command was executed.
After the server is restarted, set the shared class (change type from local to shared) using (sdxattr -C command) and set using (hvgdsetup -a command) to use the shared class in the cluster application.
6. Copy data of the local class and shared class to one case on the operation system from one case on a disaster recovery system by using the copy function of storage.
7. When copying of storage becomes synchronized, cancel copying of the storage.
8. When using the copy function per LUN (such as ACM CCM) of storage, restore object configuration information on the operation system. (sdxconfig Restore command)
In the sdxconfig Restore command, -e chkps option is set.
After executing the sdxconfig Restore command, reboot the server where the sdxconfig Restore command was executed.
After the server is restarted, set the shared class (change type from local to shared) using (sdxattr -C command) and set using (hvgdsetup -a command) to use the shared class in the cluster application.
9. When mirroring storage on the operation system, add a disk to a mirror group on the operation system. (sdxdisk -M and sdxdisk -C commands)

3.23 I/O Response Time Assurance [IOmonitor]

This section explains the general notes on designing the I/O response time assurance function provided by GDS I/O Monitor Option as an optional product of GDS.

- In the I/O Response Time Assurance Function, memory is used in each volume. For details, check the amount of memory required described in the installation guide.
- When mirroring the internal disk, mirror the disk with different paths.
When mirroring the internal disk of the same path, the failure of the single disk may also affect other disks and delay the I/O.
- When using the storage device, connect the storage device to the server in order to satisfy the following conditions:
 - Connects using FC card or FCoE card.

- Duplicates the path using ETERNUS Multipath Driver or DM-MP
 - Duplicates the path using different controller module in storage.
 - Access paths for each storage LUN is configured as double paths.
 - When using FC switch, use 2 FC switches.
- When mirroring the disk of storage device, connect the 2 storage devices to the server to mirror (mirroring between units) the disk of the different unit.
 - When using multipath software, either use ETERNUS Multipath Driver or DM-MP. Using both for this configuration is not supported.
 - When using multi-path configuration disk, it is recommended to set the I/O response time value to the default value (77 seconds) or more.
If a value smaller than the default value is set, the I/O may terminate by I/O Response Time Assurance Function before the path switches when an error of a single path of the multi-path configuration occurs.
 - In the system with mirror volume, it is recommended to set the system panic mode to disable for all the class with single volume.
When the system panic mode of I/O Response Time Assurance Function is set to enable for a class with single volume, the redundancy by mirroring has no effect, because the system will panic with just an occurrence of I/O error or I/O delay in only one storage unit.
 - Perform system operation test with the maximum load assumed in the system that uses this function, and depending on the system load check if I/O time out does not occur.
If using the following functions that affect I/O load during operation, use the following functions also to perform system operation test with the maximum load:
 - Synchronization copying of GDS
 - Advance Copy function of ETERNUS

Chapter 4 Installation

This chapter explains how to install GDS, GDS Snapshot, and GDS I/O Monitor Option.

4.1 Notes on Installing OS

- When mirroring the system disk by GDS, note the following points in installing OS.
 - Configure and set up the system disk to satisfy "Conditions for system disk configuration and settings" in "[3.5 System Disk Mirroring Preconditions](#)."
 - If you set up partitions automatically by the operating system installation program before its installation, system disks can be in LVM configuration. In this case, change the partition setup to the configuration which system disks are not in LVM configuration and install the operating system again.
- When re-installing OS, see "[9.5 OS Re-installation](#)."

4.2 Installing GDS

For how to install GDS, see the installation guide of GDS.

4.3 Installing GDS Snapshot

For how to install GDS Snapshot, see the installation guide of GDS Snapshot.

4.4 Installing GDS I/O Monitor Option

For how to install GDS I/O Monitor Option, see the installation guide of GDS I/O Monitor Option.

4.5 Setting Up DM-MP

Before setting up DM-MP, set the following.

- Setting up `user_friendly_name`

Describe the following setting to the DM-MP multipath configuration file `/etc/multipath.conf`.

```
defaults {
    user_friendly_names yes
}
```

In an FJcloud-Baremetal environment, describe the following setting to "device" in the DM-MP multipath configuration file `/etc/multipath.conf`.

```
devices {
    device {
        vendor          "NETAPP"
        product         "LUN.*"
        ...
        user_friendly_names yes
    }
}
```

- Inhibiting the unnecessary message output of `multipathd`

For inhibiting the unnecessary message output of `multipathd` when creating volumes, register the volume devices in the DM-MP black list. For details, see "[\(10\) Messages of multipathd are output when creating volumes](#)." in "[D.1.14 OS Messages](#)."

- Setting up when using the mirroring among servers

Register devices (local disks or iSCSI devices) used for the mirroring among servers in the DM-MP black list.



See

For setting up DM-MP, see the DM-MP manual.

4.6 Configuring KVM Environment

This section describes the necessary settings when using GDS on the KVM guest.

4.6.1 Adding Virtual Disk

If you use GDS on the KVM guest, pay attention to the following points when you add a virtual disk to the KVM guest:

- For a disk to be added to a guest, specify with the by-id name.
- Add the disk as a virtual disk without dividing it. If a partition or a file is added as the virtual disk, GDS cannot manage the virtual disk.
- For the device attribute of a disk to be registered with a class on a guest, set the following value:

libvirt package of the host OS	device attribute
libvirt-0.9.4-23.el6_2.3 or earlier	disk
libvirt-0.9.4-23.el6_2.4 or later	lun



Information

- You can check the version of the libvirt package with the rpm(8) command.

```
# rpm -qi libvirt
```

- When you add a virtual disk to a guest by using the Virtual Machine Manager (virt-manager), the value of the device attribute will be set to "disk."
- For virtio-SCSI device, set "lun" to the value of the device attribute.



Note

In the following cases, you need to change the value of the device attribute from disk to lun.

- When you add a virtual disk by using the Virtual Machine Manager (virt-manager) in the environments where libvirt-0.9.4-23.el6_2.4 or later is applied.
- When you upgrade the libvirt package to libvirt-0.9.4-23.el6_2.4 or later.

Set the device attribute in the guest configuration file (`/etc/libvirt/qemu/guest_name.xml`) on a host OS. When you change the device attribute, stop a guest OS beforehand. For the method of changing the device attribute is as follows:

```
# virsh edit guest_name
```

Example before change

```

:
<disk type='block' device='disk'>
  <driver name='qemu' type='raw' />
  <source dev='/dev/disk/by-id/scsi-1FUJITSU_30000085002B' />
  <target dev='vdb' bus='virtio' />
  <shareable />
  <address type='pci' domain='0x0000' bus='0x00' slot='0x06'
function='0x0' />
</disk>
:

```

Example after change

```

:
<disk type='block' device='lun'>
  <driver name='qemu' type='raw' />
  <source dev='/dev/disk/by-id/scsi-1FUJITSU_30000085002B' />
  <target dev='vdb' bus='virtio' />
  <shareable />
  <address type='pci' domain='0x0000' bus='0x00' slot='0x06'
function='0x0' />
</disk>
:

```

4.6.2 Major Number of the virtio Block Device

When using the virtio block device on the KVM guest, configure the system so that the major number of the virtio block device is the same on all nodes in the cluster system. A device type and a driver type should be the same on all nodes. If a type of the device to be connected or a driver to be used varies depending on each node, the major number may be different among the nodes. For example, if DM-MP is applied for the root partition for only some nodes, the DM-MP driver is loaded in a different order depending on each node. In this case, the major number of the virtio block device may vary.

Take the procedure below to check the major number of the virtio block device. The major number is 253 in the example below.

```

# cat /proc/devices | grep virtblk
253 virtblk

```


If the major number varies on each node, change the setting of the driver so that the major number is the same on all nodes. For how to change the major number of the driver or how to change the installing order, see the manual of the driver, or contact the supplier of the driver.

4.6.3 Setting up Shared Disk

When using a shared disk on the KVM guest, set [Shareable] and then set [none] (cache='none') to [cache model].

The setting procedure is as follows:

For virtio block device

1. Stop the guest OS.
2. Select the stopped guest OS from the Virtual Machine Manager and click the [Open] button in the tool bar.
3. Click  in the tool bar of the displayed window to display the detailed information of hardware.
4. Select a virtual disk [VirtIO Disk] from the hardware list in the left.
5. In the [Virtual disk] window, perform the following settings and click the [Apply] button.
 - Select the [Shareable] check box.

- Select [none] for the [cache model] list.

For virtio-SCSI device

Describe "shareable" and "cache='none'" in the guest configuration file (/etc/libvirt/qemu/guest_name.xml) on the host OS according to the following description while a guest OS is being stopped.

```
# virsh edit guest_name
```


Example;

```
<disk type='block' device='lun'>
  <driver name='qemu' type='raw' cache='none' />
  <source dev='/dev/disk/by-id/scsi-36000b5d0006a0000006a1296001f0000' />
  <target dev='sdh' bus='scsi' />
  <shareable />
  <address type='drive' controller='0' bus='0' target='0' unit='7' />
</disk>
```

4.6.4 Setting up virtio Block Device Used for Mirroring among Servers

When using the virtio block device for the mirroring among servers function, set up the virtio block device so that the by-id file is created on the KVM guest.

Take the following steps to set up the serial number to the virtual disk. The serial number should be a character string of up to 10 characters that does not duplicate in the virtual machine.

1. Stop the guest OS.
2. Select the stopped guest OS from the Virtual Machine Manager and click the [Open] button in the tool bar.
3. Click  in the tool bar of the displayed window to display the detailed information of hardware.
4. Select a virtual disk [VirtIO Disk] from the hardware list in the left.
5. In the [Virtual disk] window, set the serial number to [Serial number] of [Advanced options] and then click the [Apply] button.
6. Start the guest OS and make sure that the by-id file of virtual disk exists on the guest OS.
 - Make sure that the by-id files exist in all virtio block devices used for mirroring among servers.
 - Make sure that the serial number set in step 5 is included in the file name of by-id file.

```
# ls -l /dev/disk/by-id
...
lrwxrwxrwx 1 root root 9 Apr 18 08:44 virtio-disk001 -> ../../vdg
lrwxrwxrwx 1 root root 9 Apr 18 08:43 virtio-disk002 -> ../../vdh
...
                                serial number
```

4.7 Configuring VMware Environment

If you use GDS on a VMware guest, pay attention to the following points when you add a disk to a VMware guest.

- When registering a virtual disk (VMDK or RDM in virtual compatible mode) in a class, it is necessary to set true to the virtual machine configuration parameter: disk.EnableUUID. If true is not set to disk.EnableUUID, stop the virtual machine first, then change it to true. For how to set the virtual machine configuration parameter, see the manual of VMware.



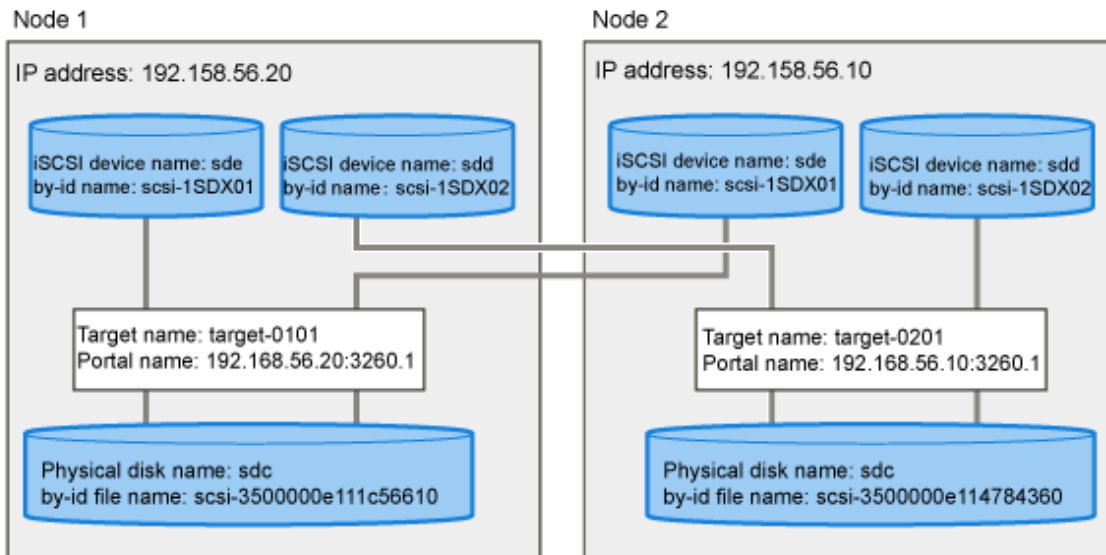
When using a shared disk on a VMware guest, see "PRIMECLUSTER Installation and Administration Guide" for necessary setups and notes.

4.8 Disk Setting for Performing Mirroring among Servers

If you perform the mirroring among servers, you need to set the target disk device to be accessed from each node as the iSCSI device before the setting of "6.3 Shared Disk Resource Registration."

This example is described based on the configuration indicated in the following figure.

Figure 4.1 Disk setting for Performing Mirroring among Servers



By setting this configuration, the target disk device can be accessed from each node just like the shared disk device.

Note

- The iSCSI target created for each node must be only one.
- As shown in the figure above, each disk can be accessed from all nodes in the same iSCSI device name. However, a different iSCSI device name for each node may be assigned to the identical disk. For how to check the iSCSI device name, refer to "4.8.5 Confirming the Correspondence of iSCSI Device Name of Each Node."
- The above figure indicates the example when mirroring the physical disk "sdc." It is also possible to mirror not a physical disk but a partition (such as "sdc2"). When mirroring a partition, confirm the partition configuration (the partition number and the size) to record it by using the parted command. The information is necessary for swapping the disk.
- The physical disk (sdc in the above figure) must be described in the Excluded Device List. For details on the Excluded Device List, see "6.1 Creating Excluded Device List."

4.8.1 Setting Tuning Parameters

Take the following steps to set the tuning parameters.

1. Add the following description to the /etc/opt/FJVSdx/sdx.cf file on both nodes.
 - For the virtual environments (KVM and VMware) or physical environment

```
SDX_NETMIRROR=on
SDX_NETMIRROR_DISK_COUNT=8
SDX_MV_STOP_WAIT_TIME=45
ED_CMD_RETRY_COUNT=60
ED_DRV_RETRY_COUNT=60
```

- For the cloud environment (except an FJcloud-Baremetal environment) or RHOSP environment

```
SDX_NETMIRROR=on
SDX_NETMIRROR_DISK_COUNT=8
SDX_MV_STOP_WAIT_TIME=45
ED_CMD_RETRY_COUNT=100
ED_DRV_RETRY_COUNT=100
```

Note

If [Timeout value of PRIMECLUSTER] below is changed to be larger than the default value, change the value of [Parameters necessary to be changed] according to [Calculation formula of parameter value].

[Timeout value of PRIMECLUSTER]

- Timeout detection time of heartbeat of CF (CLUSTER_TIMEOUT)
- Timeout value of shutdown facility

[Parameters necessary to be changed]

- ED_CMD_RETRY_COUNT
- ED_DRV_RETRY_COUNT

[Calculation formula of parameter value]

$(\langle \text{Increased CLUSTER_TIMEOUT} \rangle + \langle \text{Number of nodes} \rangle \times \langle \text{Increased timeout value of shutdown facility} \rangle) / 3 + 60$

If either one of the timeout values is changed to be smaller than the default value, calculate the increment of this timeout value as 0.

Round the value to an integer.

"Timeout value of shutdown facility" is the total value of timeout values of all the shutdown agents.

2. Restart both nodes.

4.8.2 Checking and Setting Required Packages

Set the following configuration for all nodes in which the mirroring function among servers is used.

1. Check that the targetcli package has already been installed.

```
# rpm -q targetcli
```

2. Set the start and stop order for iSCSI target service.

Create the directory, /etc/systemd/system/fjvsdx.service.d.

```
# mkdir /etc/systemd/system/fjvsdx.service.d
# chmod 755 /etc/systemd/system/fjvsdx.service.d
```

Create the setting file within the created directory.

```
# touch /etc/systemd/system/fjvsdx.service.d/netmirror.conf
# chmod 644 /etc/systemd/system/fjvsdx.service.d/netmirror.conf
```

Add the following content to the setting file.

```
[Unit]
After=target.service
```

Apply the setting of the start and stop order for iSCSI target service.

```
# systemctl daemon-reload
```

3. Check the setting of the start and stop order for iSCSI target service.

```
# systemctl show fjsvsdx.service | grep "After="  
After=iscsi-shutdown.service iscsi.service target.service ...
```

Confirm that target.service is output.

4. Set to automatically start the iSCSI target service when the system is started.

```
# systemctl enable target.service
```

4.8.3 Creating iSCSI Target

Set the following configuration for all nodes in which the mirroring function among servers is used.

The "current node" in the following procedure indicates: node 1 if you execute the procedure on the node 1; node 2 if you execute the procedure on the node 2.

1. Execute targetcli to be an interactive mode.

```
# targetcli
```

2. Set not to add a default portal automatically.

```
/> set global auto_add_default_portal=false
```

3. Register a disk or a partition of a local node for the mirroring among servers.

```
/> /backstores/block/ create storage_name device_path
```

Point

- Specify arbitrary characters (255 characters or less) to *storage_name*.
- For *device_path*, specify the path for the by-id file of a disk or a partition on a local node which is used for the mirroring among servers.
In an Azure environment, a NIFCLOUD environment, or an environment where the by-id file does not exist, do not specify the path for the by-id file, but create a partition in the disk where the gpt disk label is set, and specify the path for the by-partuuid file for that partition.
- If there are plural disks for the mirroring among servers on a local node, perform the procedure for each disk.

Execution example

```
/> /backstores/block/ create storage1 /dev/disk/by-id/scsi-3500000e111c56611
```

4. Check the setting of device.

Execute the following command to check that emulate_write_cache is 0.

```
/> /backstores/block/storage_name get attribute emulate_write_cache
```

Point

Specify the storage name set in the step 3 to *storage_name*.

Execution example

```
/> /backstores/block/storage1 get attribute emulate_write_cache
emulate_write_cache=0
```

5. Create iSCSI target.

```
/> /iscsi/ create
```

6. Check the created iSCSI target.

```
/> /iscsi/ ls
```

Execution example

```
/> /iscsi/ ls
o- iscsi ..... [Targets: 1]
  o- iqn.2003-01.org.linux-iscsi.node1.x8664:sn.2613f8620d98 ..... [TPGs: 1]
    o- tpg1 ..... [no-gen-acls, no-auth]
      o- acls ..... [ACLs: 0]
      o- luns ..... [LUNs: 0]
      o- portals ..... [Portals: 0]
```

This example shows that iqn name "iqn.2003-01.org.linux-iscsi.node1.x8664:sn.2613f8620d98" was created and the target was created in /iscsi/.

7. Register iSCSI initiators of both nodes.

```
/> /iscsi/iqn_name/tpgX/acls create initiator_name_of_node1
/> /iscsi/iqn_name/tpgX/acls create initiator_name_of_node2
```

Point

- Specify the iqn name confirmed in the step 6 to *iqn_name*.
- Specify the directory name (Example: tpg1) created in the *iqn_name* directory to *tpgX*.
- *Initiator name* can be confirmed in the /etc/iscsi/initiatorname.iscsi file.

Execution example

```
/> /iscsi/iqn.2003-01.org.linux-iscsi.node1.x8664:sn.2613f8620d98/tpg1/acls create \
iqn.1994-05.com.redhat:dee92ff9979d
/> /iscsi/iqn.2003-01.org.linux-iscsi.node1.x8664:sn.2613f8620d98/tpg1/acls create \
iqn.1994-05.com.redhat:fa7eb9cf483c
```

Note

An iSCSI initiator name needs to be differed in each node of the cluster system.

If an iSCSI initiator name in each node is identical, change it to a different iSCSI initiator name in each node.

See the following procedure to change an iSCSI initiator name.

1. Determine a new iSCSI initiator name.
Use the following command to create an initiator name with random ID.

(Example)

```
# iscsi-iname  
iqn.1994-05.com.redhat:5ea612eab618
```

2. Change a configuration file.
In the `/etc/iscsi/initiatorname.iscsi` file, rewrite a character string following "InitiatorName=" to the iSCSI initiator name of the step 1.

(Example)

[Before change]

```
InitiatorName=iqn.1994-05.com.redhat:4539cd846026
```

[After change]

```
InitiatorName=iqn.1994-05.com.redhat:5ea612eab618
```

3. Restart `iscsid` to update the configuration.

(Example)

```
# systemctl restart iscsid.service
```

8. Connect the device registered in the step 3.

```
/> /iscsi/iqn_name/tpgX/luns create /backstores/block/storage_name
```

Point

- Specify `iqn` name confirmed in the step 6 to `iqn_name`.
- Specify the directory name (Example: `tpg1`) created in the `iqn_name` directory to `tpgX`.
- Specify the storage name set in the step 3 to `storage_name`.
- If there are plural disks for the mirroring among servers on a local node, perform the procedure for each disk.

Execution example

```
/> /iscsi/iqn.2003-01.org.linux-iscsi.node1.x8664:sn.2613f8620d98/tpg1/luns create \  
/backstores/block/storage1
```

9. Register IP address for the target.

```
/> /iscsi/iqn_name/tpgX/portals create IPaddress
```

Point

- Specify the `iqn` name confirmed in the step 6 to `iqn_name`.
- Specify the directory name (Example: `tpg1`) created in the `iqn_name` directory to `tpgX`.
- Specify IP address for the mirroring among servers on a local node to `IPaddress`.
IPv4 address or IPv6 address can be set.

IP address of cip can not be specified. Use IP address assigned NIC (such as eth0) directly.
For conditions for NIC and IP address, refer to "3.9.1 Network Configuration."

Execution example

```
/> /iscsi/iqn.2003-01.org.linux-iscsi.nodel.x8664:sn.2613f8620d98/tpg1/portals create \  
192.168.125.123
```

10. End the interactive mode of targetcli. This operation records the settings.

```
/> exit
```

11. Confirm the target configuration.

```
# targetcli ls
```

Point

(*1) Confirm that the path for the by-id file of the device used for the mirroring among servers was output.

If the by-partuuid file is specified in an Azure environment, a NIFCLOUD environment, or an environment where the by-id file does not exist, confirm that the path for the by-partuuid file is output.

(*2) Confirm that iSCSI initiator name for each cluster node was output.

(*3) Confirm that the storage confirmed in (*1) was assigned for each iSCSI initiator.

(*4) Confirm that IP address of a local node for the mirroring among servers was output.

Execution example

```
# targetcli ls  
o- / ..... [...]  
  o- backstores ..... [...]  
    | o- block ..... [Storage Objects: 1]  
    | | o- storage1 . [/dev/disk/by-id/scsi-3500000e111c56611 (8.0GiB) write-thru activated] (*1)  
    | o- fileio ..... [Storage Objects: 0]  
    | o- pscsi ..... [Storage Objects: 0]  
    | o- ramdisk ..... [Storage Objects: 0]  
  o- iscsi ..... [Targets: 1]  
    | o- iqn.2003-01.org.linux-iscsi.nodel.x8664:sn.2613f8620d98 ..... [TPGs: 1]  
      | o- tpg1 ..... [no-gen-acls, no-auth]  
        | o- acls ..... [ACLs: 2]  
          | o- iqn.1994-05.com.redhat:dee92ff9979d ..... [Mapped LUNs: 1] (*2)  
            | | o- mapped_lun0 ..... [lun0 block/storage1 (rw)] (*3)  
            | o- iqn.1994-05.com.redhat:fa7eb9cf483c ..... [Mapped LUNs: 1] (*2)  
              | o- mapped_lun0 ..... [lun0 block/storage1 (rw)] (*3)  
            o- luns ..... [LUNs: 1]  
              | o- lun0 ..... [block/storage1 (/dev/disk/by-id/scsi-3500000e111c56611)]  
            o- portals ..... [Portals: 1]  
              o- 192.168.125.123:3260 ..... [OK] (*4)  
  o- loopback ..... [Targets: 0]
```

4.8.4 Establishing iSCSI Session

Set the following configuration for all nodes in which the mirroring function among servers is used.

When performing the mirroring among servers in node 1 and node 2, the "current node" and "another node" in the following procedure are indicated nodes below.

The node to perform the procedure	Current node	Another node
Node 1	Node 1	Node 2
Node 2	Node 2	Node 1

1. Search and register the target portal of the current node and confirm the portal name of the current node.

```
# iscsiadm -m discovery --type sendtargets --portal IPaddress
```

Point

- For *IPaddress*, specify the IP address for the mirroring among servers on the current node.
- Confirm the portal name displayed in the same line as the target name of the current node which was set in "[4.8.3 Creating iSCSI Target](#)."

Execution example

```
# iscsiadm -m discovery --type sendtargets --portal 192.168.56.20
192.168.56.20:3260,1 target-0101
```

In this example, the target name and the portal name of the current node are as follows:

- Target name: target-0101
- Portal name: 192.168.56.20:3260,1

2. Change the connection setup with the target portal of the current node.

```
# iscsiadm --mode node --op update --targetname targetname --name=node.startup --value=manual
# iscsiadm -m node -p portalname -o update -n node.session.timeo.replacement_timeout -v 30
# iscsiadm -m node -p portalname -o update -n 'node.conn[0].timeo.noop_out_timeout' -v 5
# iscsiadm -m node -p portalname -o update -n 'node.conn[0].timeo.noop_out_interval' -v 5
```

Point

- For *targetname*, specify the target name of the current node.
- For *portalname*, specify the portal name of the current node.

Execution example:

```
# iscsiadm --mode node --op update --targetname target-0101 --name=node.startup --value=manual
# iscsiadm -m node -p 192.168.56.20:3260,1 -o update -n node.session.timeo.replacement_timeout
-v 30
# iscsiadm -m node -p 192.168.56.20:3260,1 -o update -n 'node.conn[0].timeo.noop_out_timeout'
-v 5
# iscsiadm -m node -p 192.168.56.20:3260,1 -o update -n 'node.conn[0].timeo.noop_out_interval'
-v 5
```

3. Search and register the target portal of another node and confirm the portal name of another node.

```
# iscsiadm -m discovery --type sendtargets --portal IPaddress
```

Point

- For *IPaddress*, specify the IP address for the mirroring among servers on another node.
- Confirm the portal name displayed in the same line as the target name of another node which was set in "[4.8.3 Creating iSCSI Target](#)."

Execution example:

```
# iscsiadm -m discovery --type sendtargets --portal 192.168.56.10
192.168.56.10:3260,1 target-0201
```

In this example, the target name and the portal name of the another node are as follows:

- Target name: target-0201
- Portal name: 192.168.56.10:3260,1

4. Change the connection setup with the target portal of another node.

```
# iscsiadm --mode node --op update --targetname targetname --name=node.startup --value=manual
# iscsiadm -m node -p portalname -o update -n node.session.timeo.replacement_timeout -v 30
# iscsiadm -m node -p portalname -o update -n 'node.conn[0].timeo.noop_out_timeout' -v 5
# iscsiadm -m node -p portalname -o update -n 'node.conn[0].timeo.noop_out_interval' -v 5
```

Point

- For *targetname*, specify the target name of another node.
- For *portalname*, specify the portal name of another node.

Execution example

```
# iscsiadm --mode node --op update --targetname target-0201 --name=node.startup --value=manual
# iscsiadm -m node -p 192.168.56.10:3260,1 -o update -n node.session.timeo.replacement_timeout
-v 30
# iscsiadm -m node -p 192.168.56.10:3260,1 -o update -n 'node.conn[0].timeo.noop_out_timeout'
-v 5
# iscsiadm -m node -p 192.168.56.10:3260,1 -o update -n 'node.conn[0].timeo.noop_out_interval'
-v 5
```

5. Log in to the target of the current node.

```
# iscsiadm -m node -T targetname --portal portalname --login
```

Point

- For *targetname*, specify the target name of current node.
- For *portalname*, specify the portal name of current node.

Execution example:

```
# iscsiadm -m node -T target-0101 --portal 192.168.56.20:3260,1 --login
```


6. Log in to the target of another node.

```
# iscsiadm -m node -T targetname --portal portalname --login
```

Point

- For *targetname*, specify the target name of another node.
- For *portalname*, specify the portal name of another node.

Execution example:

```
# iscsiadm -m node -T target-0201 --portal 192.168.56.10:3260,1 --login
```

7. Confirm the status of the iSCSI session and the device name of the iSCSI device.

```
# iscsiadm -m session -P 3
```

Point

Confirm the following:

- The connection status and session status of iSCSI must be "LOGGED IN."
- For the iSCSI devices on the current node for the disks on the current node and another node, on which the mirroring among servers is performed, the status must be "running".

Execution example

```
# iscsiadm -m session -P 3
iSCSI Transport Class version xxxxxxxxx
version xxxxxxxxx
Target: target-0201
  Current Portal: 192.168.56.10:3260,1
  Persistent Portal: 192.168.56.10:3260,1
  *****
  Interface:
  *****
  ...
  iSCSI Connection State: LOGGED IN
  iSCSI Session State: LOGGED_IN
  Internal iscsid Session State: NO CHANGE
  ...
  *****
  Attached SCSI devices:
  *****
  Host Number: 12 State: running
  scsi12 Channel 00 Id 0 Lun: 0
  scsi12 Channel 00 Id 0 Lun: 1
           Attached scsi disk sdd           State: running

Target: target-0101
  Current Portal: 192.168.56.20:3260,1
  Persistent Portal: 192.168.56.20:3260,1
  *****
  Interface:
  *****
```

```

...
iSCSI Connection State: LOGGED IN
iSCSI Session State: LOGGED_IN
Internal iscsid Session State: NO CHANGE
...
*****
Attached SCSI devices:
*****
Host Number: 13 State: running
scsi13 Channel 00 Id 0 Lun: 0
scsi13 Channel 00 Id 0 Lun: 1
        Attached scsi disk sde                State: running

```

In this example, the iSCSI device name of the current node is "sde" and the iSCSI device name of another node is "sdd".

8. Check that the symbolic link file indicating the created iSCSI device has been created in /etc/opt/FJSVsdx/.sdx_by_id directory. It may take some time to create the symbolic link file. If the symbolic link file does not exist, check it again after some time.

```
# ls -l /etc/opt/FJSVsdx/.sdx_by_id/
```

Execution example:

```
# ls -l /etc/opt/FJSVsdx/.sdx_by_id/*
lrwxrwxrwx 1 root root 8 Dec  4 17:57 /etc/opt/FJSVsdx/.sdx_by_id/sdx_1SDX01 -> /dev/sde
lrwxrwxrwx 1 root root 8 Dec  4 18:23 /etc/opt/FJSVsdx/.sdx_by_id/sdx_1SDX02 -> /dev/sdd
```



Note

- After configuring this setting, do not restart the node until all the settings described in "6.3 Shared Disk Resource Registration" are configured. If you have already started the node, take the step "5. Log in to the target of the current node." and the following steps again.
- After registering the iSCSI device that is set by this procedure, do not log out of the target. When logging out of the target by mistake, restart the node.

4.8.5 Confirming the Correspondence of iSCSI Device Name of Each Node

In the configuration of "Figure 4.1 Disk setting for Performing Mirroring among Servers" each disk can be accessed from node 1 and node 2 in the same iSCSI device name. However, a different iSCSI device name for each node may be assigned.

This section describes how to confirm the iSCSI device name in each node for the disk device on which the mirroring among servers is performed.

1. Confirm the correspondence of the iSCSI device name of node 1 and the by-id link file.

Execute the following command to make sure to check that the by-id link file for the iSCSI device created in "4.8.4 Establishing iSCSI Session."

```
# ls -l /dev/disk/by-id
...
lrwxrwxrwx 1 root root  9 Apr 18 08:44 scsi-1SDX01 -> ../../sdg
lrwxrwxrwx 1 root root  9 Apr 18 08:43 scsi-1SDX02 -> ../../sdh
...
```

2. Confirm the correspondence of the iSCSI device name of node 2 and the by-id link file.

In node 2, confirm in the same procedure of step 1.

```
# ls -l /dev/disk/by-id
...
lrwxrwxrwx 1 root root  9 Apr 18 08:44 scsi-1SDX01 -> ../../sdf
```

```
lrwxrwxrwx 1 root root 9 Apr 18 08:43 scsi-1SDX02 -> ../../sdg
...
```

In this example, the iSCSI device names of each node are as follows.

Disk	by-id link file name	iSCSI device name	
		Node 1	Node 2
The local disk of node 1	scsi-1SDX01	sdg	sdf
The local disk of node 2	scsi-1SDX02	sdh	sdg

4.9 Preparation for Starting GDS Management View

GDS Management View manages and monitors objects by using the GUI screen.

This chapter explains the preparation to start the GDS Management View.

GDS Management View uses Web-Based Admin View.

For details on Web-Based Admin View, see "PRIMECLUSTER Web-Based Admin View Operation Guide."

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

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

4.9.1 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.

4.9.1.1 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 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.

4.9.1.2 Creating User Groups

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

The other user group sdxroot, explained in "4.9.1.1 User Group Types" must be created as necessary.

User group can be created using the following command.

```
# groupadd sdxroot
```

4.9.1.3 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.

4.9.2 Setting up the Client Environment

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

4.10 Settings to Use Dell EMC PowerPath

When using Dell EMC PowerPath or Dell EMC PowerPath/VE, perform the following settings before registering the disk of Dell EMC storage units to the PRIMECLUSTER resource database and the GDS classes.

1. Edit the configuration parameter file `/etc/opt/FJVSvdx/sdx.cf` by using the editor such as `vim(1)`.

- When "SDX_UDEV_USE=on" is described
Change "on" to "off."
- When "SDX_UDEV_USE=on" is not described
Add the description of "SDX_UDEV_USE=off."

```
# vim /etc/opt/FJVSvdx/sdx.cf
...
SDX_UDEV_USE=off
...
```

2. Reboot the system.

```
# shutdown -r now
```

Chapter 5 GDS Management View

This chapter explains how to start and exit GDS Management View and its screen configuration and functions.

5.1 Starting GDS Management View

5.1.1 Starting Web-Based Admin View Operation Menu

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

The Web-Based Admin View screen can be started in the following way.

- Startup from the Java application

Start the screen from the Java application installed on a client.

Procedures

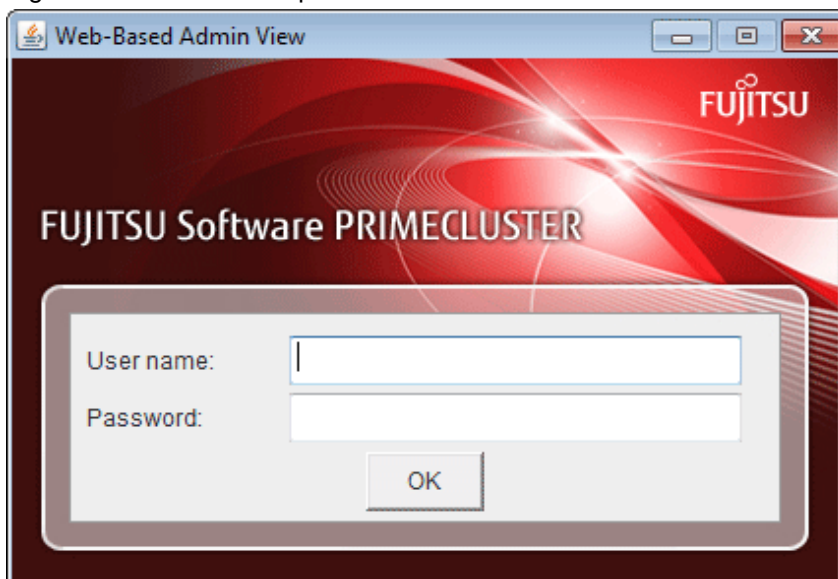
- Startup from the Java application

Start the screen from either of the following shortcuts (PRIMECLUSTER Web-Based Admin View Startup) on a client.

- Shortcut registered in the Java application program group
- Desktop shortcut

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

Figure 5.1 User Name Input Screen



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

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


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

Figure 5.2 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.


5.1.3 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 "[5.6 Operation](#)"

Figure 5.3 GDS Management Main Screen



5.2 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 5.4 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 "5.4 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 "5.5 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 "[5.3 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.

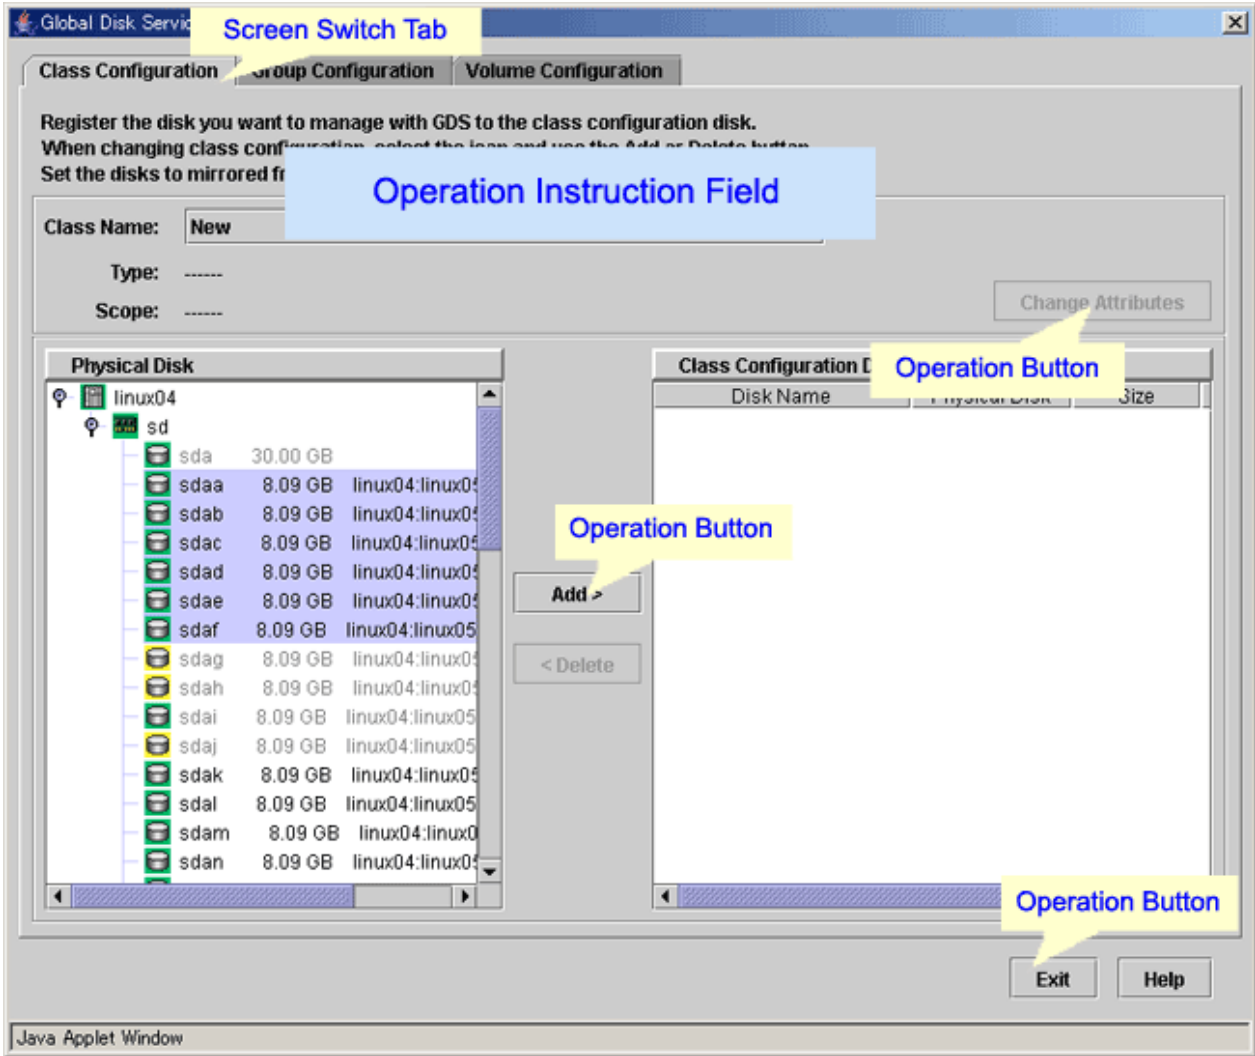
Pilot Lamp	Status	Meaning
 (Gray lit up)	Normal	Normal status.
 (Red blinking)	Abnormal	Critical abnormality (e.g. closed class, invalid volume)
 (Red lit up)	Abnormal	When red blinking warning lamp is single-clicked.
 (Yellow blinking)	Degradation	- Volume operating at degradation. (e.g. invalid slice, 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 5.5 GDS Configuration Settings Screen



For information on how to set each configuration, see "6.6.1.3 Class Configuration," "6.6.1.4 Cluster System Class Configuration," "6.6.1.5 Group Configuration," and "6.6.1.6 Volume Configuration."

5.3 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.

5.3.1 [General] Menu



Change Monitoring Intervals

Sets the monitoring interval (min) of objects.

Figure 5.6 [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 5.7 [General]: [Exit] Screen



5.3.2 [Settings] Menu

Figure 5.8 [Settings] menu




Class Configuration

Sets the class configuration.

For details, see "[6.6.1.3 Class Configuration](#)," or "[6.6.1.4 Cluster System Class Configuration](#)."

Figure 5.9 [Settings]: [Class Configuration] Screen

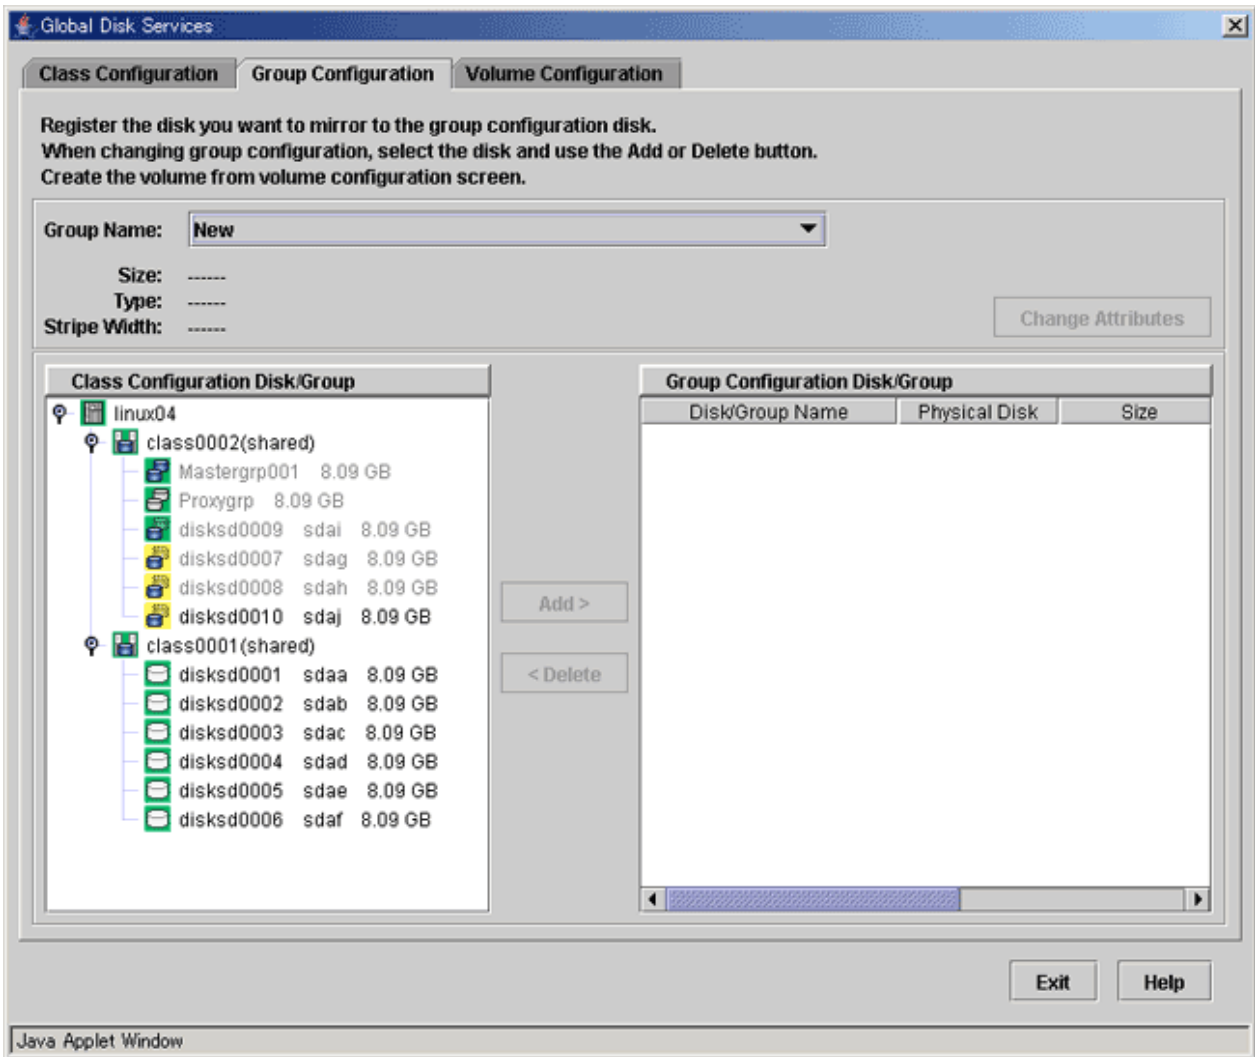


Group Configuration

Sets the group configuration.

For details, see "[6.6.1.5 Group Configuration](#)."

Figure 5.10 [Settings]: [Group Configuration] Screen

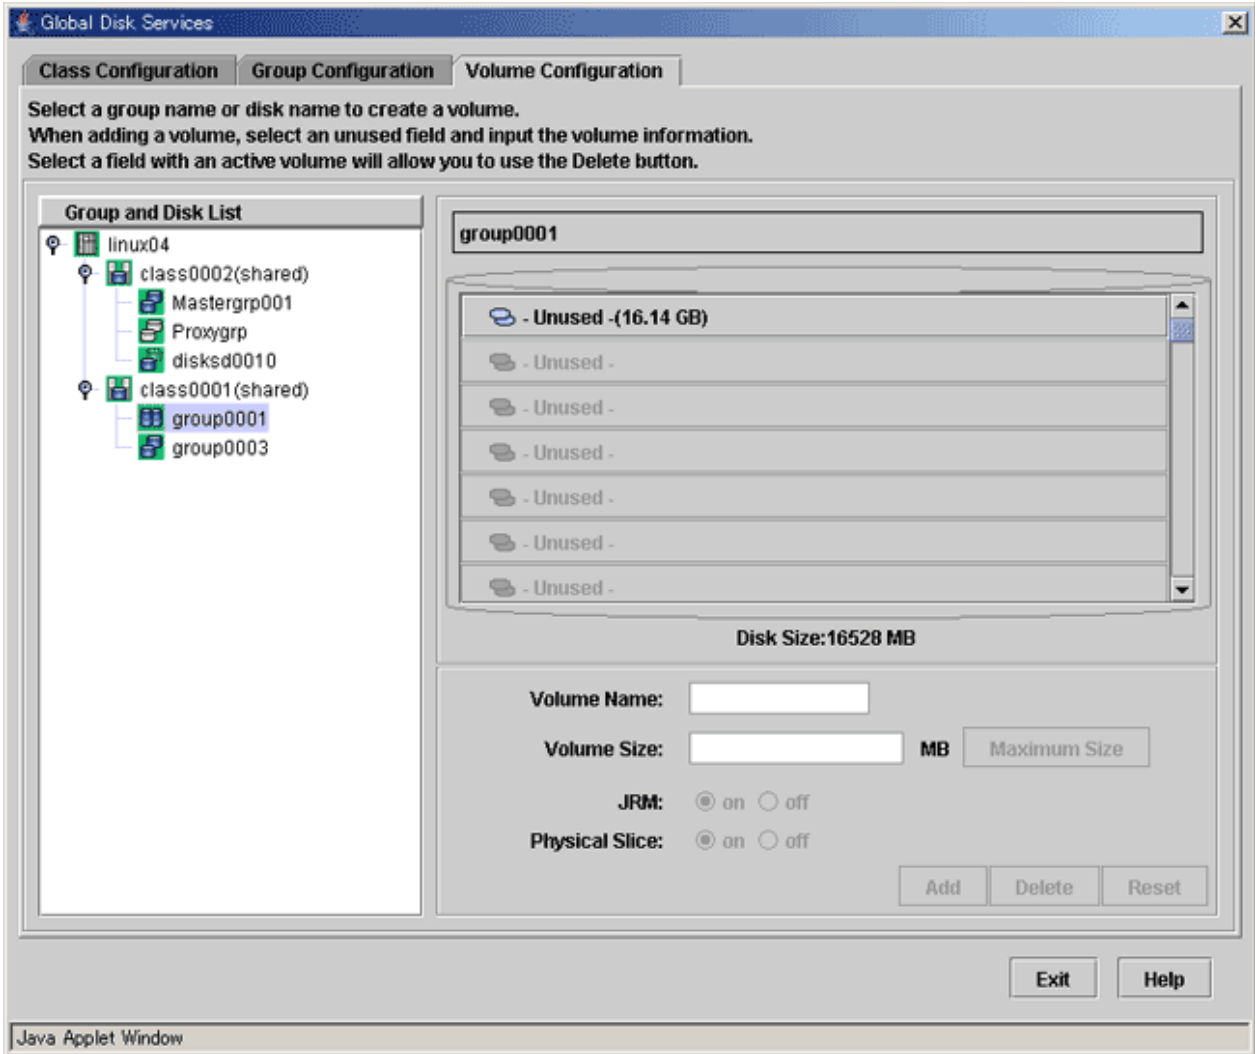


Volume Configuration

Sets the volume configuration.

For details, see "[6.6.1.6 Volume Configuration](#)."

Figure 5.11 [Settings]: [Volume Configuration] Screen




System Disk Settings [EFI]

Mirrors the system disk.

For details, see "[6.5.2.1.2 Operation Procedure.](#)"

Figure 5.12 [Settings]: [System Disk Settings] Screen




Unmirror System Disk [EFI]


Unmirrors the system disk.

For details, see "[9.4.1.2 Operation Flow.](#)"

Figure 5.13 [Settings]: [Unmirror System Disk] Screen



5.3.3 [Operation] Menu



Swap Physical Disk

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

For details, see "[7.3.1.2 Operation Procedure.](#)"

Restore Physical Disk

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

For details, see "[7.3.1.2 Operation Procedure.](#)"

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 "[7.2.1 Using GDS Management View.](#)"

Attach Slice

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

For details, see "[7.2.1 Using GDS Management View.](#)"

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 "[7.7.1.2 Operation Procedure.](#)"

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 "[7.2.1 Using GDS Management View.](#)"

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 "[7.2.1 Using GDS Management View.](#)"

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 "[6.10.1.1 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 "[Part](#)."

Rejoin

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

For details, see "[Part](#)."

Relate

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

For details, see "[6.10.1.2 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 "[7.9.1.2 Operation Procedure](#)."

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 "[7.10.1.2 Operation Procedure](#)."

Swap

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

For details, see "[8.7.1.1.2 Operation Procedure](#)."

Break

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

For details, see "[9.1 Breaking Proxy](#)."

Change Attributes

Changes the attributes of the selected object.

For details, see "[8.2.1 Using GDS Management View](#)."

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 resource registration was performed.

Note

When [Update Physical Disk Information] is executed, among the by-id files (symbolic link files created under the /dev/disk/by-id directory) managed by the udev function of Linux, the by-id files of disks which are not registered in classes may not exist temporarily because OS creates the by-id files again triggered by the GDS processes for obtaining the device information.


When access to the by-id files of disks which are not registered in classes fails, retry access.

Alternatively, when [Update Physical Disk Information] is executed, do not access to the by-id files of disks which are not registered in classes.

Check Status

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

5.3.4 [View] Menu



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 "[5.5 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.

5.3.5 [Help] Menu



Help

Displays Help information.

5.4 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) [EFI]

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




6. Groups (mirror or netmirror)

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
 (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
 (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.

Identify the hardware error on a disk unit based on the log messages for the disk driver which is output to the /var/log/messages file. For details, see "[D.1.13 Disk Unit Error](#)."

5.5 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

GDS: Global Disk Services

- [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

GDS: Global Disk Services

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

GDS Snapshot: Global Disk Services Snapshot

5.6 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.

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	See "6.6.1.3 Class Configuration."
Remove	# sdxdisk -R or # sdxclass -R	See "9.3.1.4 Removing Class."
Change attribute (Class name)	# sdxattr -C -a name= <i>classname</i>	See "8.2.1.2 Class Configuration."

Operation	Command Operation	GDS Management View Operation
Change attribute (Type)	# sdxattr -C -a type={local shared}	See "8.2.1.2 Class Configuration."
Change attribute (Scope)	# sdxattr -C -a scope= <i>node: node:...</i>	See "8.2.1.2 Class Configuration."
Change attribute (Hot spare)	# sdxattr -C -a hs={on off}	Unsupported
Change attribute (Hot spare mode)	# sdxattr -C -a hsmode={exbox bybox}	Unsupported
Change attribute (I/O Response Time Assurance)	# sdxattr -C -a iotimeout={on[: <i>time[:panic]</i>] off}	Unsupported
Restore	# sdxfix -V	Unsupported

Group operation

Operation	Command Operation	GDS Management View Operation
Connect	# sdxdisk -C or # sdxgroup -C	See "6.6.1.5 Group Configuration."
Remove	# sdxdisk -D or # sdxgroup -D or # sdxgroup -R	See "9.3.1.3 Removing Group."
Change attribute (Group name)	# sdxattr -G -a name= <i>groupname</i>	See "8.2.1.3 Group Configuration."
Change attribute (Active disk)	# sdxattr -G -a actdisk= <i>disk</i>	Unsupported

Volume operation

Operation	Command Operation	GDS Management View Operation
Create	# sdxvolume -M	See "6.6.1.6 Volume Configuration."
Remove	# sdxvolume -R	See "9.3.1.3 Removing Group."
Start	# sdxvolume -N	See "Start Volume" in "5.3.3 [Operation] Menu."
Stop	# sdxvolume -F	See "Stop Volume" in "5.3.3 [Operation] Menu."
Expand	# sdxvolume -S	Unsupported
Change attribute (Volume name)	# sdxattr -V -a name= <i>volumename</i>	See "8.2.1.4 Volume Configuration."
Change attribute (JRM)	# sdxattr -V -a jrm={on off}	See "8.2.1.4 Volume Configuration."
Change attribute (Physical slice)	# sdxattr -V -a pslice={on off}	See "8.2.1.4 Volume Configuration."
Change attribute (Lock volume)	# sdxattr -V -a lock={on off}	Unsupported
Change attribute (Access mode)	# sdxattr -V ... -a mode={rw ro}	Unsupported

Operation	Command Operation	GDS Management View Operation
Change attribute (JRM for proxy)	# sdxattr -V -a pjrm= off	Unsupported
Start copying	# sdxcopy -B	See "7.2.1 Using GDS Management View."
Cancel copying	# sdxcopy -C	See "7.2.1 Using GDS Management View."
Restore	# sdxfix -V	Unsupported

GDS: Global Disk Services

Slice operation

Operation	Command Operation	GDS Management View Operation
Detach	# sdxslice -M	See "7.7.1.2 Operation Procedure."
Attach	# sdxslice -R	See "7.7.1.2 Operation Procedure."
Activate	# sdxslice -N	See "7.7.1.2 Operation Procedure."
Stop	# sdxslice -F	See "7.7.1.2 Operation Procedure."
Take over	# sdxslice -T	Unsupported
Change attribute	# sdxattr -S	Unsupported

Disk operation

Operation	Command Operation	GDS Management View Operation
Swap	# sdxswap -O	See "7.3.1.2 Operation Procedure."
Restore	# sdxswap -I	See "7.3.1.2 Operation Procedure."
Change attribute	# sdxattr -D	See "8.2.1.2 Class Configuration."
Clear errors	# sdxfix -D	Unsupported

Proxy operation

Operation	Command Operation	GDS Management View Operation
Join	# sdxproxy Join	See "6.10.1.1 Join."
Part	# sdxproxy Part	See "Part."
Rejoin	# sdxproxy Rejoin	See "Part."
Rejoin and restore	# sdxproxy RejoinRestore	See "7.10.1.2 Operation Procedure."
Swap	# sdxproxy Swap	See "8.7.1.1.2 Operation Procedure."
Relate	# sdxpsoxy Relate	See "6.10.1.2 Relate."
Update	# sdxpsoxy Update	See "7.9.1.2 Operation Procedure."
Restore	# sdxpsoxy Restore	See "7.10.1.2 Operation Procedure."
Cancel hard copy session	# sdxpsoxy Cancel	Unsupported
Create alternative boot environment [EFI]	# sdxproxy Root	Unsupported

Operation	Command Operation	GDS Management View Operation
Break	# sdxproxy Break	See " 9.1 Breaking Proxy. "

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	See " 7.1 Viewing Configurations/Statuses and Monitoring Statuses. "

Other operations


Operation	Command Operation	GDS Management View Operation
System disk settings [EFI]	# sdxdisk -M # sdxdisk -C # sdxroot -M <reboot>	See " 6.5.2.1.2 Operation Procedure. "
Unmirror system disk [EFI]	# sdxdisk -D # sdxroot -R <reboot> # sdxvolume -F # sdxvolume -R # sdxdisk -D # sdxdisk -R	See " 9.4.1.2 Operation Flow. "
Configuration parameter operations	# sdxparam	Unsupported
Object configuration operations	# sdxconfig	Unsupported

5.7 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 5.14 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.

1. Use the following procedure to log out from the Web-Based Admin View.
 - 1-1. If the management screen of the operation management product is displayed, close all screens.
 - 1-2. When only the Web-Based Admin View screen is displayed, select <Log Out>.
2. The login screen will be displayed. With the login screen displayed, click the [X] in the top right corner of the Web-Based Admin View screen.



Note

You cannot exit the login screen by clicking the [X] in the top right corner.

5.8 Uninstalling Java Application

If you no longer need to use the Web-Based Admin View, delete the Java application.

For details, see "PRIMECLUSTER Web-Based Admin View Operation Guide."

Chapter 6 Settings

This chapter explains how to set up GDS, GDS Snapshot, and GDS I/O Monitor Option.

6.1 Creating Excluded Device List

6.1.1 Excluded Device List

The Excluded Device List (/etc/opt/FJSVsdx/lib/exdevtab file) is a list of disks which are to be excluded from GDS management.

In the Excluded Device List, describe the following disks that are to be excluded from GDS management.

- Disks such as native devices (sd devices) which compose multipath disks (mpath devices, emcpower devices)
- Physical disks which compose iSCSI devices for the mirroring among servers

Example of the Excluded Device List:

```
# cat /etc/opt/FJSVsdx/lib/exdevtab
/dev/sde
/dev/sdf
/dev/sdg
/dev/sdh
```

Disks described in the Excluded Device List are protected so that they cannot be mistakenly registered with a class.

- For the GDS Management View, disks are not displayed in the selection screen.
- If an attempt to register a disk with a class by a command, an error message "ERROR: *physical_disk_name*: no such device" is output.



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 resource registration, and do not describe these devices in the shared disk definition file. For details on resource registration and the shared disk definition file, see "6.3 Shared Disk Resource Registration."



Note

Device Name Change

When device name change occurs, change the device name described in the Excluded Device List to the name after device name change occurs.

6.1.2 To Use DM-MP (Device Mapper Multipath)

After DM-MP setup is completed, describe native devices (sd devices) which compose mpath devices to the Excluded Device List /etc/opt/FJSVsdx/lib/exdevtab.

6.1.3 To Use Dell EMC Storage Units

After completing the configuration of devices that are described in "Devices that Cannot be Managed" in "3.3 Dell EMC Storage Units" and Dell 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.

[Procedure]

1. The `syminq` command provided by SYMCLI is available for checking BCV, R2, GateKeeper, and CKD devices. Execute the `syminq` command, and describe all devices indicated as BCV, R2, GK, and CKD (`sdX`, `emcpowerX`) in the Excluded Device List.
2. The `syminq` command is unavailable for checking VCMDB devices. When using Dell EMC's SAN management software (Volume Logix, ESN Manager, SAN Manager and so on), ask your Dell 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. Describe all native devices (`sdX`) in the Excluded Device List.
emcpower0 (Target)
sdb (Nontarget)
sdc (Nontarget)
4. 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.

```
# cat /etc/opt/FJSVsdx/lib/exdevtab
/dev/emcpowerc BCV
/dev/emcpowerd BCV
/dev/emcpowere GK
/dev/emcpowerf GK
/dev/emcpowerg CKD
/dev/emcpowerh R2
/dev/sdb PP
/dev/sdc PP
...
/dev/sdp PP
/dev/sdq PP
#
```

 **Information**

exdevtab.sh

The script sample `"/etc/opt/FJSVsdx/bin/exdevtab.sh"` for simply creating the Excluded Device List `"/etc/opt/FJSVsdx/lib/exdevtab"` is provided.

To use the script, open it with an editor and modify the following parameter (`syminq` command path) according to the execution environment.

```
SYMINQ=/usr/symcli/bin/syminq
```

By executing `exdevtab.sh`, all BCV, GateKeeper, and CKD devices are included in the Excluded Device List. The native devices of `emcpower` devices, R2 devices, and VCMDB 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 4.

 **Note**

Using Dell EMC TimeFinder or Dell EMC SRDF in a Proxy Configuration

For using Dell EMC TimeFinder or Dell EMC SRDF in a proxy configuration, do not describe BCV and R2 devices connected with proxy groups in the Excluded Device List, but describe the native devices that compose those devices in the list. For details, see ["3.19 Using Dell EMC TimeFinder or Dell EMC SRDF in Proxy Configuration."](#)

6.1.4 To Use Mirroring among Servers

After performing "4.8 Disk Setting for Performing Mirroring among Servers", describe the physical disk where the mirroring among servers is performed to the Excluded Device List /etc/opt/FJVSdx/lib/exdevtab.

The configuration example in "Figure 4.1 Disk setting for Performing Mirroring among Servers" shows that the physical disk sdc is described in the Excluded Device List.

When performing mirroring per partition, describe the disk where the target partition exists in the Excluded Device List.

6.2 Using iSCSI Disk

The following tuning parameters can set the waiting time for the startup of GDS.

Parameter	Description
SDX_STARTUP_WAIT_TIME= <i>time</i>	The parameter sets the time (seconds) to wait for the iSCSI device to be recognized by the system. The value is set to 1800 by default. Values from 0 to 43200 can be specified. If the value is set to 0, the waiting time is not set, either.
SDX_ISCSI_INTERVAL= <i>time</i>	The parameter sets the time interval (seconds) to check the iSCSI device when the system is started. The value is set to 2 by default. Values from 1 to 3600 can be specified.

How to set the tuning parameter

1. Add the tuning parameter to the /etc/opt/FJVSdx/sdx.cf file.
2. After adding the parameter, restart the system. The setting is enabled after restarting the system.

Point

Generally, these tuning parameters are not necessary to be set.

Set the tuning parameters if the following message is output when the system is started.

```
ERROR: check the iSCSI device has timeout, iSCSI=num
```

For details, see Explanation and Resolution of "Message 42055" in "C.3.2 Error Messages (42000-42099)."

6.3 Shared Disk Resource Registration

6.3.1 What Is Shared Disk Resource Registration?

In a cluster system, before configuring GDS classes, groups, and volumes, it is required to register shared disk units with the PRIMECLUSTER resource database. The resource database is used to manage information of node and shared disk unit physical connection configuration (how each shared disk unit is connected to each node). GDS recognizes shared disk unit physical connection configuration based on information in the resource database to manage shared disk units.

Note

Shared disk resource registration is necessary also when using the mirroring among servers.

When using the mirroring among servers, register the iSCSI devices that are created in "4.8 Disk Setting for Performing Mirroring among Servers."

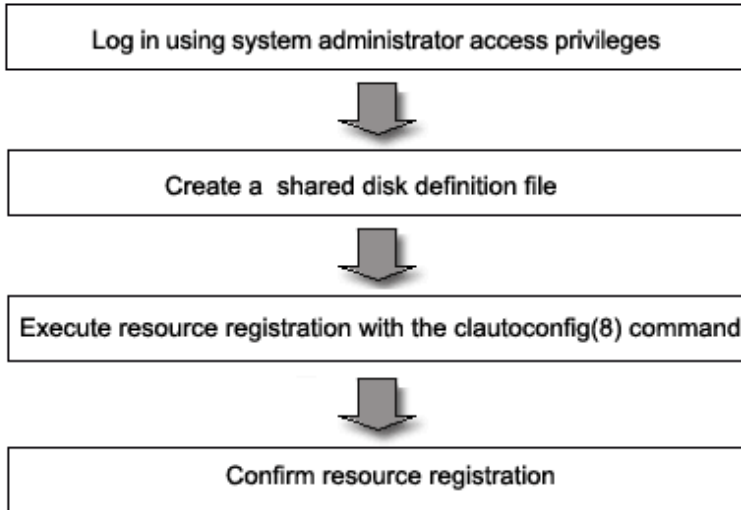
Maximum 16 iSCSI devices can be registered.

For a single node, there are no disks to be shared by each node in the cluster system, so resource registration of shared disks do not need.

6.3.2 Flow of Shared Disk Resource Registration

The following figure outlines the procedures for shared disk unit resource registration.

Figure 6.1 The Flow of Shared Disk Unit Resource Registration



6.3.3 Preconditions and Important Points

- Before performing resource registration, make sure the following particulars.

1. CF, CIP, and resource database settings are complete on all nodes.



For details, see "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide."

2. All nodes and disk units are physically connected.
 3. All nodes are up in multi-user mode.
 4. To use a disk unit in multi-path configuration, multi-path setting is complete.
- Perform resource registration only on a single node that composes the cluster system.
 - Do not execute any GDS or GDS Snapshot command during clautoconfig(8) command execution. If executed, the GDS or GDS Snapshot command may terminate abnormally.
 - During clautoconfig(8) command execution, do not execute the clautoconfig(8) command on any other node that composes the cluster system. If executed, disk unit resource registration is not done properly. In the event that the command was executed, it is necessary to perform the recovery procedure as follows on all nodes that compose the cluster system and then perform resource registration again. In a situation where hardware unit resources other than disk units are registered, it is also necessary to register those resources again.

[Recovery procedure]

1. Resetting the resource database

Reset the resource database using the clinitreset(8) command.



For details, see the manual for the `clinitreset(8)` command.

2. Rebooting the node

Reboot the node.



For details, see "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide."

6.3.4 Shared Disk Unit Resource Registration Procedure

Perform the following procedure on a node that composes the cluster system.

1. Logging in with system administrator access privileges

Log in to a node that composes the cluster system using system administrator access privileges.

2. Creating a shared disk definition file

A shared disk definition file is a file to declare physical connection configuration of nodes in a cluster system and shared disk units. Use the following format.

```
Resource key name Device name Node identification name
Resource key name Device name Node identification name
:
```

- Describe "Resource key name, Device name, and Node identification name" in one line for every shared disk unit connected to each node.
- "Resource key name", "Device name", and "Node identification name" are space delimited.
- Specify values as follows for "Resource name", "Device name", and "Node identification name".

Resource key name

For a resource key name, specify a resource name that indicates a shared relationship for every shared disk unit. The name should be the same by device shared between nodes. The resource key name is specified in "shd number" format. "shd" is a fixed text string. "number" can contain any 4-digit figure. When multiple shared disks exist, specify different "number"s.

[Example] When `/dev/sdb` and `/dev/sdc` are connected to and shared between multiple nodes

```
Resource key name of /dev/sdb: shd0001
Resource key name of /dev/sdc: shd0002
```

Device name

For a device name, specify the device path of a shared disk unit in the full path.

When using the mirroring among servers, specify the device path of the iSCSI devices that are created in "[4.8 Disk Setting for Performing Mirroring among Servers](#)" in the full path. Maximum 16 iSCSI devices can be described.

[Example] When the device is `/dev/sdb`.

```
/dev/sdb
```



- When using DM-MP
 - Describe a device name with `/dev/mapper/mpathX` format.
 - Do not describe a device name with `/dev/dm-X` format.

- Do not describe a native device (sd device) which composes mpath devices.
- For a guest in the virtual environment
 - Describe a device name for a guest. For example, for the virtio block device of the KVM guest, describe the device name for the KVM guest /dev/vdX, not the device name for the host OS /dev/sdX.

Node identification name

For a node identification name, specify the ID number of a node where the shared disk unit resides. The node identification names can be checked with the clgettree(1) command.



See

For details on the clgettree(1) command, see the clgettree(1) command manual.

[Example] In the following example, node1 and node2 are node identification names.

```
# /etc/opt/FJSVcluster/bin/clgettree
Cluster 1 cluster
      Domain 2 PRIME
            Shared 7 SHD_PRIME
            Node 3 node1 ON
            Node 5 node2 ON
```

The following example shows a shared disk definition file created when "/dev/sdb" and "dev/sdc" are shared between node1 and node2.

```
shd0001 /dev/sdb node1
shd0001 /dev/sdb node2
shd0002 /dev/sdc node1
shd0002 /dev/sdc node2
```

- Save the shared disk definition file without deleting it.
- When adding a shared disk device and registering the added shared disk device on the resource database, define only the information of the added shared disk device.

Example: When registering the added disk device /dev/sdd (*1) on the resource database after shd0001 and shd0002 are already registered on the resource database:

```
shd0003 /dev/sdd node1
shd0003 /dev/sdd node2
```

*1) Note

The device name of the added shared disk device may not follow the device name of the registered device in alphabetical order. Make sure to check the device name of the added shared disk device before defining the information of the added disk device.

3. Resource registration with the clautoconfig(8) command

Execute the clautoconfig(8) command to register information in the shared disk definition file created in step 2. with the resource database using the following format:

[Format]

```
/etc/opt/FJSVcluster/bin/clautoconfig -f Full_path_of_the_shared_disk_definition_file
```

[Example]

```
# /etc/opt/FJSVcluster/bin/clautoconfig -f /var/tmp/diskfile
```

Note

- If the clautoconfig(8) command terminates abnormally, take corrective actions as indicated in a displayed message. For clautoconfig(8) command messages, see "PRIMECLUSTER Messages."
- If a device name of a shared disk unit varies depending on nodes, execute the clautoconfig(8) command on the node where all the device files described in the shared disk definition file exist. If the device files described in the shared disk definition file do not exist on the node where the clautoconfig(8) command was executed, the following error messages are output and the resource registration fails.

```
FJSVcluster: ERROR: clautoconfig: 6900: Automatic resource registration processing terminated abnormally. (detail: /dev/sdX)
```

Example) The case where the following device files exist in each node is described.

Node	Internal disk	Shared disk
node1	sda, sdb, and sdc	sdd and sde
node2	sda and sdb	sdc and sdd

In this case, the contents of the shared disk definition file is as follows:

```
shd0001 /dev/sdd node1
shd0001 /dev/sdc node2
shd0002 /dev/sde node1
shd0002 /dev/sdd node2
```

In the shared definition file, device files of sdc, sdd, and sde are described.

Device files of sdc, sdd, and sde exist on node node1, however the device file of sde does not exist on node node2. Therefore, in this example, you need to execute the clautoconfig(8) command on node node1.

4. Confirming resource registration

Check information of the shared disk units registered with the resource database using the clgettree(1) command.

[Example]

The following example shows clgettree(1) command output generated when resource registration is performed using the shared disk definition file shown as an example in step 2.

The resources in the "DISK" class are shared disk unit resources. The resources in the "SHD_DISK" class are resources that indicate shared disk unit shared relationships.

```
# /etc/opt/FJSVcluster/bin/clgettree
Cluster 1 cluster
  Domain 2 PRIME
    Shared 7 SHD_PRIME
      SHD_DISK 14 SHD_Disk14 UNKNOWN
        DISK 15 sdb UNKNOWN node1
        DISK 16 sdb UNKNOWN node2
      SHD_DISK 17 SHD_Disk17 UNKNOWN
        DISK 18 sdc UNKNOWN node1
        DISK 19 sdc UNKNOWN node2
    Node 3 node1 ON
      DISK 15 sdb UNKNOWN
      DISK 18 sdc UNKNOWN
    Node 5 node2 ON
      DISK 16 sdb UNKNOWN
      DISK 19 sdc UNKNOWN
```

Note

The `clautoconfig(8)` command does not check shared disk units in a specified shared disk definition file whether they are physically connected to nodes. If any error is found in the shared disk definition file after `clautoconfig(8)` command execution, it is necessary to perform the recovery procedure as follows on all nodes that compose the cluster system and then perform resource registration again. In a situation where hardware unit resources other than shared disk units are registered, it is also necessary to register those resources again.

[Recovery procedure]

1. Resetting the resource database

Reset the resource database using the `clinitreset(8)` command.

See

For details, see the `clinitreset(8)` command manual.

2. Rebooting the node

Reboot the node.

6.4 Updating Physical Disk Information

After registering the shared disk device as a resource and before defining the configuration of GDS objects such as classes and volumes, select [Update Physical Disk Information] from the [Operation] menu in GDS Management View.

Note

When [Update Physical Disk Information] is executed, among the by-id files (symbolic link files created under the `/dev/disk/by-id` directory) managed by the `udev` function of Linux, the by-id files of disks which are not registered in classes may not exist temporarily because OS creates the by-id files again triggered by the GDS processes for obtaining the device information.

When access to the by-id files of disks which are not registered in classes fails, retry access.

Alternatively, when [Update Physical Disk Information] is executed, do not access to the by-id files of disks which are not registered in classes.

6.5 System Disk Mirroring Settings [EFI]

Take the following procedure to set the system disk mirroring:

- 1) Checking the partition configuration of the system disk
- 2) Setting the system disk mirroring
- 3) Checking system disk settings
- 4) Checking physical disk information and the slice numbers
- 5) Collecting backups of the system disk
- 6) Backing up EFI configuration information

Note

Backing up system disk

- Data backed up before system disk mirroring cannot be restored back to the mirrored system disk.
- If the system disk mirroring is once canceled and then configured again, you cannot restore the backed up data of the system disk that was collected when the system disk was mirrored previously.

For these reasons, make sure to collect the backed up data of the system disk if the system disk mirroring is configured.

Also make sure to collect the backed up data of the system disk again if the system disk mirroring is configured again.

6.5.1 Checking Partition Configuration of System Disk

Check the partition configuration of the system disk with the parted command and make a note of them. You must perform this step before setting the system disk mirroring.

These details are required for restoration when the system cannot start due to a certain failure.

```
# parted /dev/sda print
Disk /dev/sdf: 34.4GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start   End     Size    File system  Name                Flags
  1      1049kB  630MB  629MB   fat32        EFI System Partition boot, esp
  2      630MB   15.0GB 14.4GB   xfs
  3      15.0GB  17.0GB 2000MB   xfs
  4      17.0GB  19.0GB 2001MB   xfs
  5      19.0GB  21.0GB 2000MB   xfs
  6      21.0GB  23.0GB 2001MB   xfs
  7      23.0GB  25.0GB 2000MB   linux-swap(v1)      swap
```


6.5.2 Setting System Disk Mirroring

6.5.2.1 Setting with GDS Management View

6.5.2.1.1 Operation Flow

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

Figure 6.2 System Disk Settings Operation



Note

The hot spare function cannot be used for the system disk in this version.

6.5.2.1.2 Operation Procedure

This subsection explains how to mirror a system disk.



Note

Disk that will be mirrored by system disk settings

For system disk settings from the GDS Management View, disks with `/`, `/var`, `/usr`, `/boot`, `/boot/efi`, 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`, `/boot`, `/boot/efi`, and a swap area, disks with only swap areas can be excluded from mirroring, but disks with `/`, `/var`, `/usr`, `/boot`, and `/boot/efi` 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.

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

1. Confirming original disks

Figure 6.3 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/fstab, 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 /, /var, /usr, /boot, /boot/efi and swap is displayed in the [Mount Point] field.
- It is impossible to uncheck system disks for which /, /var, /usr, /boot, and /boot/efi are displayed in the [Mount Point] field.

If not mirroring, click <Cancel>.

If mirroring, click <Next>.

2. Creating root class

Figure 6.4 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/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 "2.3.1 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 6.5 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

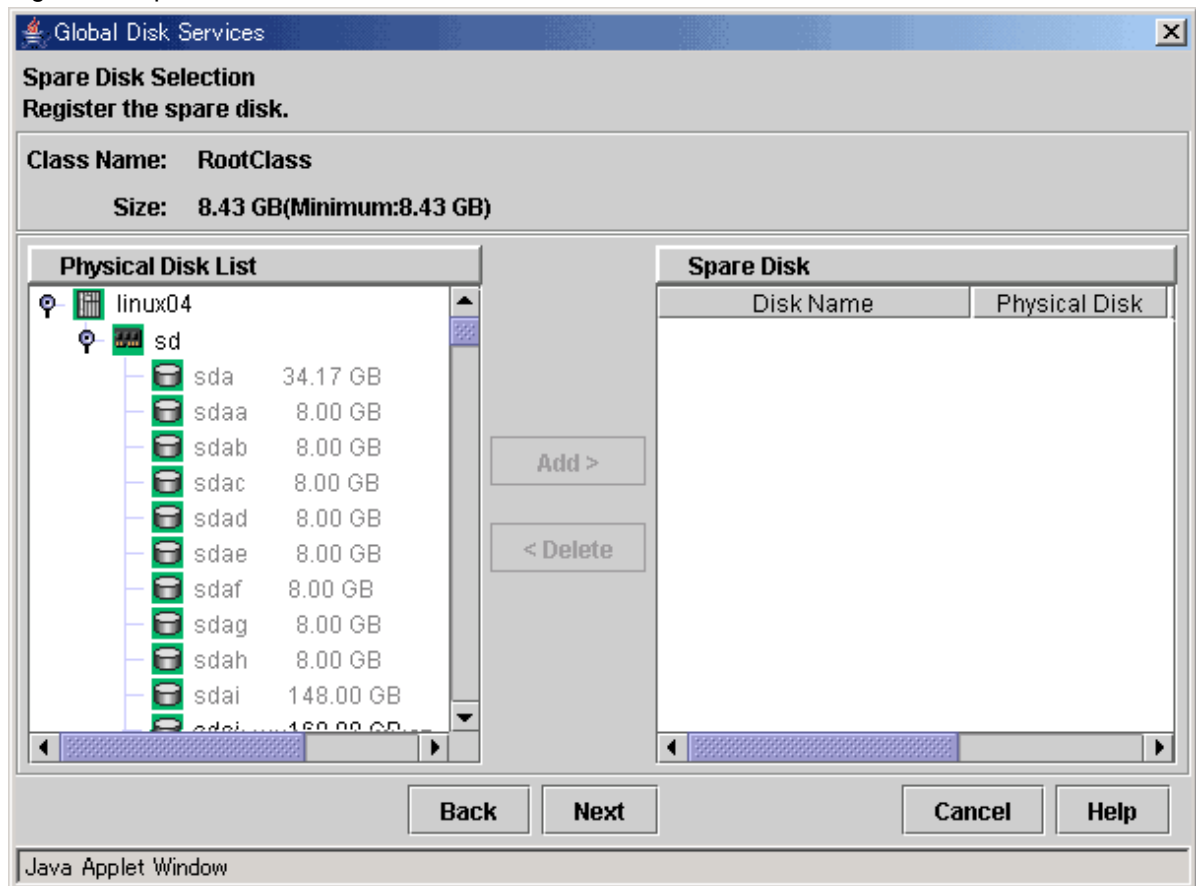
GDS Management's system disk settings do not allow you to add a physical disk capacity that is smaller than the original disk. Users should always add a physical disk that is larger than the original disk.

4. Registering spare disk

Note

The hot spare function cannot be used for the system disk in this version.

Figure 6.6 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>.

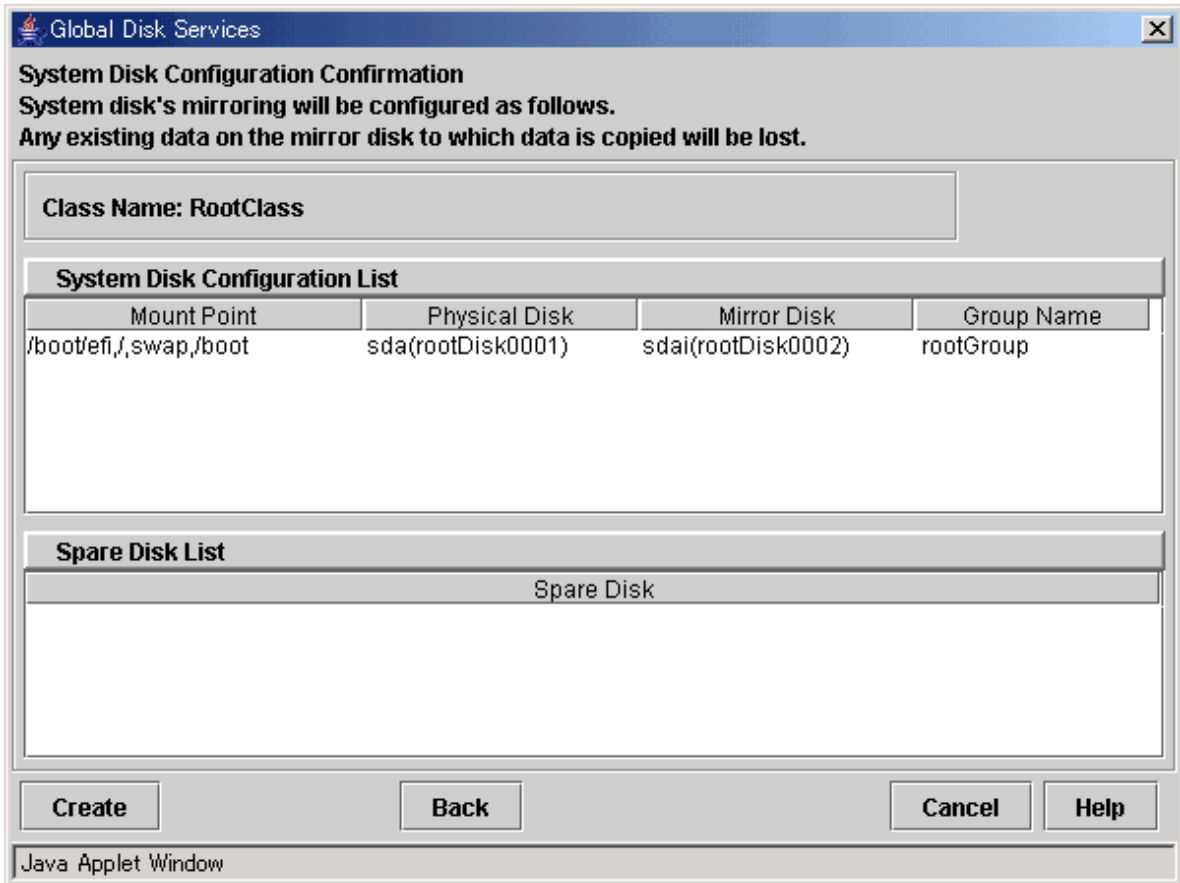
Point

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 6.7 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/fstab, mount point will be displayed in the [Mount Point] field.

Click <Create> to continue. It may take a few minutes for processing.

 **Note**

For RHEL7.4 or later environment

If the grub2-efi-x64-modules package of OS is not installed, the following error message will be displayed and the system disk settings may fail.

5000

Command line output: SDX:sdxroot:ERROR: internal error, sdxfunc=update_initrd_grub

At this time, the following message is output to the log information of the console and main screen.

```
node_name: ERROR: update_initrd_grub: internal error: cmd = '/etc/opt/FJSVsdx/bin/sdxmkinitrd -c RootClass' status = 1
node_name: ERROR: internal error, sdxfunc=update_initrd_grub
```

In this situation, install the grub2-efi-x64-modules package from the OS CD-ROM (DVD), and then re-execute the system disk settings.

Information

Automatically Generated Volume Names

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

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

6. Completing system disk configuration

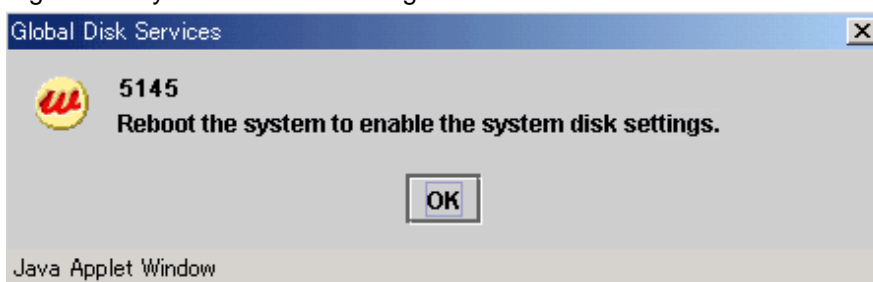
Figure 6.8 Setting System Disk Mirroring Complete



*The file path displayed in the message varies depending on the version of OS.

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

Figure 6.9 System Reboot Message



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

Note

Rebooting the System after System Disk Settings are Completed

After the system disk settings are completed, if you change the system volume names before rebooting the system, the system may not be started. After the system disk settings are completed, reboot the system immediately without performing GDS setup, and so on.

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.

6.5.2.2 Setting with Commands

Assuming the disk on which / (root), /usr, /var, /boot, and /boot/efi 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.

Note

- In this version, to mirror / (root), /usr, /var, /boot, /boot/efi, or swap areas, these areas should be located in the same disk.
- In this version, set the system disk mirroring from the GDS Management View, not by using the command. For details, see "[6.5.2.1 Setting with GDS Management View.](#)"

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, /var, /boot, and /boot/efi is "sda", and the disk allocated as a swap area is "sdb."

```
# sdxdisk -M -c System -a type=root -d sda=Root1:keep,  
sdc=Root2:undef,sdb=Swap1:keep,sdd=Swap2:undef
```

3. Connect the system disks to a group.

```
# sdxdisk -C -c System -g Group1 -d Root1,Root2 -v 1=root:on,  
2=usr:on,3=var:on,4=home:on,5=boot:on,6=efi:on  
# sdxdisk -C -c System -g Group2 -d Swap1,Swap2 -v 1=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 not displayed with the mount(8) command may not be open.

4. Check that mirroring definition is completed.

```
# sdxroot -M -c System -d Root1,Swap1
```

5. Reboot the system.

```
# shutdown -r now
```

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.

6.5.3 Checking System Disk Settings

This section explains how to check that the system disk mirroring is set properly.

[Procedure]

1) Check the GDS configuration.

Execute the `sdxinfo -c` command to check the GDS configuration (class information, disk information, group information, volume information, and slice information).

```
# sdxinfo -c Class_Name
```



For detailed information on contents displayed by the `sdxinfo` command, refer to "[B.1.6 sdxinfo - Display Object Configuration and Status Information.](#)"

(Example)

```
# sdxinfo -c RootClass
OBJ  NAME          TYPE      SCOPE      SPARE
-----
class RootClass root      (local)    0

OBJ  NAME          TYPE      CLASS      GROUP      DEVDNAM  DEVBLKS  DEVCONNECT  STATUS
-----
disk  rootDisk0001  mirror   RootClass  rootGroup  sda      286749488 *           ENABLE
disk  rootDisk0002  mirror   RootClass  rootGroup  sdb      286749488 *           ENABLE

OBJ  NAME          CLASS     DISKS              BLKS      FREEBLKS  SPARE
-----
group rootGroup  RootClass  rootDisk0001:rootDisk0002  286749488  256344880  0

OBJ  NAME          CLASS     GROUP      SKIP  JRM  1STBLK  LASTBLK  BLOCKS  STATUS
-----
volume *          RootClass  rootGroup  *      *      0      2047      2014    FREE
volume efiVolume  RootClass  rootGroup  off    on      2048   411647   409600  ACTIVE
volume bootVolume RootClass  rootGroup  off    on      411648 1435647 1024000 ACTIVE
volume rootVolume  RootClass  rootGroup  off    on      1435648 22407167 20971520 ACTIVE
volume swapVolume  RootClass  rootGroup  off    on      22407168 30795775 8388608 ACTIVE
volume *          RootClass  rootGroup  *      *      30795776 30816255 20480    PRIVATE
volume *          RootClass  rootGroup  *      *      30816256 286749454 255933199 FREE

OBJ  CLASS     GROUP      DISK          VOLUME      STATUS
-----
slice RootClass  rootGroup  rootDisk0001  efiVolume   ACTIVE
slice RootClass  rootGroup  rootDisk0002  efiVolume   ACTIVE
slice RootClass  rootGroup  rootDisk0001  bootVolume  ACTIVE
slice RootClass  rootGroup  rootDisk0002  bootVolume  ACTIVE
slice RootClass  rootGroup  rootDisk0001  rootVolume  ACTIVE
slice RootClass  rootGroup  rootDisk0002  rootVolume  ACTIVE
slice RootClass  rootGroup  rootDisk0001  swapVolume  ACTIVE
slice RootClass  rootGroup  rootDisk0002  swapVolume  ACTIVE
```

2) Check `fstab`.

Execute the following command to check that the device names have been modified.

```
# cat /etc/fstab
```

(Example)

[Before setting up the GDS configuration]

```
# cat /etc/fstab
#
/dev/sda3      /                ext4  defaults                1      1
/dev/sda2      /boot            ext4  defaults                1      2
/dev/sda1      /boot/efi        vfat  umask=0077,shortname=winnt 0      0
/dev/sda4      swap             swap  defaults                0      0
tmpfs          /dev/shm         tmpfs defaults                0      0
devpts         /dev/pts         devpts gid=5,mode=620         0      0
sysfs          /sys             sysfs defaults                0      0
proc           /proc            proc  defaults                0      0
```

[After setting up the GDS configuration]

```
# /etc/fstab
#
/dev/sfdisk/qdssys2  /...
/dev/sfdisk/qdssys5  /boot ...
/dev/sfdisk/qdssys6  /boot/efi ...
/dev/sfdisk/qdssys32 swap             swap  defaults                0      0
tmpfs               /dev/shm         tmpfs defaults                0      0
devpts              /dev/pts         devpts gid=5,mode=620         0      0
sysfs                /sys             sysfs defaults                0      0
proc                /proc            proc  defaults                0      0
```

3) Check the swapon device.

Execute the following command to check that the swap device name has been modified.

```
# swapon -s
```

(Example)

[Before setting up the GDS configuration]

```
# swapon -s
Filename                                Type      Size      Used      Priority
/dev/sda4                               partition 4194296  0         -1
```

[After setting up the GDS configuration]

```
# swapon -s
Filename                                Type      Size      Used      Priority
/dev/sfdisk/qdssys32                    partition 4194296  0         -1
```

6.5.4 Checking Physical Disk Information and Slice Numbers

If system disks have been registered with the root class, check the following details using this procedure and make a note of them.

- System volume physical disk information
- System volume slice numbers

These details are required for performing system disk backup, restore, and recovery from failure.



When Using the System Volume Snapshot Function

Check also the following details.

- Physical disk information on the proxy volumes of system volumes

1) Check the root class name and the system volume names.

```
# mount
/dev/sfdsk/gdssys2 on / ...
/dev/sfdsk/gdssys4 on /var type ...
/dev/sfdsk/gdssys3 on /usr type ...
/dev/sfdsk/gdssys5 on /boot type ...
/dev/sfdsk/gdssys6 on /boot/efi type ...
...
# swapon -s
Filename                Type                ...
/dev/sfdsk/gdssys32     partition          ...

# ls -l /dev/sfdsk/gdssys*
brw-rw---- 1 root disk 231, 2 Jan 5 18:40 /dev/sfdsk/gdssys2
brw-rw---- 1 root disk 231, 3 Jan 5 18:40 /dev/sfdsk/gdssys3
brw-rw---- 1 root disk 231, 32 Jan 5 18:40 /dev/sfdsk/gdssys32
brw-rw---- 1 root disk 231, 4 Jan 5 18:40 /dev/sfdsk/gdssys4
brw-rw---- 1 root disk 231, 5 Jan 5 18:40 /dev/sfdsk/gdssys5
brw-rw---- 1 root disk 231, 6 Jan 5 18:40 /dev/sfdsk/gdssys6

# ls -l /dev/sfdsk/*/dsk/*
brw-r--r-- 1 root root 231, 5 Jan 5 18:41 /dev/sfdsk/System/dsk/bootVolume
brw-r--r-- 1 root root 231, 6 Jan 5 18:41 /dev/sfdsk/System/dsk/efiVolume
brw-r--r-- 1 root root 231, 2 Jan 5 18:41 /dev/sfdsk/System/dsk/rootVolume
brw-r--r-- 1 root root 231, 32 Jan 5 18:41 /dev/sfdsk/System/dsk/swapVolume
brw-r--r-- 1 root root 231, 3 Jan 5 18:41 /dev/sfdsk/System/dsk/usrVolume
brw-r--r-- 1 root root 231, 4 Jan 5 18:41 /dev/sfdsk/System/dsk/varVolume
```

Find the appropriate device with the same major number and the same minor number. The major or minor number is displayed in each device.

For example, in the case of /dev/sfdsk/gdssys2, the major number is 231 and the minor number is 2. Therefore, /dev/sfdsk/RootClass/dsk/rootVolume is the appropriate device.

Based on this correspondence, match the use and the volume of devices.

In this example, the root class name is System, and the system volume names are as follows:

Use	Volume name
/	rootVolume
/var	varVolume
/usr	usrVolume
/boot	bootVolume
/boot/efi	efiVolume
Swap area	swapVolume

2) Check the group names and slice numbers of the system volumes.

```
# sdxinfo -v -c System -e long
OBJ    NAME          TYPE    CLASS  GROUP  ...  SNUM  PJRM
-----
volume rootVolume  mirror System Group1 ...  1    *
volume varVolume  mirror System Group1 ...  2    *
volume usrVolume  mirror System Group1 ...  3    *
```

```

volume bootVolume mirror System Group1 ... 4 *
volume efiVolume mirror System Group1 ... 5 *
volume swapVolume mirror System Group1 ... 6 *
...

```

For the -c option, specify the root class name confirmed in step 1).

The group names are displayed in the GROUP fields. In this example, the group name is Group1.

The slice numbers are displayed in the SNUM fields. In this example, the slice numbers are as follows.

Use	Volume name	Slice number
/	rootVolume	1
/var	varVolume	2
/usr	usrVolume	3
/boot	bootVolume	4
/boot/efi	efiVolume	5
Swap area	swapVolume	6



Note

When Using the System Volume Snapshot Function

If the sdxinio command is executed as above, information on the proxy volumes of the system volumes is given additionally. Check also the group name of the proxy volumes. If the proxy is joined through group operation, the proxy volume slice numbers are the same as those of the corresponding system volume slice numbers.

3) Check the SDX disk names of disks composing the system volumes.

```

# sdxinio -G -c System
OBJ  NAME      CLASS   DISKS      ...
-----
group Group1 System  Root1:Root2 ...

```

For the -c option, specify the root class name confirmed in step 1).

Check the DISKS field in the line showing the group name confirmed in step 2) in its NAME field.

In this example, the SDX disk names are Root1 and Root2.



Note

When Using the System Volume Snapshot Function

If the sdxinio command is executed as above, proxy volume group information is given additionally. Check also the SDX disk names of disks composing the proxy volume group.

4) Check the physical disk names of disks composing the system volumes.

```

# sdxinio -D -c System
OBJ  NAME      TYPE    CLASS   GROUP  DEVNAM  ...
-----
disk  Root1     mirror System Group1  sda    ...
disk  Root2     mirror System Group1  sdb    ...

```

For the -c option, specify the root class name confirmed in step 1).

The physical disk names are displayed in the DEVNAM fields.

In this example, the physical disk names are as follows.

SDX disk name	Physical disk name
Root1	sda
Root2	sdb

Note

When Using the System Volume Snapshot Function

If the `sdxinfo` command is executed as above, information on disks composing the proxy volumes is given additionally. Check also the physical disk names of those disks.

5) Check information on the physical disks composing the system volumes.

```
# readlink -f /sys/block/sda/device
/sys/devices/pci0000:02/0000:02:1f.0/0000:06:02.0/host0/target0:0:0/0:0:0:0
# readlink -f /sys/block/sdb/device
/sys/devices/pci0000:02/0000:02:1f.0/0000:06:02.0/host0/target0:0:2/0:0:2:0
```

For the `readlink` command arguments, specify `/sys/block/physical_disk_name/device`.

Among symbolic link destination paths displayed, elements before `hostX` and the last `X:Y:Z` is physical disk information.

In this example, it is as follows.

SDX disk name	Physical disk name	Physical disk information	
Root1	sda	0000:06:02.0	0:0:0
Root2	sdb	0000:06:02.0	0:2:0

Note

When Using the System Volume Snapshot Function

Using a similar method, check also physical disk information on the proxy volumes of the system volumes.

6.5.5 Collecting Backups of System Disk

Collect backups of the system disk.

See

For details on the backup method of the system disk, see "7.4 Backing Up and Restoring System Disk."

6.5.6 Backing up EFI Configuration Information

Back up EFI configuration information.

See

For details on the backup method, see the manuals for PRIMEQUEST.

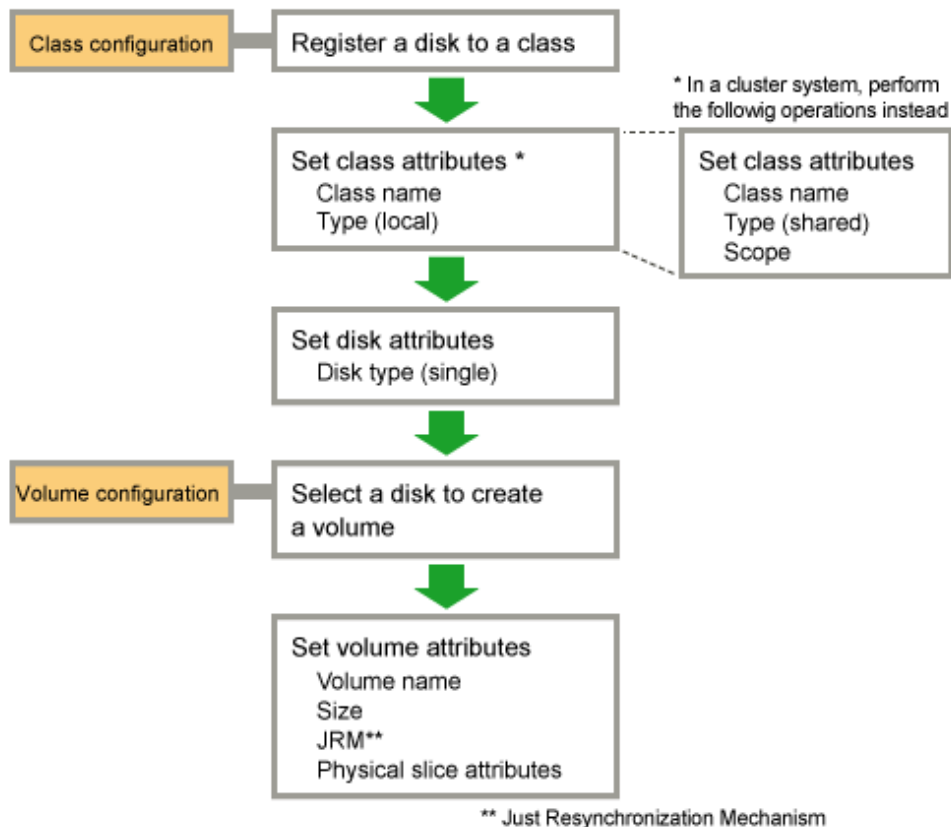
6.6 Class, Group, Volume Settings

6.6.1 Setting with GDS Management View

6.6.1.1 Single Volume Configuration Settings

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

Figure 6.10 Single Volume Configuration Settings Operation



 See

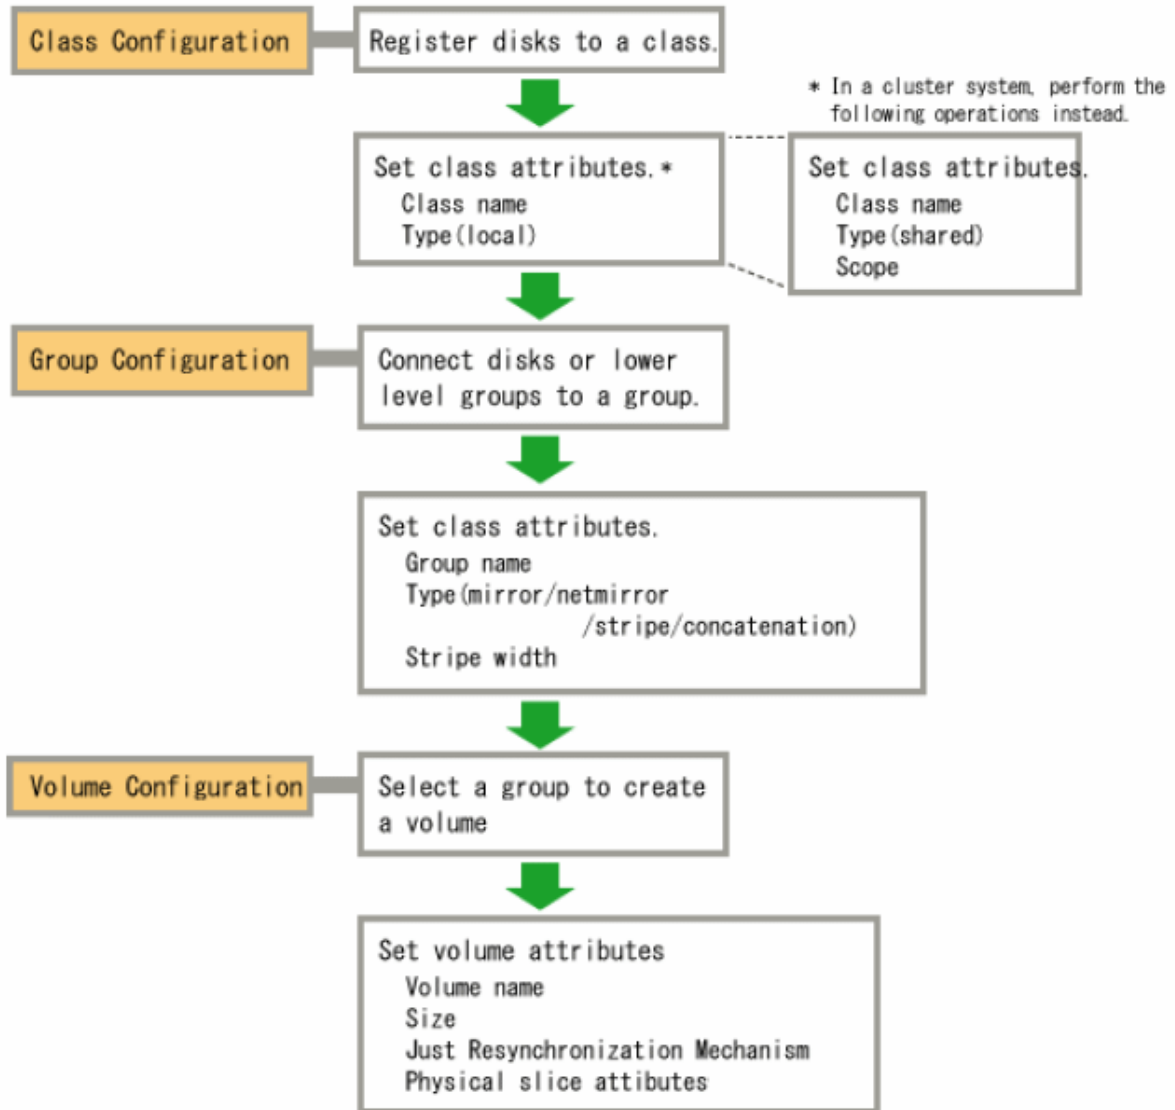
For details on the operation methods, see "[6.6.1.3 Class Configuration](#)", "[6.6.1.4 Cluster System Class Configuration](#)", and "[6.6.1.6 Volume Configuration](#)."

6.6.1.2 Other Volume Configuration Settings

The operation outline of the configuration settings for volumes other than single volume (mirror volume, netmirror 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 6.11 Volume Configuration Settings Operation



See


For details on the operation methods, see "6.6.1.3 Class Configuration", "6.6.1.4 Cluster System Class Configuration", "6.6.1.5 Group Configuration", and "6.6.1.6 Volume Configuration."

6.6.1.3 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 6.12 Class Configuration



1. Selecting a class name

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

2. Selecting a 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.

Note

If all physical disks registered with the class are not the same size, register the largest size disk first. For details, see ["2.3.9 Disk Size."](#)

Figure 6.13 Selecting Physical Disk to Configure Class



Information

Up to 400 physical disks can be selected by one-time operation.

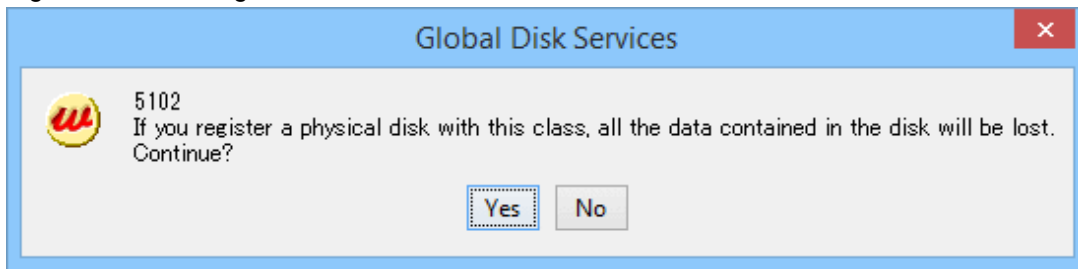
Point

When building the mirroring among servers in the class, select the iSCSI device names confirmed in ["4.8.5 Confirming the Correspondence of iSCSI Device Name of Each Node."](#)

3. Creating a class

Click <Add>, and the message screen below appears.

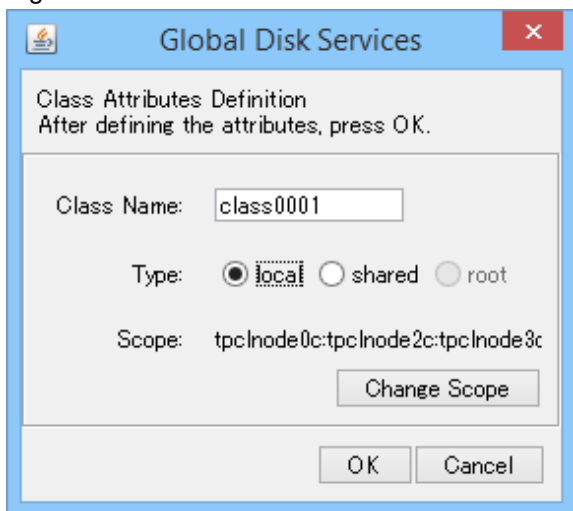
Figure 6.14 Warning



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

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

Figure 6.15 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>.

Information

Inputting Class Name

Class name will be used for the device path 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 "2.3.1 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(8) 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 "[6.6.1.4 Cluster System Class Configuration](#)."

6.6.1.4 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 6.16 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

Node Name

A node identifier of PRIMECLUSTER is displayed in the [Node Name] of the [Change Scope] screen and a host name (mip) of Web-Based Admin View is displayed in the GDS Configuration Tree field of the Main screen.

A node identifier and a host name may be different. In this case, the both names are displayed in the form of "node identifier (host name)" in the [Node Name].

Note

Class Resource

- 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(8) command.
- Cluster applications that use resources of a class should be set after the volume configuration is complete.

See

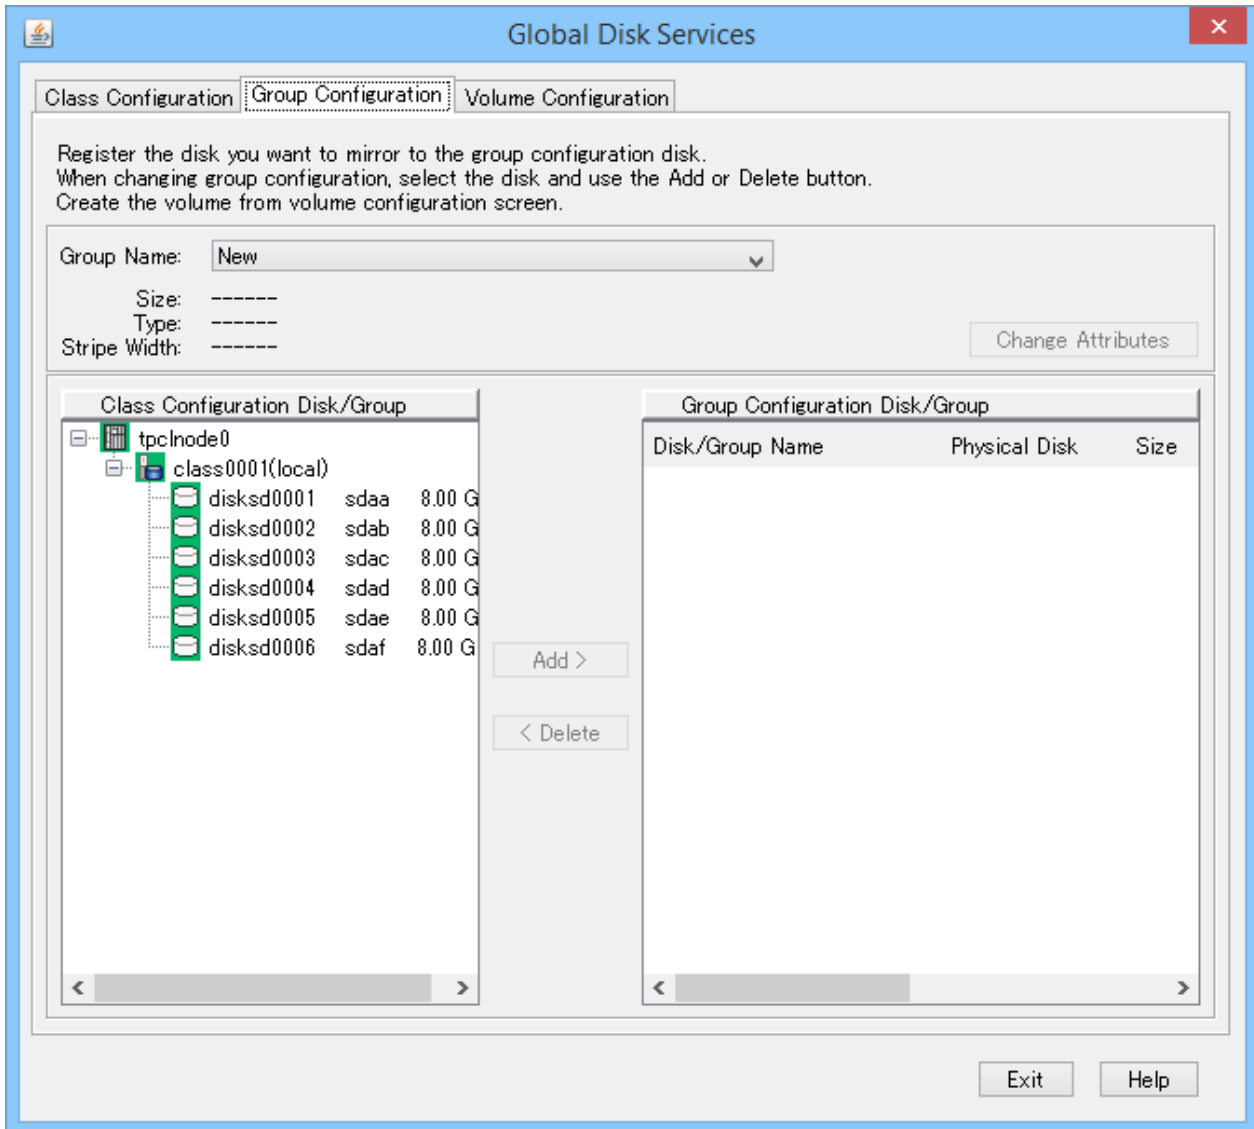
For information on installation and initial setting of cluster systems, refer to "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide."

6.6.1.5 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 6.17 Group Configuration



Follow the procedures below to create a new group.

1. Selecting group name

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

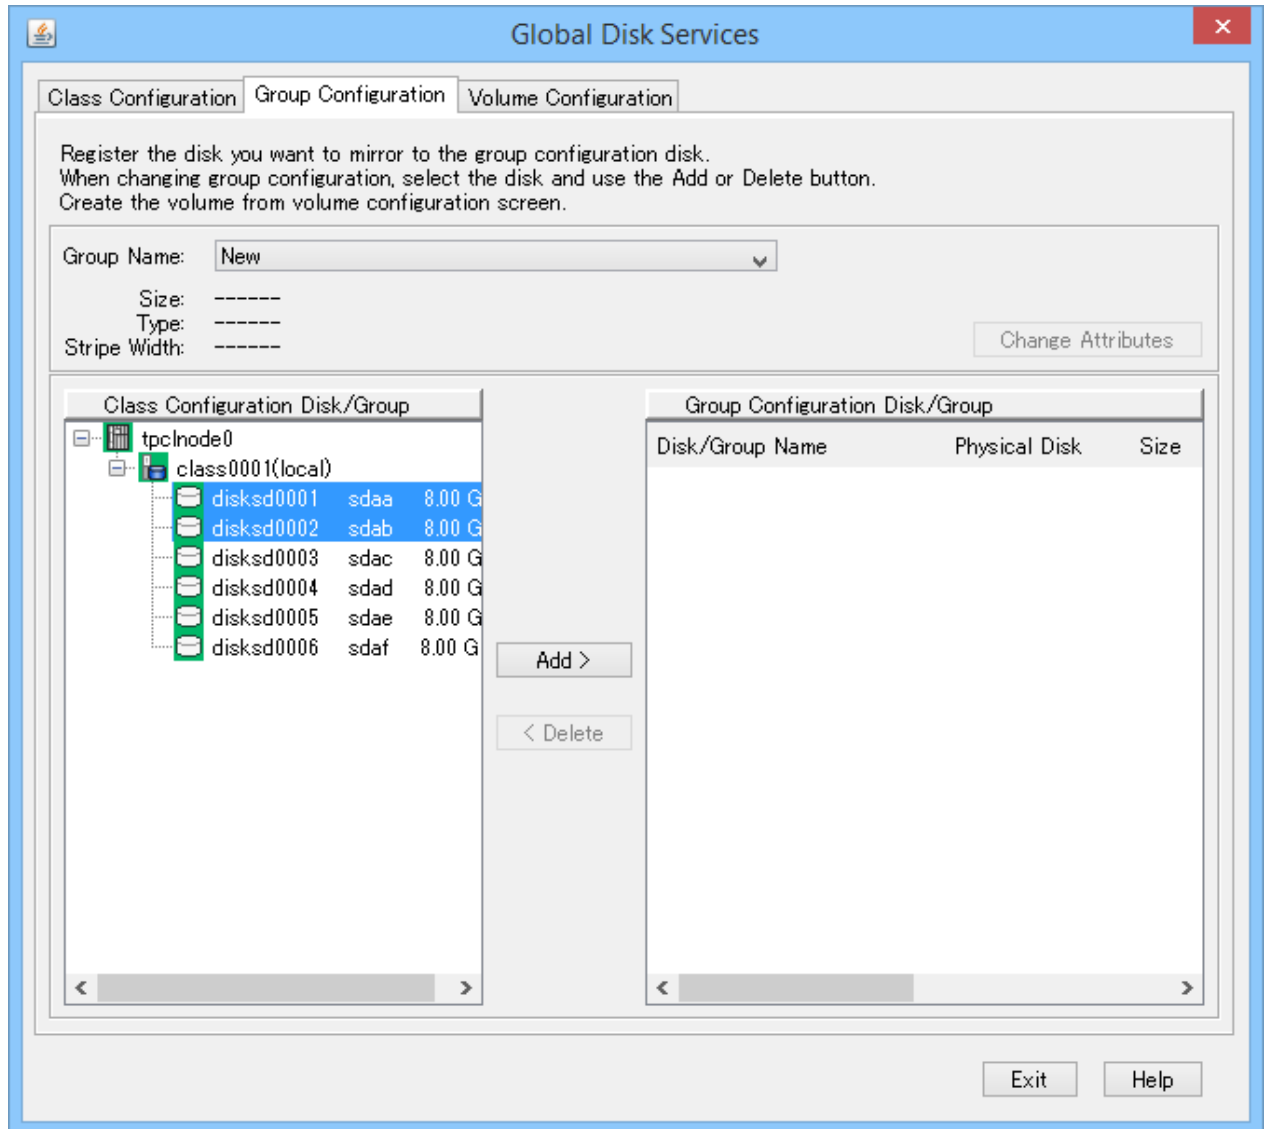
2. Selecting disk/lower level group

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.

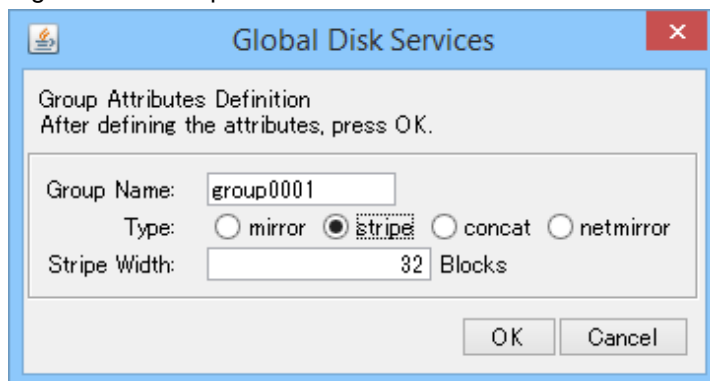
Figure 6.18 Selecting Disk/Lower level group to Configure Group



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 6.19 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 the following types. The default setting is "mirror."

- For mirroring: mirror
- For striping: stripe
- For concatenating: concat
- For mirroring among servers: netmirror

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 "2.3.1 Object Name."
.....



Note

.....
Stripe groups and concatenation groups cannot be created in RHEL8.
.....

4. **Completing group creation**


After creating all groups, click <Exit> and close Group Configuration screen.

6.6.1.6 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 6.20 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.

 **Information**

Disk Size

The [Disk Size] displayed below the volume figure, displays available size of the group or disk selected in the [Group and Disk List].

Available size is the size of an area (including the area where volume has already been created) that can be used as volume area.

For details about available size of a group, see "2.3.10 Group Size."

For details about available size of a disk, see "2.3.9 Disk Size."

2. **Selecting unused volume**

Click the <Unused> field with a volume icon, and select an unused volume.

 **Note**

Size displayed in the <Unused> field

The size displayed in the <Unused> field is the maximum size that can be created as a single volume.

If any volume has ever been deleted in the relevant group or on the relevant disk, the sum of the sizes displayed in the <Unused> fields and that of the sizes of the free spaces may not match.

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 "[2.3.1 Object Name](#)."

b. Setting [Volume Size]

Type the volume size in MB units, using numbers only.



Note

If Volume Size Does Not Match Cylinder Boundary

- When the specified volume size does not match the cylinder boundary of the disk, it will be automatically adjusted by rounding up. For details, see "[2.3.11 Volume Size](#)."
 - Since volume sizes are adjusted to fit into the cylinder boundary, volume sizes that can be created may be smaller than the value displayed in [Disk Size].
-

c. [Maximum Size] button

Sets the value in the [Volume Size] field to the maximum available size.

d. Setting [JRM]

Change the setting when you want to disable the just resynchronization feature. Default is set to "on."

If you have selected a stripe group or concatenation group in step 1., the setting cannot be changed.

If you have selected the netmirror group in step 1., do not change the setting to "off." The volume cannot be created if the setting is "off".

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 6.21 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/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.

6.6.2 Setting with Commands

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.

6.6.2.1 Creating Single Volume

The following example shows the procedures for creating a single volume using a physical disk named sda.

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 sda to shared class "Class1", which is shared on node1 and node2, and assigning the name "Disk1."

```
# sdxdisk -M -c Class1 -a type=shared,scope=node1:node2 -d sda=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
```

6.6.2.2 Creating Mirror Volume

The following example shows the procedures for creating a volume by mirroring physical disks named sda and sdb.

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 sda and sdb with local class "Class1", and name these disks Disk1 and Disk2.

```
# sdxdisk -M -a type=local -c Class1 -d sda=Disk1,sdb=Disk2
```



When registering disks with a shared class, it is necessary to use the -a option to specify the scope attribute. See an example below.

```
# sdxdisk -M -c Class1 -a type=shared,scope=node1:node2 -d  
sda=Disk1,sdb=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 by using the `sdxinfo -S` command.

If 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 "[B.1.6 sdxinfo - Display Object Configuration and Status Information.](#)"

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
```

6.6.2.3 Creating Netmirror Volume

The following example shows the procedures for creating a volume mirroring the iSCSI devices "sda" and "sdb" among servers.

Note that "sda" and "sdb" must be created according to "[4.8 Disk Setting for Performing Mirroring among Servers.](#)"

1) Registering disks to a class

Register the iSCSI devices with a shared class. When the specified shared class does not exist, it will be created automatically.

Example) Register iSCSI devices "sda" and "sdb" with the shared class "Class1", and name these devices "Disk1" and "Disk2."

```
# sdxdisk -M -c Class1 -a type=shared,scope=node1:node2 -d sda=Disk1,sdb=Disk2
```

2) Connecting the disks to a netmirror group

Connect the disks to a netmirror group. When the specified netmirror group does not exist, it will be created automatically.

Example) Connecting "Disk1" and "Disk2" to the netmirror group "Group1."

```
# sdxdisk -C -c Class1 -g Group1 -d Disk1,Disk2 -a type=netmirror
```

3) Creating a netmirror volume

Create a volume within the netmirror group.

Example) Creating a volume of 1,000 blocks within the netmirror group "Group1," and assigning a volume name "Volume1."

```
# sdxvolume -M -c Class1 -g Group1 -v Volume1 -s 1000
```

After returning from the command, the synchronization copying will automatically begin.

4) Confirming the completion of procedure

Confirm that the synchronization copying is complete by using the `sdxinfo -S` command.

If 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 "[B.1.6 sdxinfo - Display Object Configuration and Status Information.](#)"

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
```

6.6.2.4 Creating Large-Capacity Volume Using Concatenation

The following example shows the procedures for creating a volume by concatenating physical disks named `sda` and `sdb`.

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 `sda` and `sdb` with shared class "Class1", which is shared on `node1` and `node2`, and naming those disks `Disk1` and `Disk2`.

```
# sdxdisk -M -c Class1 -a type=shared,scope=node1:node2 -d
sda=Disk1,sdb=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
```



Concatenation groups cannot be created in RHEL8.

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
```

6.6.2.5 Creating Stripe Volume

The following example shows the procedures for creating a volume by striping physical disks named `sda` and `sdb`.

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 sda and sdb with shared class "Class1", which is shared on node1 and node2, and naming those disks Disk1 and Disk2.

```
# sdxdisk -M -c Class1 -a type=shared,scope=node1:node2 -d
sda=Disk1,sdb=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
```



Note

Stripe groups cannot be created in RHEL8.

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
```

6.6.2.6 Creating Mirror Volume by Combining Striping and Mirroring

The following example shows the procedures for creating a volume by constructing stripe groups with physical disks sda and sdb and other physical disks sdc and sdd respectively and then mirroring the 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 sda, sdb, sdc and sdd to shared class "Class1", which is shared on node1 and node2, and assigning the names "Disk1", "Disk2", "Disk3" and "Disk4" respectively.

```
# sdxdisk -M -c Class1 -a type=shared,scope=node1:node2 \
-d sda=Disk1,sdb=Disk2,sdc=Disk3,sdd=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
```

Note

Stripe groups cannot be created in RHEL8.

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 by using the `sdxinfo -S` command.

If 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 ["B.1.6 sdxinfo - Display Object Configuration and Status Information."](#)

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
```

6.7 File System Configuration

6.7.1 Creating the File System

This section describes the procedure to create the file system in a volume. If the volume was not started, after starting the volume, perform the following procedure. Create the file system by using the commands of OS or the file system. As the device to create the file system, specify the device file of the volume `/dev/sfdsk/Class_Name/dsk/Volume_Name`.

Below is the example when creating the ext4 file system in the volume (class name: class1, volume name, volume1).

```
# mkfs -t ext4 /dev/sfdsk/class1/dsk/volume1
```



See

For details on how to create the file system, see the manual of the file system.

6.7.2 File System Auto Mount

File systems created on volumes in classes other than the root class (local classes and shared classes) cannot be mounted with the OS auto mount feature at OS startup.

This is because the OS auto mount process is executed before the GDS startup script is executed.

When including file systems created on volumes in classes other than the root class in the `/etc/fstab` file, be sure to describe the "noauto" option in the fourth field.

If the noauto option is not described, the following message is displayed at OS startup, and OS is started in emergency mode.

- For RHEL7

```
[ TIME ]Dependency failed for device dev-sfdsk-class_name-dsk-volume_name.device.
```

- For RHEL8 or later

```
[ TIME ]Dependency failed for Local File Systems.
```



Information

For Shared Classes

To mount or unmount file systems on volumes in shared classes when starting or exiting cluster applications, `/etc/fstab.pcl` file setting and Fsystem resource setting are necessary. When using GFS, see "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide." For details, see "PRIMECLUSTER Global File Services Configuration and Administration Guide."



Information

For Local Classes

To automatically mount or unmount file systems created on volumes in local classes on starting or exiting OS, create and set the script and unit files in order to satisfy the following conditions:

- Execute the mounting process after all of the GDS service (`fjvsdx*.service`) when starting OS.
- Execute the unmounting process before all of the GDS service (`fjvsdx*.service`) when stopping OS.



See

For details about GDS services, see "F.2 systemd Service Lists."

Similarly, to export file systems created on volumes in local classes to an NFS client, create and set the script and unit files in order to satisfy the following conditions:

[RHEL7]

- Execute the mounting process before `nfs.service` when starting OS.
- Execute the unmounting process after `nfs.service` when stopping OS.

[RHEL8]

- Execute the mounting process before `nfs-server.service` when starting OS.

- Execute the unmounting process after `nfs-server.service` when stopping OS.



For the startup script creation method, see "6.9 Volume's Block Special File Access Permission."

6.7.3 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 Installation and Administration Guide."

6.8 raw Device Binding

For using a volume or a temporarily detached slice as a character (raw) device, use the raw device bound to the volume's or the temporarily detached slice's block device with the `raw(8)` command.

Example) Bind raw device "raw1" to the block device of Volume1 in Class1.

```
# raw /dev/raw/raw1 /dev/sfdsk/Class1/dsk/Volume1
```

By describing the setting for raw device bounding in the `/etc/sysconfig/rawdevices` file, the raw device can automatically be bound at OS startup. For details, see the `raw(8)` and `rawdevices` manuals.



The setting file for binding raw devices is `/usr/lib/udev/rules.d/60-raw.rules`. However, when using a volume or a temporarily detached slice as a raw device, describe the setting to the `/etc/sysconfig/rawdevices` file.

When changing the access permission for the character device special file (`/dev/raw/raw<N>`) of the raw device, create a startup script in which the command that changes the access permission is described, and create a unit file that executes the startup script so that the script is executed later than the GDS service (`fjvsdx*.service`).

If the setting is described in the `rawdevices` file to bind the raw device for the volume or for the temporarily detached slice, the following message may be output when OS is started. Even this message is output, the raw device is normally bound after GDS is initialized. The system is not affected.

```
Cannot locate block device '/dev/sfdsk/Class_Name/dsk/Volume_Name' (No such file or directory)
```

After the raw device binding, if the relevant volume or temporarily detached slice is re-created, it is necessary to bind the raw device again.



The OS may or may not support raw devices or the `raw(8)` command. For details, see the OS manual.

6.9 Volume's Block Special File Access Permission

The block special file `/dev/sfdsk/Class_Name/dsk/Volume_Name` of the volume is created when a volume is created and re-created every time the node is rebooted.

The volume's block special file access permission is set as follows.

- Owner: root

- Group: root
- Mode: 0600

When changing the access permission, create a startup script in which the command that changes the access permission is described, and create a unit file that executes the startup script so that the script is executed later than all of the GDS services (fjvsdx*.service).



See

For details about GDS services, see "F.2 systemd Service Lists."

Descriptions in the Startup Script for Access Permission Setting

An example of script description for RHEL is shown below.

```
#!/bin/bash

# description: chgperm - change GDS volume permission      ...(1)

. /etc/init.d/functions

start() {
    /bin/chown gdsusr:gdsgrp /dev/sfdsk/Class1/dsk/Volume1  ...(2)
    /bin/chmod 0644 /dev/sfdsk/Class1/dsk/Volume1          ...(2)
    return
}

stop() {
    return
}

case "$1" in
    start)
        start
        ;;
    stop)
        stop
        ;;
    restart)
        stop
        start
        ;;
    *)
        echo "Usage: /etc/init.d/chgperm {start|stop|restart}"  ...(3)
        ;;
esac
exit 0
```

At GDS installation, a sample of this startup script is installed to /etc/opt/FJVSdx/etc/chgperm.sample. Copy and customize this sample before using it.

Explanation of the Descriptions:

(1)

To the right of "description:," give a description of this startup script.

(2)

Describe commands executed at node startup.

The commands shown in this example set the owner to gdsusr, the group to gdsgrp, and the mode to 0644 with respect to the block special file of Volume1 in Class1.

(3)

Describe the process to show the usage of this startup script.
The startup script name in this example is chgperm.



See

- For details on (1), see chkconfig(8).
- For details on the commands to set the file access permission, see chown(1) and chmod(1).

Descriptions of the unit file

See the example below.

```
[Unit]
Description=change GDS volume permission          ... (1)
After=fjvsdxmon.service                          ... (2)

[Service]
ExecStart=/bin/bash -c '/home/bin/chgperm start' ... (3)
Type=oneshot

[Install]
WantedBy=multi-user.target
```

Explanation of the Descriptions:

(1)

To the right of Description=, describe the description of this unit file.

(2)

Describe the script to be started later then the GDS service "fjvsdxmon.service."

(3)

Describe the start script that changes the access permission (/home/bin/chgperm in this example).

How to Set the Startup Script for Access Permission Setting

Configure the following setting.

1. Locate the unit file in the path /usr/lib/systemd/system/*unit_file_name*.
2. Execute the following command to enable the service.

```
# /usr/bin/systemctl enable unit_file_name
```



Information

By executing the command in the above procedure 2, the following symbolic link file is created.

```
/etc/systemd/system/multi-user.target.wants/unit_file_name
```

6.10 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.

6.10.1 Setting with GDS Management View

6.10.1.1 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 joined in the Main Screen.


 **Information**

If the number of volumes in a group is more than 400, you cannot select the group as the target to join the proxy group.

2. **Selecting the [Join] menu**

Select [Operation]:[Proxy Operation]:[Join] in the Main screen.

Figure 6.22 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 6.23 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 6.24 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 "2.3.1 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.4 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 6.25 Confirming "Join" (Volumes)



To continue the process, click <Yes>. Clicking <No> displays the Select Proxy dialog box shown in step 3. again.

Figure 6.26 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 6.27 Information Message of "Join" Completion



Click <OK> to close the information message window.

6.10.1.2 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 mirror volume, a single volume, or a mirror 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 6.28 Relate



3. Selecting a proxy to be related

The Select Proxy dialog box appears.

Figure 6.29 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 6.30 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 6.31 Information Message of "Relate" Completion



Click <OK> to close the information message window.

6.10.2 Setting with Commands

6.10.2.1 Join

Join proxy volumes or proxy groups by using the sdxproxy Join command.



See

- For details on the sdxproxy command, see "[B.2.1 sdxproxy - Proxy Object Operations.](#)"
- In the description of how to back up by using the proxy volume, the procedure to join proxy by using the sdxproxy Join command is explained. For details, see the following sections:
 - "[7.5 Backing Up and Restoring System Disk through Alternative Boot Environment \[EFI\]](#)"
 - "[7.8.2 Using Commands](#)" in "[7.8 Backup \(by Synchronization\)](#)"

6.10.2.2 Relate

Relate proxy volumes or proxy groups by using the sdxproxy Relate command.



See

- For details on the `sdxproxy` command, see "B.2.1 `sdxproxy` - Proxy Object Operations."
- In the description of how to back up by using the proxy volume, the procedure to relate proxy by using the `sdxproxy Relate` command is explained. For details, see the following section:
 - "7.9.2 Using Commands" in "7.9 Backup (by OPC)"

6.11 Cloning

Cloning is to create another system by copying the existing system.

This section explains how to clone the existing system on which GDS is installed and set.



See

In Cluster Systems

For how to clone the cluster system, see PRIMECLUSTER Installation and Administration Guide.

6.11.1 Summary

See below for the procedure of cloning.

1. Preparation

On the original system, prepare for cloning such as collecting the configuration information of GDS.

2. Copying the disk data

Copy the disk data of the original system to the disk of the copy destination system.

Copy the data by using software such as SystemcastWizard, which provides the cloning function, or by using the function to copy the disk data such as the copy function of storage.

3. Setting up the copy destination system

On the copy destination system, change the GDS configuration depending on the hardware configuration.

If the system disk is mirrored on the copy source system, you need to cancel the system disk mirroring once on one of the following systems; the copy source system or the copy destination system.

If you do not want to cancel the system disk mirroring on the copy source system, clone the system disk with the following procedure "a."

For other cases, clone the system disk with the procedure "b."

In particular, it is recommended to take "b." if there are a great number of systems on the copy destination system.

a. Clone the system disk while mirroring it and then cancel and reconfigure the system disk mirroring on the copy destination system.

In step 3, start the system from an installation CD of OS, and then delete the configuration information of the system disk mirroring on the copy destination system. After that, start the system from the system disk to set up the system disk mirroring.

b. Clone the system disk after canceling the system disk mirroring once.

In step 1, cancel the system disk mirroring on the original system.

After step 2, mirror the system disk again on the original system.

In step 3, set up the system disk mirroring on the copy destination system.

Take either of the following procedures to clone the disk of the local class.

- Copying the whole disk data including the private slice.
- Copying only the data of the private slice.

- Copying only the data in volume areas.
- Not copying the disk data.

Determine the copy range by the software and the specification of the copy function (which areas that can be copied) to use the data copy and by the necessity of copying data of volume areas.

6.11.2 Preconditions

When copying the system by cloning, the following precondition is required.

- Capacity of the disk to be registered in the local class should be the same in the original system and the copy destination system.

6.11.3 Preparation

Take the following procedure on the original system.

This section explains the example when the class name is Class1.

1. Backing up object configuration of local classes

Back up object configuration of all local classes.

```
# sdxconfig Backup -c Class1 -o /var/tmp/Class1.conf
```

2. Saving configuration information of local classes

When copying the data of the local disk, save the configuration information of the target local classes.

```
# sdxinfo -c Class1 -e long > /var/tmp/Class1.info
```

3. Canceling the system disk mirroring

When cloning the system after canceling the system disk mirroring once, cancel the mirroring of the system disk.

For how to cancel the system disk mirroring, see "[9.4 System Disk Mirroring Cancellation \[EFI\]](#)."

6.11.4 Copying Disk Data

1. Copying the disk data

Copy the disk data of the original system to the disk of the copy destination system.

Copy the data by using software such as SystemcastWizard, which provides the cloning function, or by using the function to copy the disk data such as the copy function of storage.

For details, see the manuals of these functions.

2. Setting up the system disk mirroring

If the system disk mirroring is canceled in "[6.11.3 Preparation](#)", mirror the system disk again on the original system.

For how to set up, see "[6.5 System Disk Mirroring Settings \[EFI\]](#)."

6.11.5 Setting up Copy Destination System

Take the following procedure on the copy destination system.

1. Deleting the configuration information of the system disk mirroring

When copying the system disk data while mirroring the system disk, start the system from an installation CD of OS to delete the configuration information of the system disk mirroring.

For how to delete, see "Resolution" of "[\(4\) System cannot be booted. \(Failure of all boot disks\)](#)" in "[D.1.5 System Disk Abnormality \[EFI\]](#)" In this step, you do not need to replace the physical disk and restore the backup data. Additionally, when starting the system after exiting the rescue mode, start the system in single user mode.

2. Start the system in single user mode.

Because the console cannot be used in a cloud environment (except an FJcloud-Baremetal environment), do not select single user mode. In this case, take the following procedure to restrict the automatic startup of the PRIMECLUSTER service.

1. Restrict the automatic startup of the PRIMECLUSTER service.

Restrict the automatic startup of the PRIMECLUSTER service by executing the following command on all nodes.

```
# /opt/FJSVpclinst/bin/pclservice off
```

2. Restart the system.

Restart the system on all nodes.

```
# /sbin/shutdown -r now
```

3. Deleting the configuration information of the local class

Take the following procedure on all nodes.

This section explains the example when the class name is Class1.

1. Delete the configuration of the class name from the /etc/opt/FJSVsdx/sysdb.d/class.db file.

```
# cd /etc/opt/FJSVsdx/sysdb.d
# vim class.db
Class1    <-Delete all the entries of the class name
...
# Disk Class List
```

2. Delete all the files of the class name under the /etc/opt/FJSVsdx/sysdb.d directory.

```
# rm Class1
```

3. Delete the configuration information of the class from the /etc/sysconfig/devlabel file.

```
# cd /etc/sysconfig
# vim devlabel
...
/etc/opt/FJSVsdx/.devlabel/Class1/sdx_dev...    <- Delete the class
configuration information
```

4. Delete all the directories of the class name under the /etc/opt/FJSVsdx/.devlabel directory.

```
# cd /etc/opt/FJSVsdx/.devlabel
# rm -rf Class1
```

4. Deleting the management information of the physical disk

```
# cat /dev/null > /etc/opt/FJSVsdx/sdx.udev
```

5. Correcting the excluded device list

If the physical disk name differs between the original system and the copy destination system, correct the physical disk name of the disk that is excluded from GDS management described in the excluded device list (/etc/opt/FJSVsdx/lib/exdevtab file) to be the same as the physical disk name in the copy destination system.

For details on the excluded device list, see "[6.1 Creating Excluded Device List.](#)"

6. Restart the system in multi-user mode.

In a cloud environment (except an FJcloud-Baremetal environment), take the following procedure to cancel restricting the automatic startup of the PRIMECLUSTER service.

1. Configure the automatic startup of the PRIMECLUSTER service.

Execute the following command on all nodes and change the PRIMECLUSTER service settings back to the state they were in before they were restricted in step 2.

```
# /opt/FJSVpclinst/bin/pclservice on
```

2. Restart the system.

Restart the system on all nodes.

```
# /sbin/shutdown -r now
```

From this procedure, the example shows that the class name is Class1, and the configuration file path is the /var/tmp/Class1.conf.

7. Changing the physical disk name in the configuration file

If the physical disk name to be registered in the local class differs between the original system and the copy destination system, change the physical disk name in the configuration file to the physical disk name of the copy destination system.

Here is the example when the physical disk name of the original system is sdf, and the physical disk name of the copy destination system is sdh.

```
# sdxconfig Convert -e replace -c Class1 -p sdf=sdh -i /var/tmp/Class1.conf -o /var/tmp/Class1.conf -e update
```

8. Restoring the object configuration information in the local class

1. When copying the data of the private slice

```
# sdxconfig Restore -c Class1 -i /var/tmp/Class1.conf -e chkps
```

2. When not copying the data of the private slice

```
# sdxconfig Restore -c Class1 -i /var/tmp/Class1.conf
```

9. Restart the system.

10. Setting up the system disk mirroring

Set up the system disk mirroring.

For how to set up, see "[6.5 System Disk Mirroring Settings \[EFI\]](#)."

6.12 I/O Response Time Assurance [IOmonitor]

Here the I/O Response Time Assurance Function provided as an option product "GDS I/O Monitor Option" for GDS, will be explained.

6.12.1 Setting I/O Response Time Assurance

Here the method of setting the I/O Response Time Assurance Function for an already created class will be explained.

1. Set the I/O Response Time Assurance.

Example) If setting the I/O response time to 77 seconds and the system panic mode to on for the class Class.

```
# sdxattr -C -c Class -a iotimeout=on:77:on
```

2. Check the settings of the I/O Response Time Assurance.

Example) If checking the settings of the I/O Response Time Assurance of the class Class.

```
# sdxinfo -I -c Class
CLASS  MODE TIME PANIC
```

```
-----
Class      on      77      on
```

If the value set in procedure 1 is displayed in the MODE field, the TIME field and the PANIC field of the specified class the setup of the I/O Time Response Assurance is complete.

3. Change the settings of the Hang-up Monitoring Function if a value other than the default value was set for the I/O Response Time. For details, see "[6.12.2 Hang-up Monitoring Function.](#)"

 **Information**

The settings for the I/O Response Time Assurance can be done when creating a class using the `sdxdisk -M` command. For details, see "[B.1.2 sdxdisk - Disk Operations.](#)"

6.12.2 Hang-up Monitoring Function

Because the GDS I/O Monitor Option assures the I/O response time, it is equipped with hang-up monitoring functions to monitor whether or not the GDS daemon or driver is hung-up.

The hang-up monitoring function monitors whether or not the `sdxservd` daemon and the `sfdsd` driver of the current node are hung-up, and if a hang-up is detected it panics the current node.

The hang-up monitoring function checks the status of the daemon and driver at a 154 seconds monitoring interval, if there is no response within the 77 second conclusion time it will be determined to be hung-up.

If a value other than the default value is set for the I/O response time, change the monitoring interval and conclusion time parameter of the hang-up monitoring function as follows:

Item	Parameter Name	Explanation
Monitoring Interval	SDX_HCHK_INTERVAL	The time interval (seconds) after monitoring has finished before the next monitoring start. Set to double the time of the conclusion time. Default Value: 154 Possible setting values: 12 or more, 2592000 or less, and a value greater than that of the conclusion time (seconds).
Conclusion Time	SDX_HCHK_CRITERION	Time (seconds) after monitoring has been started before it is judged to be hung-up. Set to the same value as the I/O response time. If the value of the I/O response time differs depending on the class, set the shortest time. Default value: 77 Possible setting values: 1 or more, 3600 or less, and a value less than the monitoring interval (seconds).

How to change parameters

1. Change the setting parameters in the GDS settings file `/etc/opt/FJSVsdx/sdx.cf`, or add parameter settings.

Example) If setting the monitoring time interval to 156 seconds and the conclusion time to 78 seconds

```
SDX_HCHK_INTERVAL=156
```

```
SDX_HCHK_CRITERION=78
```

2. Restart the system. Or execute the following command.

```
# /usr/bin/systemctl restart fjsvdxhealth.service
```

3. After restarting the system or service with step 2, make sure that the service is operating.

```
# /usr/bin/systemctl status fjsvdxhealth.service
* fjsvdxhealth.service - PRIMECLUSTER Global Disk Services hang up monitor service
...
Active: active (running) ...
...
```

Note

.....

No monitoring is done in single user mode.

.....

Chapter 7 Operation and Maintenance

This section explains how to operate and maintain GDS, GDS Snapshot, and GDS I/O Monitor Option.

7.1 Viewing Configurations/Statuses and Monitoring Statuses

In the Main Screen of GDS Management View, object configuration and status viewing and status monitoring can be performed.

7.1.1 Confirming SDX Object Configuration

Select [SDX Object] in the [View]:[Details] menu to view object configurations according to the following units.

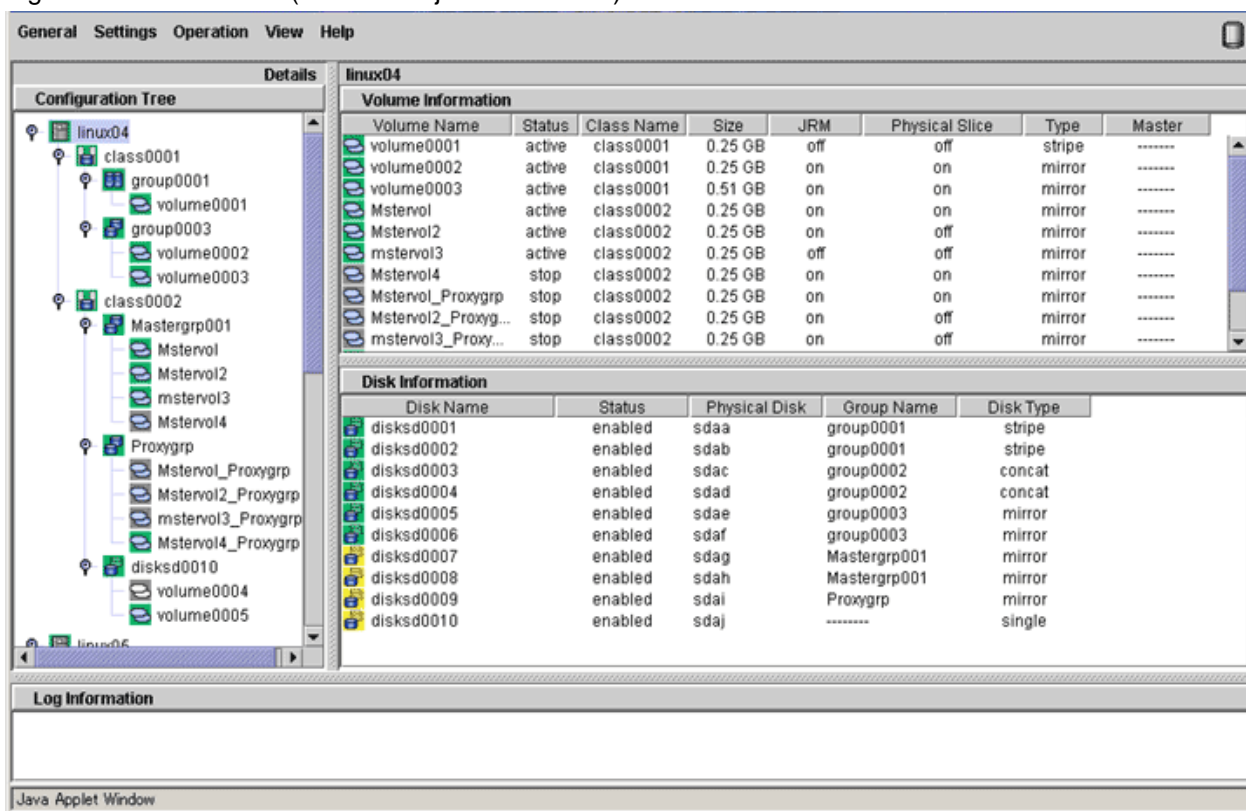
- Object configuration within a node
- Object configuration within a class
- Object configuration within a group
- Object configuration within a single disk
- Object configuration within a volume

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.

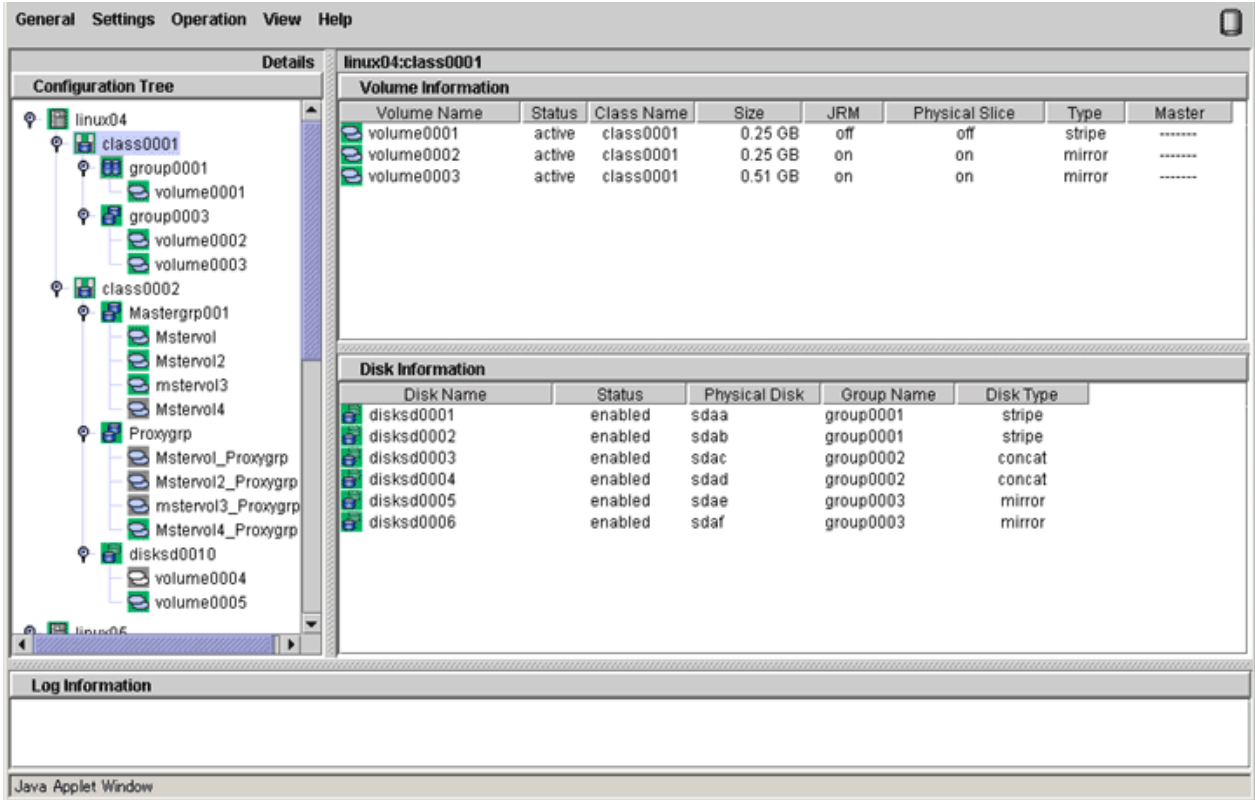
Figure 7.1 Main Screen (for SDX Objects of a Node)



Object configuration within a class

Click the class icon in the Configuration Tree field, and all volumes and disks within the class appear.


Figure 7.2 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 7.3 Main Screen (for SDX Objects of a Group)



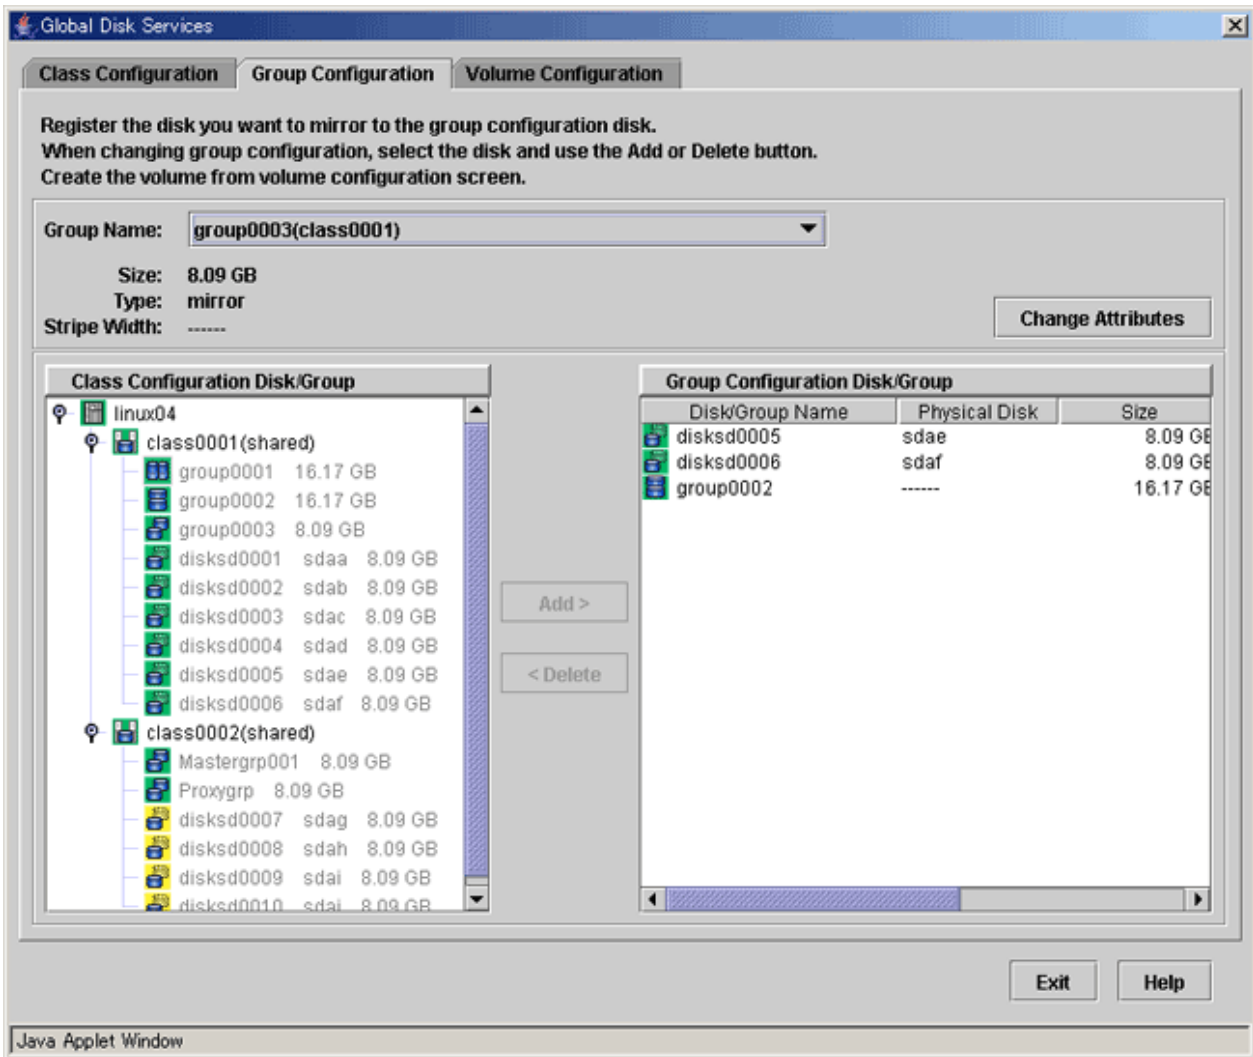
Groups displayed in the GDS configuration tree field are only the highest level groups. You can view the disks and lower level groups constituting their lower level groups in the Group Configuration screen.

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 `disksd0005`, `disksd0006` and lower level group `group0002`. In a similar manner, view the configuration of lower group `group0002`.


Figure 7.4 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.

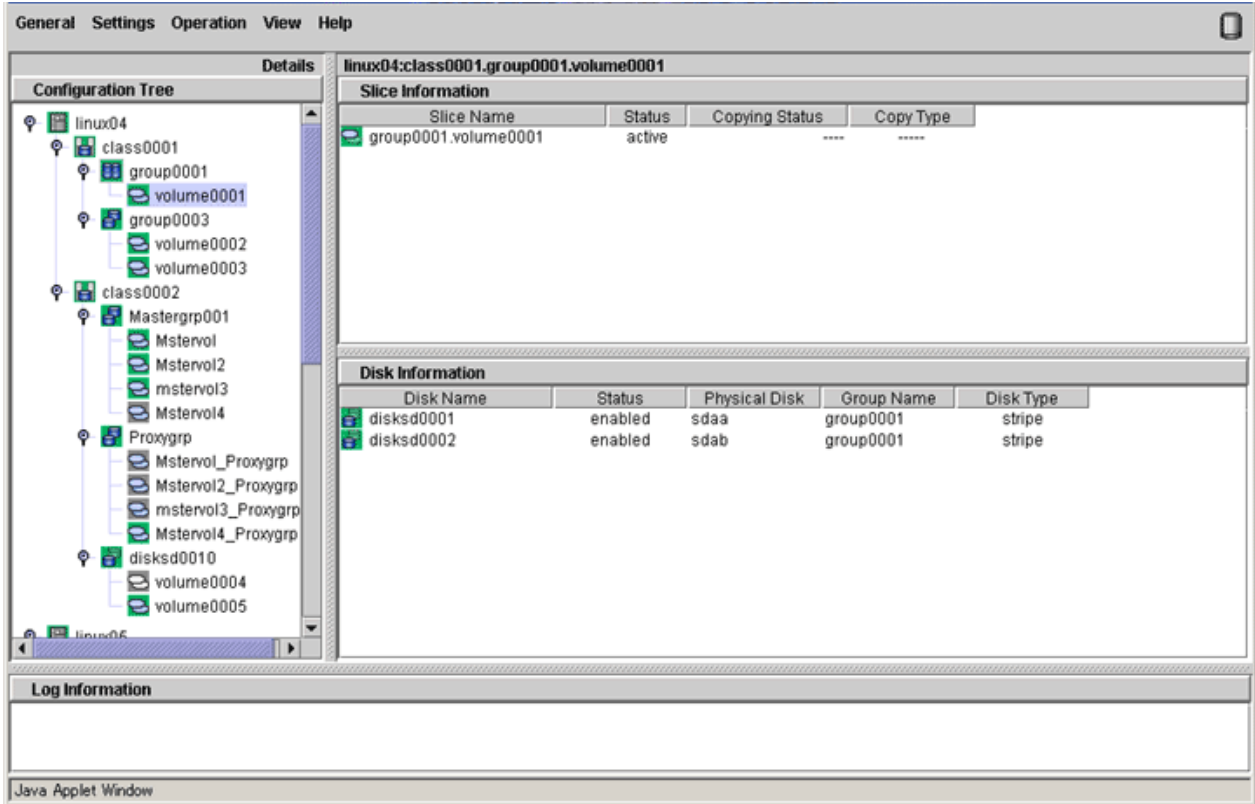
Figure 7.5 Main Screen (for SDX Objects of a Single Disk)



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 7.6 Main Screen (for SDX Objects of a Volume)



7.1.2 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 7.7 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 7.8 Main Screen (for Proxy Objects of a Class)

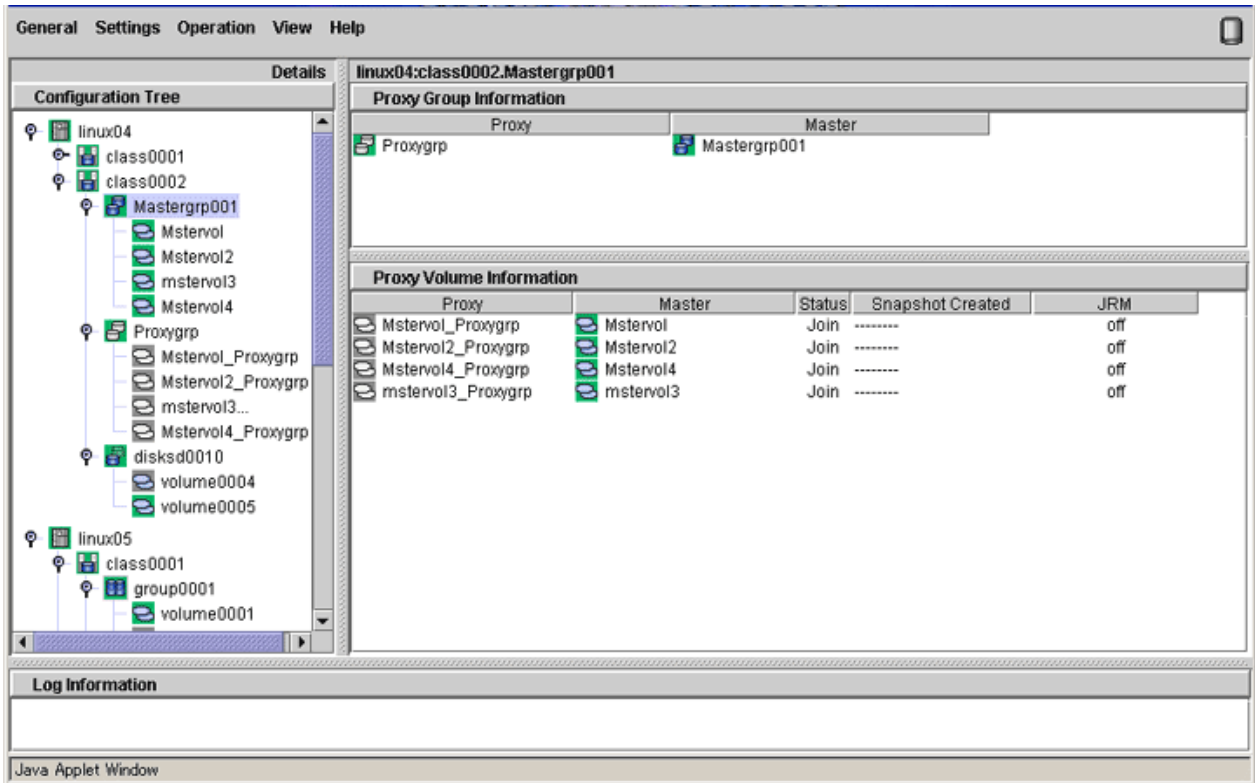


Proxy object configuration related to a group

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

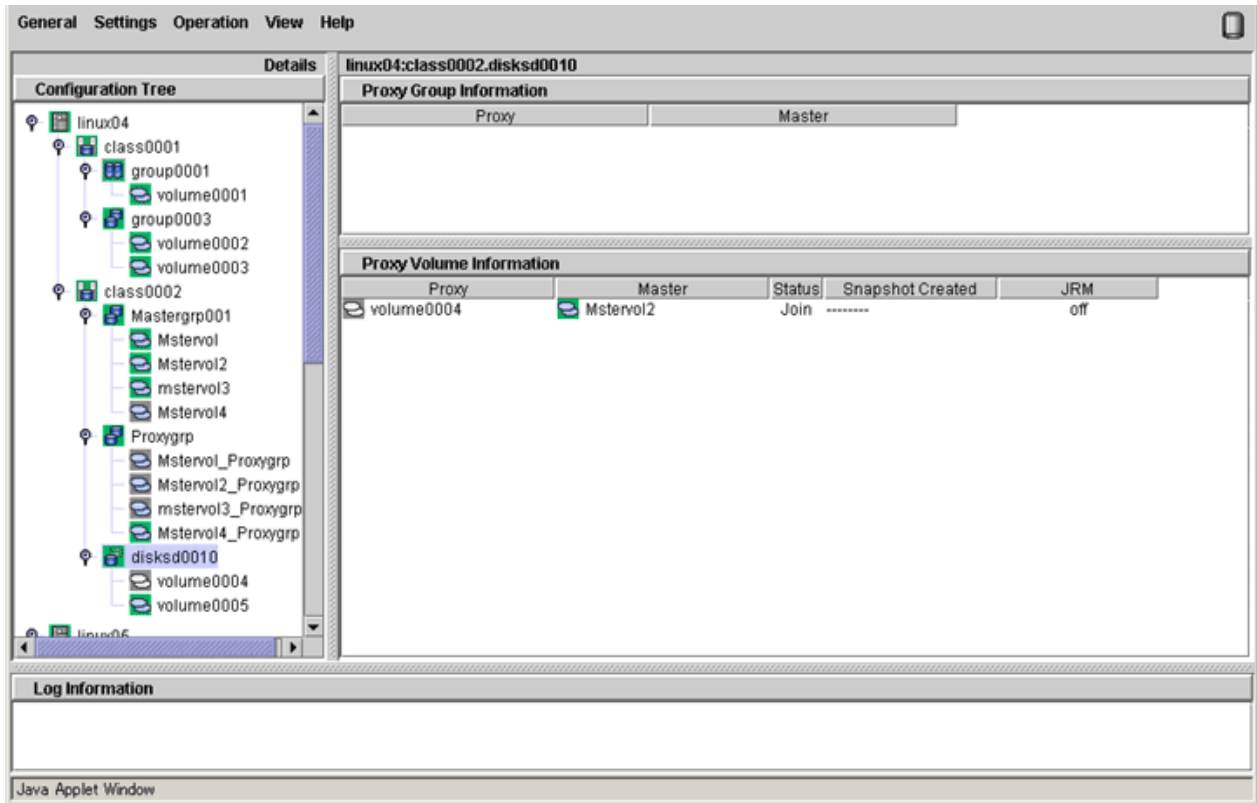
Figure 7.9 Main Screen (for Proxy Objects of a Group)



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 7.10 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 7.11 Main Screen (for Proxy Objects of a Volume)



7.1.3 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.



For details on warning lamps, see "5.2 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.




For details on icons, see "5.4 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 "7.3.1.2 Operation Procedure."

Figure 7.12 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 7.13 Main Screen when [View]: [Abnormal Object] is selected



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.

7.1.4 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 "5.4 Icon Types and Object Status."

1. Checking the state of an object

Using one of the following methods, check the state of an object.

- [Check Status] in the popup menu
In the Main Screen, right click the desired object and select [Check Status] from the popup menu.
- [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 7.14 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.

7.2 Synchronization Copy Operation of Mirror Volume

7.2.1 Using GDS Management View

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 7.15 Start Copying



The confirmation screen shown below appears.

To continue the process, click <Yes>. Clicking <No> cancels the copy start process.


Figure 7.16 Confirming Start Copying



3. **[Start Copying] completion screen**

Click <OK>, and the message notifying the [Start Copying] completion screen shown below appears.

Figure 7.17 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

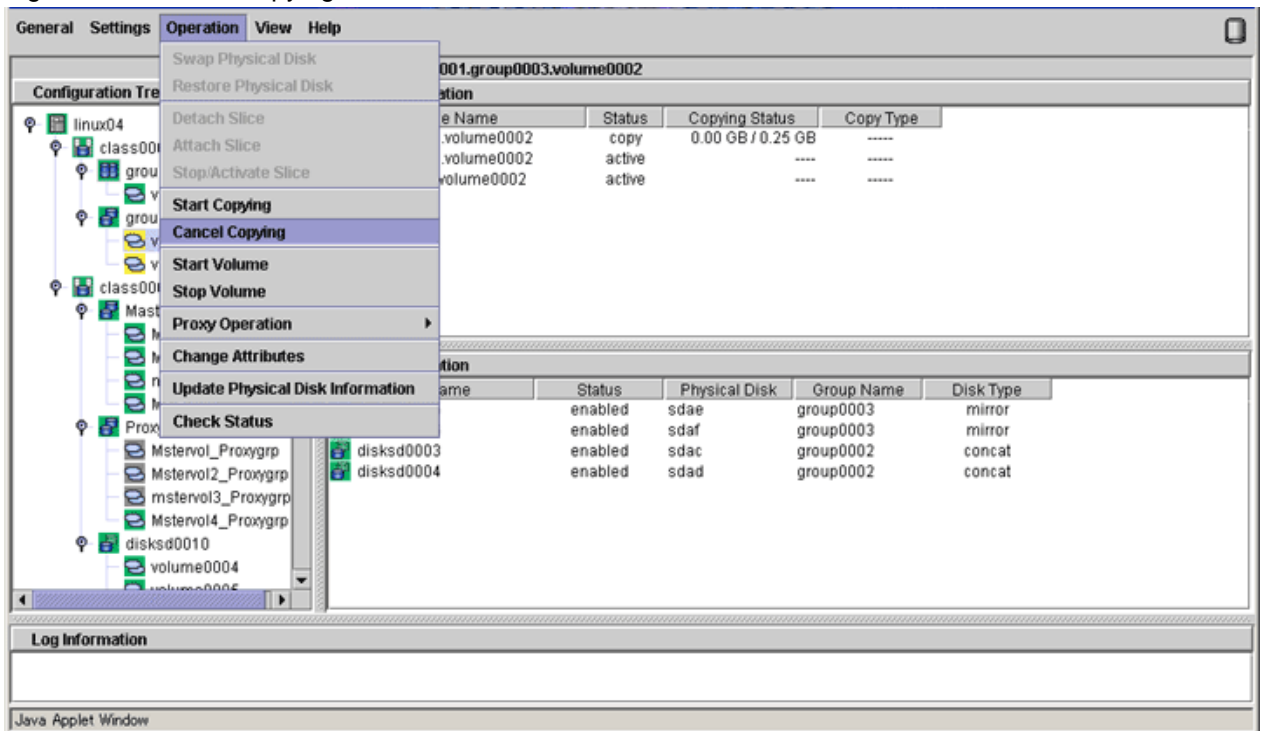
1. **Select the volume to cancel copying**

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 7.18 Cancel Copying



The confirmation screen shown below appears.

To continue the process, click <Yes>. Clicking <No> cancels the copy cancellation process.


Figure 7.19 Confirming Cancel Copying



3. [Cancel Copying] completion screen

Click <OK>, and the message notifying the [Cancel Copying] completion screen shown below appears.

Figure 7.20 Notifying Cancel Copying Completion



Note

[Cancel Copying] operation unavailable

When more than one volume is selected, [Cancel Copying] operation cannot be performed.

7.2.2 Using Commands

You can start, cancel, interrupt the synchronization copying, as well as to change parameters of the synchronization copying by using the `sdxcopy` command.

See

For details on the `sdxcopy` command, see "[B.1.10 sdxcopy - Synchronization Copying Operation.](#)"

7.3 Disk Swap

This section explains how to swap the physical disk that is registered in GDS.

Disk to be swapped	Reference
A disk connected with the netmirror group	"7.3.3 Swapping Disks of Netmirror Group"
A disk registered in the shadow class	"7.3.4 Swapping Physical Disks Registered with Shadow Classes"
Other disks	"7.3.1 Setting with GDS Management View" or "7.3.2 Using Commands"


GDS: Global Disk Services

7.3.1 Setting with GDS Management View

7.3.1.1 Operation Flow

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

Figure 7.21 Disk Swap and Restoration Operation



7.3.1.2 Operation Procedure

When a disk unit abnormality occurs, contact field engineers 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/log/messages` file. For details, see "D.1.13 Disk Unit Error."

Note

Notes on Physical Disk Swap

See "7.3.4 Swapping Physical Disks Registered with Shadow Classes," and "7.3.5 Notes."

Swap Physical Disk

In order to swap the disk units, you must take the physical disk offline.

The procedures are explained below.

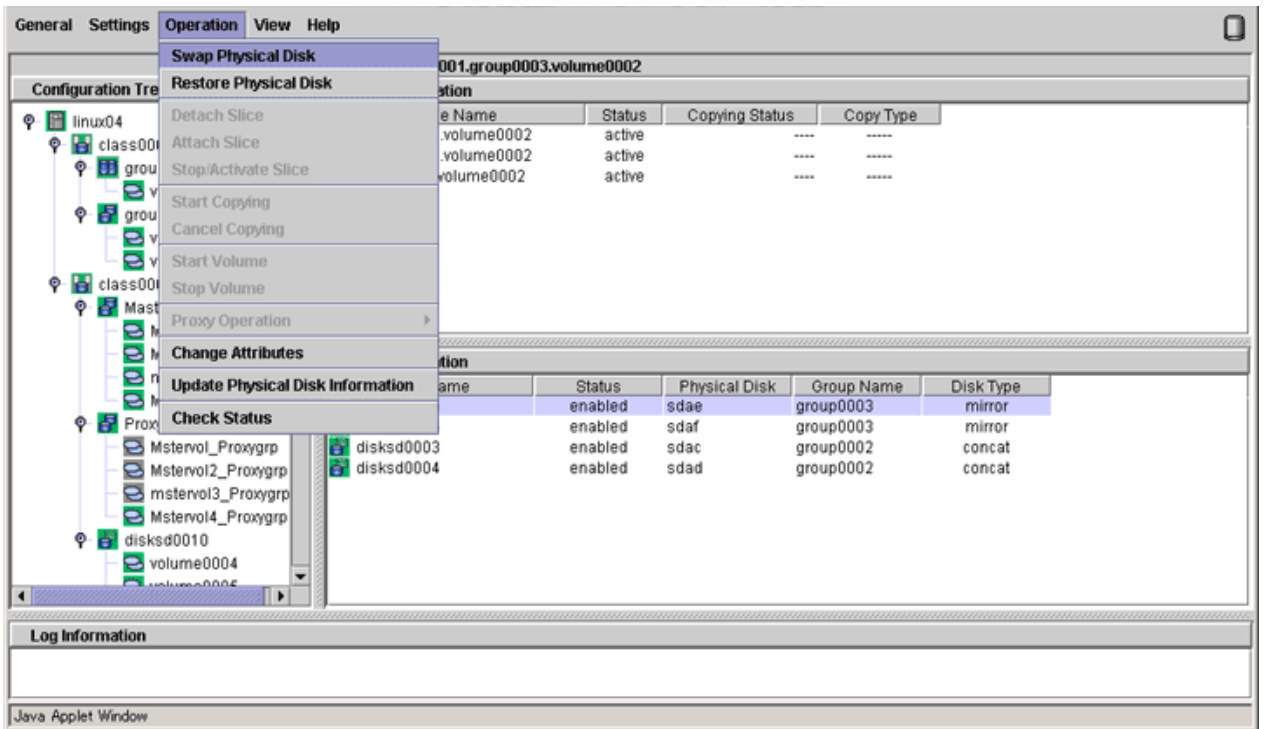
1. **Displaying the status of physical disk**

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 7.22 Swap Physical Disk



The confirmation screen shown below appears.

To continue the process, click <Yes>. Clicking <No> cancels the physical disk swapping process.

Figure 7.23 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 field engineers to swap the disk unit.

Figure 7.24 Notifying Offline Completion



Restore Physical Disk

After swapping the disk units, you must put the swapped physical disk back online.

The procedures are explained below.

1. Check the device name change

When an internal disk registered in the root class or local class swapped, a physical disk cannot be restored if there is a device name change that means the physical disk name is different from the name at the disk registration. Check that there is no difference between the device name of the swapped internal disk and the device name managed by GDS.



See

.....
For the method to check the device name change, see "[7.3.5.4 Swapping Internal Disks Registered with Root Classes or Local.](#)"
.....


2. Selecting the physical disk to restore

Select the physical disk you want to restore.

3. Selecting [Restore Physical Disk] menu

In the Main screen [Operation] menu, select [Restore Physical Disk].


Figure 7.25 Restore Physical Disk



The confirmation screen shown below appears.

To continue the process, click <Yes>. Clicking <No> cancels the physical disk restore process.

Figure 7.26 Confirming Restore Physical Disk



4. Notifying online status

Click <OK>, and a message notifying the online status appears.

Figure 7.27 Notifying Online Completion



Note

If the I/O Response Time Assurance Function is set

If there is not enough memory the physical disk restoration might become error.

In this case, the following messages are displayed:

5000

Command line output: ERROR: sfdsk driver returned an error, errno=12

Also, a warning message 22018 (WARNING: sfdsk: failed to open disk, no enough address space: ...) is output to the console.

In this situation, perform any of the following actions before starting the physical disk restoration.

- (a) Start the system again.
- (b) Wait until the memory is freed up.
- (c) Increase the memory.

Refer to the installation guide for the necessary memory.

Plan the memory increase if selecting (a) or (b).

7.3.2 Using Commands

When replacing the disk by using the command, execute the `sdxswap -O` command instead of "Swap physical disk", and execute the `sdxswap -I` command instead of "Restore physical disk" in "7.3.1 Setting with GDS Management View."

See

For details on the `sdxswap` command, see "B.1.8 `sdxswap` - Swap Disk."

7.3.3 Swapping Disks of Netmirror Group

7.3.3.1 Hot Swap

To perform hot swap (swap without stopping the system) of a disk connected to the netmirror group, perform the following procedure.

1. Make a disk exchangeable.

- For shared classes

Execute "Swap Physical Disk" (or the `sdxswap -O` command) on the node where a disk is to be swapped.

If multiple partitions on the disk to be swapped are the targets of the mirroring among servers, execute "Swap Physical Disk" (or the `sdxswap -O` command) for iSCSI devices for all the partitions.

- For local classes (when object configurations were restored with the `sdxconfig Restore` command)

Delete the entire object configuration (volumes, groups, and disks) of a local class to delete the local class.

As for the order of deleting disks, delete a disk to be swapped with the `sdxdisk -R` command first.

If multiple partitions on the disk to be swapped are the targets of the mirroring among servers, delete the object configurations for iSCSI devices for all the partitions.

Note

If a netmirror group exists in a local class, "Swap Physical Disk" (or the `sdxswap -O` command) cannot be executed.

2. Stop iSCSI devices for disks to be swapped on both nodes.

If multiple partitions on disks to be swapped are the target of the mirroring among servers, execute it for iSCSI devices for all the partitions.

```
# echo offline > /sys/block/iSCSI_device_name/device/state
```

Execution example

```
# echo offline > /sys/block/sdd/device/state
```

3. Back up the target information on a node where a disk is swapped.

```
# cp /etc/target/saveconfig.json backup_file
```

Execution example

```
# cp /etc/target/saveconfig.json /var/tmp/targetfile.json
```

4. Confirm storage and LUN for a device to be swapped on a node where a disk is swapped.

If multiple partitions on disks to be swapped are the target of the mirroring among servers, confirm all the partitions.

```
# targetcli ls
```

Output example

```
# targetcli ls
o- / ..... [ ... ]
  o- backstores ..... [ ... ]
    | o- block ..... [Storage Objects: 2]
    | | o- storage1 . [/dev/disk/by-id/scsi-3500000e114784363 (8.0GiB) write-thru activated] (*1)
    | | o- storage2. [/dev/disk/by-id/scsi-3500000e1109cc830 (16.0GiB) write-thru activated]
    | o- fileio ..... [Storage Objects: 0]
    | o- pscsi ..... [Storage Objects: 0]
    | o- ramdisk ..... [Storage Objects: 0]
  o- iscsi ..... [Targets: 1]
    | o- iqn.2003-01.org.linux-iscsi.node1.x8664:sn.6e665c7c1be0 ..... [TPGs: 1]
    |   o- tpg1 ..... [no-gen-acls, no-auth]
    |     o- acls ..... [ACLs: 2]
    |       | o- iqn.1994-05.com.redhat:dee92ff9979d ..... [Mapped LUNs: 2]
    |         | | o- mapped_lun0 ..... [lun0 block/storage1 (rw)] (*2)
    |         | | o- mapped_lun1 ..... [lun1 block/storage2 (rw)]
    |         | o- iqn.1994-05.com.redhat:fa7eb9cf483c ..... [Mapped LUNs: 2]
    |         |   o- mapped_lun0 ..... [lun0 block/storage1 (rw)] (*2)
    |         |   o- mapped_lun1 ..... [lun1 block/storage2 (rw)]
    |     o- luns ..... [LUNs: 2]
    |       | o- lun0 ..... [block/storage1 (/dev/disk/by-id/scsi-3500000e114784363)] (*3)
    |       | o- lun1 ..... [block/storage2 (/dev/disk/by-id/scsi-3500000e1109cc830)]
    |     o- portals ..... [Portals: 1]
```

```
| o- 192.168.125.123:3260 ..... [OK]
o- loopback ..... [Targets: 0]
```

- (*1) Confirm that the storage name for /dev/disk/by-id/scsi-3500000e114784363 is "storage1".
 - (*2) It is LUN that is mapped to the initiator allowed to connected to. The information related to lun0 was output to the line.
 - (*3) It shows that the storage name for lun0 is "storage1" and the by-id file of the device is /dev/disk/by-id/scsi-3500000e114784363.
- In the example, if you swap the device indicated by /dev/disk/by-id/scsi-3500000e114784363, the storage name is "storage1" and LUN is "lun0".

5. Delete all the LUNs to be swapped on a node where a disk is to be swapped.

If multiple partitions on disks to be swapped are the target of the mirroring among servers, execute this step for storages for all the partitions.

```
# targetcli /backstores/block/ delete storage_name
```

 **Point**

Specify the storage name confirmed in step 4 to *storage_name*.

Execution example

```
# targetcli /backstores/block/ delete storage1
Deleted storage object storage1.
```

6. Delete storage and LUN for a device to be swapped on a node where a disk is to be swapped.

If multiple partitions on disks to be swapped are the target of the mirroring among servers, confirm storage and LUN for all the partitions.

```
# targetcli ls
```

 **Point**

Confirm that the information of the deleted storage and LUN is not output.

Output example

```
# targetcli ls
o- / ..... [...]
  o- backstores ..... [...]
    | o- block ..... [Storage Objects: 1]
    | | o- storage2 . [/dev/disk/by-id/scsi-3500000e1109cc830 (16.0GiB) write-thru activated]
    | o- fileio ..... [Storage Objects: 0]
    | o- pscsi ..... [Storage Objects: 0]
    | o- ramdisk ..... [Storage Objects: 0]
  o- iscsi ..... [Targets: 1]
    | o- iqn.2003-01.org.linux-iscsi.nodel.x8664:sn.6e665c7c1be0 ..... [TPGs: 1]
    | | o- tpg1 ..... [no-gen-acls, no-auth]
    | | | o- acls ..... [ACLs: 2]
    | | | | o- iqn.1994-05.com.redhat:dee92ff9979d ..... [Mapped LUNs: 1]
    | | | | | o- mapped_lun1 ..... [lun1 block/storage2 (rw)]
    | | | | o- iqn.1994-05.com.redhat:fa7eb9cf483c ..... [Mapped LUNs: 1]
    | | | | | o- mapped_lun1 ..... [lun1 block/storage2 (rw)]
    | o- luns ..... [LUNs: 1]
    | | o- lun1 ..... [block/storage2 (/dev/disk/by-id/scsi-3500000e1109cc830)]
```

```
|      o- portals ..... [Portals: 1]
|      o- 192.168.125.123:3260 ..... [OK]
o- loopback ..... [Targets: 0]
```

7. Swap the disk.

For the procedures for swapping disks, refer to the manuals related to the servers and the disks.

8. If the mirroring among servers is performed not per the physical disk but per partition, create partitions on the swapped disk on the node where the disk was swapped.

The partition size and the order to create the partition must be consistent with the settings before and after the disk is swapped.

In an Azure environment, a NIFCLOUD environment, or an environment where the by-id file does not exist, set the gpt disk label in the swapped disk, and create a partition.

The partition size must be consistent with the size before the disk is swapped.

9. Edit the target information that was backed up in step 3 and modify the path of the by-id file for the original disk to that for the new disk.

If multiple partitions on the disk to be swapped are the targets of the mirroring among servers, take this step for the path of the by-id file for all the partitions.

In an Azure environment, a NIFCLOUD environment, or an environment where the by-id file does not exist, modify the path of the by-partuuid file for the original disk to by-partuuid for the partition of the new disk.

```
# vim backup_file

{
  "fabric_modules": [],
  "storage_objects": [
    {
      ...
      "dev": "original_disk",      (*1)
      "name": "storage_name",
      "plugin": "block",
      "readonly": false,
      "write_back": false,
      "wwn": "fae93a41-0d59-4843-9137-9c39854e5388"
    }
  ],
}
```

(*1) Modify to the path of the by-id file or the path of the by-partuuid file for the new disk.

 Point

Specify the backup file that was backed up in step 3 to *backup_file*.

[Example before modification]

```
# vim /var/tmp/targetfile.json

{
  "fabric_modules": [],
  "storage_objects": [
    {
      ...
      "dev": "/dev/disk/by-id/scsi-3500000e114784363",
      "name": "storage1",
      "plugin": "block",
      "readonly": false,
      "write_back": false,
    }
  ],
}
```

```
        "wwn": "279b8505-418f-4449-8987-f3acaad34618"
    }
],
```

[Example after modification]

```
# vim /var/tmp/targetfile.json

{
  "fabric_modules": [],
  "storage_objects": [
    {
      ...
      "dev": "/dev/disk/by-id/scsi-3500000e111c56611",
      "name": "storaget1",
      "plugin": "block",
      "readonly": false,
      "write_back": false,
      "wwn": "279b8505-418f-4449-8987-f3acaad34618"
    }
  ],
}
```

10. Restore the target information that was edited in step 9 on a node where a disk is to be swapped.

```
# targetctl restore backup_file
```

 **Point**

Specify the backup file that was created in step 9 to *backup_file*.

Execution example

```
# targetctl restore /var/tmp/targetfile.json
```

11. Confirm that the storage and LUN was added to a node where a disk is to be swapped.

If multiple partitions on the disk to be swapped are the targets of the mirroring among servers, confirm the storage and LUN for all the partitions.

```
# targetcli ls
```

 **Point**

- (*1) Confirm that the device correspond to the storage that was deleted in step 5 is changed to the swapped device.
- (*2) Confirm that LUN which was deleted in step 5 is mapped to the initiator allowed to connected to.
- (*3) Confirm the following:
 - LUN and the storage that were deleted in step 5 are added.
 - The path of the device is specified to the device that was swapped.

Output example

```
# targetcli ls
o- / ..... [ ... ]
  o- backstores ..... [ ... ]
```

```

| o- block ..... [Storage Objects: 2]
| | o- storage1 . [/dev/disk/by-id/scsi-3500000e111c56611 (8.0GiB) write-thru activated] (*1)
| | o- storage2.[/dev/disk/by-id/scsi-3500000e1109cc830 (16.0GiB) write-thru activated]
| o- fileio ..... [Storage Objects: 0]
| o- pscsi ..... [Storage Objects: 0]
| o- ramdisk ..... [Storage Objects: 0]
o- iscsi ..... [Targets: 1]
| o- iqn.2003-01.org.linux-iscsi.node1.x8664:sn.6e665c7c1be0 ..... [TPGs: 1]
|   o- tpg1 ..... [no-gen-acls, no-auth]
|     o- acls ..... [ACLs: 2]
|       | o- iqn.1994-05.com.redhat:dee92ff9979d ..... [Mapped LUNs: 2]
|         | | o- mapped_lun0 ..... [lun0 block/storage1 (rw)] (*2)
|           | | o- mapped_lun1 ..... [lun1 block/storage2 (rw)]
|           | o- iqn.1994-05.com.redhat:fa7eb9cf483c ..... [Mapped LUNs: 2]
|             | o- mapped_lun0 ..... [lun0 block/storage1 (rw)] (*2)
|               | o- mapped_lun1 ..... [lun1 block/storage2 (rw)]
|             o- luns ..... [LUNs: 2]
|               | o- lun0 ..... [block/storage1 (/dev/disk/by-id/scsi-3500000e111c56611)] (*3)
|                 | o- lun1 ..... [block/storage2 (/dev/disk/by-id/scsi-3500000e1109cc830)]
|               o- portals ..... [Portals: 1]
|                 o- 192.168.125.123:3260 ..... [OK]
o- loopback ..... [Targets: 0]

```

12. Store the target information that was restored in step 10 on a node where a disk is to be swapped.

```
# targetctl save
```

13. Rescan an iSCSI session on both nodes.

```
# iscsiadm -m session --rescan
```

14. Enable iSCSI devices on both nodes.

If the mirroring among servers is performed per partition, enable iSCSI devices for all the partitions.

```
# echo running > /sys/block/iSCSI_device_name/device/state
```

Execution example

```
# echo running > /sys/block/sdd/device/state
```

At the time, the following OS message may be output to the system log. An action is not required.

```
kernel: sd X:X:X:X: Warning! Received an indication that the LUN assignments on this target have
changed. The Linux SCSI layer does not automatically remap LUN assignments.
```

15. Restore the GDS configuration.

- For shared classes

On the node where the disk was swapped, execute "Restore Physical Disk" (or the `sdxxswap -I` command). Execute "Restore Physical Disk" (or the `sdxxswap -I` command) for all the iSCSI devices where "Swap Physical Disk" (or the `sdxxswap -O` command) was executed in step 1.

- For local classes

Re-create the object configuration (volumes, groups, disks, and classes) of the deleted local class.

- When the configuration information file of the class exists

When the class configuration information obtained with the `sdxcnfig Backup` command exits, re-create the object configuration with the `sdxcnfig Restore` command.

- When the configuration information file of the class does not exist
Specify a shared class to re-create the object configuration.

7.3.3.2 Cold Swap

To perform cold swap (swap in a stopped system) of a disk connected to the netmirror group, perform the following procedure.

1. Make a disk exchangeable.

- For shared classes

Execute "Swap Physical Disk" (or the `sdxswap -O` command) on the node where a disk is to be swapped.

If multiple partitions on the disk to be swapped are the targets of the mirroring among servers, execute "Swap Physical Disk" (or the `sdxswap -O` command) for iSCSI devices for all the partitions.

- For local classes (when object configurations were restored with the `sdxconfig Restore` command)

Delete the entire object configuration (volumes, groups, and disks) of a local class to delete the local class.

As for the order of deleting disks, delete a disk to be swapped with the `sdxdisk -R` command first.

If multiple partitions on the disk to be swapped are the targets of the mirroring among servers, delete the object configurations for iSCSI devices for all the partitions.

Note

If a netmirror group exists in a local class, "Swap Physical Disk" (or the `sdxswap -O` command) cannot be executed.

2. If you swap the disk of the operating node while the service continues, switch the cluster application to make a node, which is not connected to the disk to be swapped, as an operating node.
3. Modify the settings for starting the system on the node where the disk is to be swapped so that the system is restarted in single user mode.

```
# systemctl set-default rescue.target
```

Because the console cannot be used in a cloud environment (except an FJcloud-Baremetal environment), do not select single user mode. In this case, perform step 3-1 to 3-3 to restrict the automatic startup of the PRIMECLUSTER service.

In an environment other than a cloud environment (except an FJcloud-Baremetal environment), follow the procedure from step 4.

1. Stop RMS.

If RMS is running, execute the following command on any one node to stop RMS.

```
# hvshut -a
```

Note

If RMS is stopped on all nodes during the synchronization copying of the GDS volume, the synchronization copying of the entire volume area is performed after the corrections are applied and all nodes are restarted.

If you do not want to perform the synchronization copying of the entire area of volume, stop RMS after the synchronization copying is completed.

To check the slice status of the GDS volume, execute the following command.

Execute the following command on any one node to check the value of the STATUS field of the command output.

The status of the copy destination slice is COPY during the synchronization copying, and after copying is complete, the status becomes ACTIVE or STOP.

```
# sdxinfo -S
```

2. Restrict the automatic startup of the PRIMECLUSTER service.

Restrict the automatic startup of the PRIMECLUSTER service by executing the following command on the node where the disk is to be swapped.

```
# /opt/FJSPclinst/bin/pclservice off
```

3. Restart the system.

Restart the system on the node where the disk is to be swapped.

```
# /sbin/shutdown -r now
```

4. Stop the node where the disk is to be swapped.

5. Swap the disk.

For the procedures for swapping the disk, refer to the manuals related to the servers and the disks.

6. Start the node where the disk was swapped.

7. If the mirroring among servers is performed not per the physical disk but per partition, create partitions on the swapped disk on the node where the disk was swapped.

The partition size and the order to create the partition must be consistent with the settings before and after the disk is swapped.

In an Azure environment, a NIFCLOUD environment, or an environment where the by-id file does not exist, set the gpt disk label in the swapped disk, and create a partition.

The partition size must be consistent with the size before the disk is swapped.

8. Back up the setting information of iSCSI target on the node where the disk was swapped.

```
# cp /etc/target/saveconfig.json backup_file
```

Execution example

```
# cp /etc/target/saveconfig.json /var/tmp/targetfile.json
```

9. Modify the path of the by-id file for the original disk to that for the new disk in the setting information of the iSCSI target.

If multiple partitions on the disk to be swapped are the targets of the mirroring among servers, take this step for the path of the by-id file for all the partitions.

In an Azure environment, a NIFCLOUD environment, or an environment where the by-id file does not exist, modify the path of the by-partuuid file for the original disk to by-partuuid for the partition of the new disk.

```
# vim /etc/target/saveconfig.json

{
  "fabric_modules": [],
  "storage_objects": [
    {
      ...
      "dev": "original_disk", (*1)
      "name": "storage_name",
      "plugin": "block",
      "readonly": false,
      "write_back": false,
      "wwn": "fae93a41-0d59-4843-9137-9c39854e5388"
    }
  ],
}
```

(*1) Modify to the path of the by-id file or the path of the by-partuuid file for the new disk.

[Example before modification]


```
# vim /etc/target/saveconfig.json

{
  "fabric_modules": [],
  "storage_objects": [
    {
    ...
      "dev": "/dev/disk/by-id/scsi-3500000e111e68e00",
      "name": "store1",
      "plugin": "block",
      "readonly": false,
      "write_back": false,
      "wwn": "4a98bfb0-7d7e-4bc8-962c-0b3cf192b214"
    }
  ],
}
```

[Example after modification]

```
# vim /etc/target/saveconfig.json

{
  "fabric_modules": [],
  "storage_objects": [
    {
    ...
      "dev": "/dev/disk/by-id/scsi-3500000e111c56610", * The line to be modified
      "name": "store1",
      "plugin": "block",
      "readonly": false,
      "write_back": false,
      "wwn": "4a98bfb0-7d7e-4bc8-962c-0b3cf192b214"
    }
  ],
}
```

10. Undo the settings for starting the system modified in step 3 so that the system is restarted in multi-user mode.

Example) When setting back to multi-user.target

```
# systemctl set-default multi-user.target
```

11. Restart the node where the disk was swapped.

In a cloud environment (except an FJcloud-Baremetal environment), take the following procedure to cancel restricting the automatic startup of the PRIMECLUSTER service.

1. Configure the automatic startup of the PRIMECLUSTER service.

Execute the following command on the node where the disk was swapped and change the PRIMECLUSTER service settings back to the state they were in before they were restricted in 2 of step 3.

```
# /opt/FJsvpclinst/bin/pclservice on
```

2. Restart the system.

Restart the system on the node where the disk was swapped.

```
# /sbin/shutdown -r now
```

12. Check that iSCSI target has been set normally on the node where the disk was swapped.

```
# targetcli /backstores/block/ ls ; targetcli /iscsi/ ls
```

(*1) Confirm that the modified path is effective.

Output example

```
o- block ..... [Storage Objects: 1]
  o- store1 ..... [/dev/disk/by-id/scsi-3500000e111c56610 (16.0GiB) write-thru activated] (*1)
o- iscsi ..... [Targets: 1]
  o- iqn.2003-01.org.linux-iscsi.node1.x8664:sn.6e665c7c1be0 ..... [TPGs: 1]
    o- tpg1 ..... [no-gen-acls, no-auth]
      o- acls ..... [ACLs: 2]
        | o- iqn.1994-05.com.redhat:dee92ff9979d ..... [Mapped LUNs: 1]
        | | o- mapped_lun0 ..... [lun0 block/store1 (rw)]
        | o- iqn.1994-05.com.redhat:fa7eb9cf483c ..... [Mapped LUNs: 1]
        | | o- mapped_lun0 ..... [lun0 block/store1 (rw)]
      o- luns ..... [LUNs: 1]
        | o- lun0 ..... [block/store1 (/dev/disk/by-id/scsi-3500000e111c56610)] (*1)
      o- portals ..... [Portals: 2]
        o- 192.168.56.20:3260 ..... [OK]
```

13. On the node where the disk was swapped, check the state of the iSCSI session.

Check that both the iSCSI connection state and the session state indicate "LOGGED IN".

Execution example

```
# iscsiadm -m session -P 1
...
Target: iqn.2003-01.org.linux-iscsi.node1.x8664:sn.6e665c7c1be0
Current Portal: 192.168.56.20:3260,1
Persistent Portal: 192.168.56.20:3260,1
*****
Interface:
*****
...
iSCSI Connection State: LOGGED IN
iSCSI Session State: LOGGED_IN
...
```

14. Rescan an iSCSI session on both nodes.

```
# iscsiadm -m session --rescan
```

15. Restore the GDS configuration.

- For shared classes

On the node where the disk was swapped, execute "Restore Physical Disk" (or the `sdxswap -I` command). Execute "Restore Physical Disk" (or the `sdxswap -I` command) for all the iSCSI devices where "Swap Physical Disk" (or the `sdxswap -O` command) was executed in step 1.

- For local classes

Re-create the object configuration (volumes, groups, disks, and classes) of the deleted local class.

- When the configuration information file of the class exists

When the class configuration information obtained with the `sdxconfig Backup` command exits, re-create the object configuration with the `sdxconfig Restore` command.

- When the configuration information file of the class does not exist

Specify a shared class to re-create the object configuration.

16. Delete the backup file created in step 8.

```
# rm backup_file
```

Execution example

```
# rm /var/tmp/targetfile.json
```

7.3.4 Swapping Physical Disks Registered with Shadow Classes



Shadow objects are not supported in this version.

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 "[B.2 GDS Snapshot Commands](#)."

The subsequent operations vary depending on whether the disks 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.

An error in a disk unit may not cause failures on both the related SDX and shadow objects. Even if only either of the objects fails, shadow disk removal and physical disk swap 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 "[7.3.1 Setting with GDS Management View](#)" and "[7.3.2 Using Commands](#)."

After removing the shadow disk, perform disk swap referring to the manual of the disk unit's copy function.

7.3.5 Notes

7.3.5.1 Physical Disk Size

Physical disk swap cannot be performed using a physical disk whose size is smaller than the original disk size.

7.3.5.2 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.

7.3.5.3 LinuxSwapping Physical Disks when the Proxy Volume Is in a Mirroring Configuration

Since the resynchronization copying is not performed, if the disk which configures the proxy volume is swapped when the proxy volume is in a mirroring configuration, the slice of the proxy volume on the swapped disk will become INVALID. In such a case, temporarily part the proxy objects and then rejoin them.

For an explanation on how to part and rejoin proxy objects, see "B.2.1 sdxproxy - Proxy Object Operations."

7.3.5.4 Swapping Internal Disks Registered with Root Classes or Local

After swapping disks, if there is a device name change that means the physical disk name is different from the name at the disk registration, "Restore Physical Disk" in GDS Management View cannot be performed and also the sdxswap -I command cannot be executed.

Before performing "Restore Physical Disk" in GDS Management View or executing the sdxswap -I command, check that there is no change between the device name of the new internal disk and the device name which is managed by GDS according to the following steps.

Do not swap multiple internal disks at one time. Swap each internal disk one by one, and according to the following steps, confirm that the device name of each internal disk remains unchanged:

1) Obtain information of the physical disk to be swapped

Before swapping the internal disk, obtain the following command output.

```
# ls -l /dev/disk/by-id
```

2) Obtain information of the swapped physical disk

After swapping the internal disk, obtain the following command output.

```
# ls -l /dev/disk/by-id
```

3) Check the device name of swapped internal disk

Compare information 1) and 2) to check the device name of the swapped internal disk.

The swapped internal disk is displayed as a by-id disk name that did not exist before swapping.

The following example shows when the physical disk of /dev/sda is swapped.

Do not compare the partition information among the information of ls command output.

[Before swapping]

```
# ls -l /dev/disk/by-id
lrwxrwxrwx. 1 ... scsi-aaaaaaaaaaaaaaaa -> ../../sda
lrwxrwxrwx. 1 ... scsi-bbbbbbbbbbbbbbbb -> ../../sdb
lrwxrwxrwx. 1 ... scsi-cccccccccccccccc -> ../../sdc
```

[After swapping]

```
# ls -l /dev/disk/by-id
lrwxrwxrwx. 1 ... scsi-xxxxxxxxxxxxxxxx -> ../../sda
lrwxrwxrwx. 1 ... scsi-bbbbbbbbbbbbbbbb -> ../../sdb
lrwxrwxrwx. 1 ... scsi-cccccccccccccccc -> ../../sdc
```

4) Check the device name of the internal disk managed by GDS

Using the following command, check what device name is used for the original internal disk managed by GDS.

In <Class_name>, specify a class name registered to the original internal disk. In <Disk_name>, specify the disk name of the original internal disk.

```
# /etc/opt/FJVSvdx/bin/sdxdevinfo -c <Class_name> -d <Disk_name>
```

Example)

```
# /etc/opt/FJVSvdx/bin/sdxdevinfo -c RootClass -d rootDisk0001
class      disk      device  by-id
-----
RootClass  rootDisk0001  sda     3500000e111c01810
```

5) Check device name change

The device name of the swapped internal disk checked in step 3) remains unchanged if it is consistent with the device name managed by GDS that was checked in step 4).

If the device name is changed, restart the system and check that the device name is returned to its original name.

If the device name remains changed, identify the cause and take the corrective action.

7.4 Backing Up and Restoring 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 breakdowns and so on or by an operation mistake.

This section discusses the following content:

- Necessary settings before and after backing up and restoring system disk.
For details, see "[7.4.1 Settings Before Backing Up](#)," "[7.4.3 Settings After Backing Up](#)," "[7.4.4 Settings Before Restoring](#)," and "[7.4.7 Settings After Restoring](#)."

Note

When the shared class or local class of GDS is used, the settings are necessary regardless of whether the system disk is mirrored by GDS.

- 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.
For details, see "[7.4.2 Backing Up \[EFI\]](#)," "[7.4.5 Restoring \(When System Can Be Booted\) \[EFI\]](#)," and "[7.4.6 Restoring \(When System Cannot Be Booted\) \[EFI\]](#)."

Note

- Data backed up before system disk mirroring cannot be restored back to the mirrored system disk.
- If the system disk mirroring is once canceled and then configured again, you cannot restore the backed up data of the system disk that was collected when the system disk was mirrored previously.

For these reasons, by using this procedure, make sure to collect the backed up data of the system disk if the system disk mirroring is configured.

Also make sure to collect the backed up data of the system disk again if the system disk mirroring is configured again.

See

For backing up and restoring volumes (e.g. /opt, /home) other than system volumes (/ , /usr, /var, /boot, /boot/efi, swap area) among volumes in the root class, see "[7.6 Backup \(Offline\)](#)," "[7.7 Backup \(by Slice Detachment\)](#)," "[7.8 Backup \(by Synchronization\)](#)," and "[7.9 Backup \(by OPC\)](#)."

7.4.1 Settings Before Backing Up

This section explains about the settings to be done before backing up the system disk.

When the shared class or local class of GDS is used, perform the following settings regardless of whether the system disk is mirrored by GDS.

[Setting procedure]

1. Set the configuration parameter of GDS.
Add the setting of SDX_DB_FAIL_NUM=0 to the end of /etc/opt/FJSVsdx/sdx.cf configuration file.

```
# vim /etc/opt/FJSVsdx/sdx.cf
...
SDX_DB_FAIL_NUM=0    <- Add
```

The reason to perform this setting is to prevent the shared/local class from closing down after restoring even though the shared/local class configuration database layout differs from backing up to restoring the system disk.

2. When using the mirroring among servers function, restrict the automatic startup of RMS.

2-1. Check if the automatic startup of RMS is restricted.

If 0 is displayed, automatic startup is restricted.

If 1 is displayed, automatic startup is not restricted.

```
# hvsetenv HV_RCSTART
```

2-2. When 1 is displayed in step 2-1, restrict the automatic startup of RMS.

```
# hvsetenv HV_RCSTART 0
# hvsetenv HV_RCSTART
0
```

7.4.2 Backing Up [EFI]

This section explains the procedure of backing up the data of the system disk mirrored with GDS to the tape.

For securing consistency of backup target files, boot the system from a CD-ROM device or boot the system in single user mode to create backups. To ensure consistency, it is recommended to boot from a CD-ROM device.

a) When booting the system from a CD-ROM device 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 mount(8) or fsck(8) command is executed for the backup target file system, this command may write data to the backup target slice. In these circumstances, unmirror system disks temporarily in advance using this procedure.

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 sda     35368272 node1     ENABLE
disk  Root2  undef  System  *      sdb     35368272 node1     ENABLE
```

Information

If disk Root1 has an INVALID slice, disconnect the Root1. When a keep disk was disconnected (if the disconnected disk's TYPE field value is keep in the `sdxinfo -D` command output), to connect the disk to the group later, change the type attribute to undef (or remove the disk from the class once and register it as an undef disk). For changing the disk type attributes, see "[Changing the Disk Attributes](#)" in "[8.2.1.2 Class Configuration](#)" or "[B.1.7 sdxattr - Change 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 "[8.2.1.3 Group Configuration](#)."
- For details on the `mount(8)` command and the `fsck(8)` command, see the Linux 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 "[7.4.6 Restoring \(When System Cannot Be Booted\) \[EFI\]](#)," restore the backup target volumes.

a2) When using the shared class or local class, set the necessary settings according to "[7.4.1 Settings Before Backing Up](#)."

a3) Shut down the system.

```
# shutdown -h now
```

a4) Turn on the power of the node, and insert the OS installation CD into the CD-ROM drive.

a5) From boot devices displayed in the boot option selection screen of the EFI boot manager, select the CD-ROM device, and boot the system in rescue mode.

For details, see the OS manual.

a6) Check the backup target physical slice name.

Check the backup target physical disk name as follows.

```
# ls -l /sys/block/sd* | grep 0000:06:02.0 | grep 0:0:0
lrwxrwxrwx 1 root root 0 Jun 1 2011 /sys/block/sda ->\
../devices/pci0000:00/0000:00:09.0/0000:01:00.0/0000:02:00.0/0000:03:00.0/\
0000:04:03.0/0000:06:02.0/host1/port-1:0/end_device-1:0/target1:0:0/1:0:0:0/block/sda
```

For the `grep` command arguments, specify physical disk information on the backup target disk (Root1 in this example) confirmed as described in "[6.5.4 Checking Physical Disk Information and Slice Numbers](#)."

In this example, the physical disk name is `sda`.

By combining the physical disk name and the slice numbers confirmed as described in "[6.5.4 Checking Physical Disk Information and Slice Numbers](#)," you can get the physical slice names.

In this example, the backup target physical slice names are as follows.

Use	Physical slice name
/	sda1

Use	Physical slice name
/var	sda2
/usr	sda3
/boot	sda4
/boot/efi	sda5

a7) Back up data of a file system to a tape medium.

The command to be used varies depending on the type of a file system to be backed up. Back up data using the command appropriate for the file system type.

The following example shows the procedure for backing up data to a tape medium of tape device /dev/st0 with the dump(8) command.

```
# dump 0uf /dev/st0 /dev/sda2
```

For the dump command's argument, specify the physical slice displayed in step a6).



For details on the backup methods, see the manuals of file systems to be backed up and used commands.

a8) Exit the rescue mode, and start the system.

```
# exit
```

a9) When using shared class or local class, set the necessary settings according to "7.4.3 Settings After Backing Up."

a10) 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  DEVDNAM  DEVBLKS  DEVCONNECT  STATUS
-----
disk  Root1  mirror System  Group1 sda      35368272 node1       ENABLE
disk  Root2  mirror System  Group1 sdb      35368272 node1       ENABLE
```

Synchronization copying will automatically take place, and when it is completed, the mirroring status is restored.



When using GDS Management View, see "8.2.1.3 Group Configuration."



If a keep disk is disconnected from the group in step a1) and if the type attribute is not changed to undef, step a10) 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 a10).

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) Boot the system in single user mode.

b3) When using shared class or local class, set the necessary settings according to "[7.4.1 Settings Before Backing Up](#)."

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 | grep " / "  
/dev/sfdsk/gdssys2 on / type ext4 (rw)
```

In this example, the device special file for the volume of the root (/) file system is /dev/sfdsk/gdssys2.

b5) Back up data of the file system to a tape medium.

The command to be used varies depending on the type of a file system to be backed up. Back up data using the command appropriate for the file system type.

The following example shows the procedure for backing up data to a tape medium of tape device /dev/st0 with the dump(8) command.

```
# dump 0uf /dev/st0 /dev/sfdsk/gdssys2
```

For the dump command's argument, specify the device special file of the volume displayed in step b4).



See

.....
For details on the backup methods, see the manuals of file systems to be backed up and used commands.
.....

b6) When using shared class or local class, set the necessary settings according to "[7.4.3 Settings After Backing Up](#)."

b7) Reboot the system.

```
# shutdown -r now
```

7.4.3 Settings After Backing Up

This section explains about the settings to be done after backing up the system disk.

When the shared class or local class of GDS is used, perform the following settings regardless of whether the system disk is mirrored by GDS.

[Setting procedure]

1. Undo the configuration parameter setting of GDS.

Comment out the line of SDX_DB_FAIL_NUM=0 in the /etc/opt/FJSVsdx/sdx.cf configuration file.

```
# vim /etc/opt/FJSVsdx/sdx.cf  
...  
# SDX_DB_FAIL_NUM=0
```

2. When using mirroring among servers function, undo the automatic startup setting of RMS.

2-1. If the setting to restrict automatic startup of RMS is done in "[7.4.1 Settings Before Backing Up](#)," set the automatic startup of RMS.

```
# hvsetenv HV_RCSTART 1  
# hvsetenv HV_RCSTART  
1
```

2-2. Start up RMS.

7.4.4 Settings Before Restoring

This section explains about the settings to be done before restoring the system disk.

When mirroring among servers function is used, perform the following settings regardless of whether the system disk is mirrored by GDS.

[Setting procedure]

- When restoring after stopping all the nodes

When all the nodes are started, perform the following settings.

If a node is stopped, the following settings are not necessary.

Execute the following command on any node for all the classes where the netmirror group exists. The latest disk is selected by executing this command.

```
# /etc/opt/FJSVsdx/bin/sdxnetdisk -S -c class_name
```

- When restoring in each node one by one

When restoring in one node while the service continues in the operating node and then switching to another operating node to continue the service for restoring in another node, perform the following settings in advance.

1. Confirm the status of slice which belongs to the netmirror volume.
Execute the following command on any node.

```
# sdxinfo -S
```

2. If the slice is INVALID, perform the synchronization copying for the netmirror volume.
Execute the following command on any node.

```
# sdxcopy -B -c class_name -v volume_name
```

3. Wait until the synchronization copying is completed.
The execution status of synchronization copying can be confirmed in the STATUS field of the sdxinfo -S command output. If "COPY" is displayed, the synchronization copying is still in progress.
When the synchronization copying is completed, "ACTIVE" or "STOP" is displayed.

```
# sdxinfo -S
```

7.4.5 Restoring (When System Can Be Booted) [EFI]

This section explains the procedure of restoring the data of the system disk mirrored with GDS from the tape 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 "7.4.2 Backing Up [EFI]."

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     DEVDNAM  DEVBLKS  DEVCONNECT  STATUS
```

```

-----
disk  Root1  mirror System  Group1  sda      35368272  node1      ENABLE
disk  Root2  undef  System  *        sdb      35368272  node1      ENABLE
-----

```

Information

If disk Root1 has an INVALID slice, disconnect the Root1. When a keep disk was disconnected (if the disconnected disk's TYPE field value is keep in the `sdxinfo -D` command output), to connect the disk to the group later, change the type attribute to undef (or remove the disk from the class once and register it again as an undef disk). For changing the disk type attributes, see "[Changing the Disk Attributes](#)" in "[8.2.1.2 Class Configuration](#)" or "[B.1.7 sdxattr - Change 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 "[8.2.1.3 Group Configuration](#)."

3) When using the mirroring among servers function, set the necessary settings according to "[7.4.4 Settings Before Restoring](#)."

4) Shutdown the system.

```
# shutdown -h now
```

5) Turn on the power of the node, and insert the OS installation CD into the CD-ROM drive.

6) From boot devices displayed in the boot option selection screen of the EFI boot manager, select the CD-ROM device, and boot the system in rescue mode.

For details, see the OS manual.

7) Check the restore destination physical slice name.

Check the restore destination physical disk name.

```
# ls -l /sys/block/sd* | grep 0000:06:02.0 | grep 0:0:0
lrwxrwxrwx 1 root root 0 Jun 1 2011 /sys/block/sda ->\
../devices/pci0000:00/0000:00:09.0/0000:01:00.0/0000:02:00.0/0000:03:00.0/\
0000:04:03.0/0000:06:02.0/host1/port-1:0/end_device-1:0/target1:0:0/1:0:0:0/block/sda
```

For the `grep` command arguments, specify physical disk information on the restore destination disk (Root1 in this example) confirmed as described in "[6.5.4 Checking Physical Disk Information and Slice Numbers](#)."

In this example, the physical disk name is `sda`.

By combining the physical disk name and the slice numbers confirmed as described in "[6.5.4 Checking Physical Disk Information and Slice Numbers](#)," you can get the physical slice names.

In this example, the restore destination physical slice names are as follows.

Use	Physical slice name
/	sda1
/var	sda2
/usr	sda3
/boot	sda4
/boot/efi	sda5

8) Restore backup data on a tape medium back to the file system.

The following example shows the procedure for restoring the root file system using data backed up with the `dump(8)` command. In this example, the file system type is `ext4`, and a temporary mount point is the `/work` directory.

```
# mkdir /work
# mkfs.ext4 /dev/sda1
# mount -t ext4 /dev/sda1 /work
# cd /work
# restore rf /dev/st0 .
# cd /
# umount /work
```

For the `mkfs.ext4(8)` command's and `mount(8)` command's arguments, specify the physical slice displayed in step 7).

Note

Do not perform restoration using data backed up before system disk mirroring.

See

For details on the restore methods, see the manuals of file systems to be restored and used commands.

9) Exit the rescue mode.

```
# exit
```

10) When using shared class or local class, set the necessary settings according to "[7.4.7 Settings After Restoring](#)."

11) Start up the system in multi-user mode.

12) Reconnect the disk disconnected in step 2) 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  DEVDNAM  DEVBLKS  DEVCONNECT  STATUS
-----
```

```

disk Root1 mirror System Group1 sda 35368272 node1 ENABLE
disk Root2 mirror System Group1 sdb 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 "[8.2.1.3 Group Configuration](#)."

 **Information**

If a keep disk is disconnected from the group in step 2) and if the type attribute is not changed to undef, step 12) 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 12).

For the disk type attribute setting method, see "Information" in step 2).

7.4.6 Restoring (When System Cannot Be Booted) [EFI]

This section explains the procedure of restoring the data of the system disk mirrored with GDS from the tape when the system cannot be booted.

- 1) Turn on the power of the node, and insert the OS installation CD into the CD-ROM drive.
- 2) From boot devices displayed in the boot option selection screen of the EFI boot manager, select the CD-ROM device, and boot the system in rescue mode.

For details, see the OS manual.

- 3) Check the restore destination physical slice names.

Check the restore destination physical disk names.

```

# ls -l /sys/block/sd* | grep 0000:06:02.0 | grep 0:0:0
lrwxrwxrwx 1 root root 0 Jun 1 2011 /sys/block/sda ->\
../devices/pci0000:00/0000:00:09.0/0000:01:00.0/0000:02:00.0/0000:03:00.0/\
0000:04:03.0/0000:06:02.0/host1/port-1:0/end_device-1:0/target1:0:0/1:0:0/block/sda

# ls -l /sys/block/sd* | grep 0000:06:02.0 | grep 0:2:0
lrwxrwxrwx 1 root root 0 Jun 1 2011 /sys/block/sdb ->\
../devices/pci0000:00/0000:00:09.0/0000:01:00.0/0000:02:00.0/0000:03:00.0/\
0000:04:03.0/0000:06:02.0/host1/port-1:0/end_device-1:0/target1:0:0/1:0:2:0/block/sdb

```

For the grep command arguments, specify physical disk information on the restore destination disks (Root1 and Root2 in this example) confirmed as described in "[6.5.4 Checking Physical Disk Information and Slice Numbers](#)."

In this example, the physical disk names are sda and sdb.

By combining the physical disk names and the slice numbers confirmed as described in "[6.5.4 Checking Physical Disk Information and Slice Numbers](#)," you can get the physical slice names.

In this example, the restore destination physical slice names are as follows.

Use	Physical slice name	
/	sda1	sdb1
/var	sda2	sdb2
/usr	sda3	sdb3

Use	Physical slice name	
/boot	sda4	sdb4
/boot/efi	sda5	sdb5

Note

When Using the System Volume Snapshot Function

Restore also the joined proxy volumes of system volumes. It is not necessary to restore the parted proxy volumes. If the proxy volume status is unknown, all the proxy volumes of system volumes should be restored.

4) Restore backup data on a tape medium back to one of the two slices.

The following example shows the procedure for restoring the root file system using data backed up with the dump(8) command. In this example, the file system type is ext4, and a temporary mount point is the /work directory.

```
# mkdir /work
# mkfs.ext4 /dev/sda2
# mount -t ext4 /dev/sda2 /work
# cd /work
# restore rf /dev/st0 .
# cd /
# umount /work
```

For the mkfs.ext4(8) command and mount(8) command arguments, specify the device special file for one of the two slices confirmed in step 3).

Note

- If restoration to a slice fails due to an I/O error and so on, restore data to the other slice.
- Do not perform restoration using data backed up before system disk mirroring.

See

For details on the restore methods, see the manuals of file systems to be restored and used commands.

5) Copy data from the slice restored in step 4) to other slices.

The following example shows the procedure for copying data from sda2 to sdb2.

```
# dd if=/dev/sda2 of=/dev/sdb2 bs=1M
```

Note

If the mirroring multiplicity is n, among n mirrored slices, copy data to all "n-1" slices other than the slice restored in step 4).

See

Specify a correct option referring to the dd(1) command manual.

6) Exit the rescue mode.

```
# exit
```

7) When using shared class or local class, set the necessary settings according to "[7.4.7 Settings After Restoring](#)."

8) Start up the system in multi-user mode.

7.4.7 Settings After Restoring

This section explains about the settings to be done after restoring the system disk.

When using the shared class or local class of GDS, perform the following settings regardless of whether the system disk is mirrored by GDS.



After restoring the system disk, perform the following settings before starting up the system in multi-user mode.

[Setting procedure]

1) If "[7.4.1 Settings Before Backing Up](#)" is not executed, start up the system in single-user mode and set according to "[7.4.1 Settings Before Backing Up](#)."

In case of the following, this setting is not necessary.

- When set according to "[7.4.1 Settings Before Backing Up](#)"
- When not using shared class and local class

2) Start up the system in multi-user mode.

3) Check if there are no errors in the configuration and state of the shared and local classes.

```
# sdxinfo
```

4) Restore the configuration of mirroring among servers.

This procedure is performed only when using the mirroring among servers function.

4-1) Delete the iSCSI device information used in mirroring among servers.

This procedure is performed when restoring after stopping all the nodes.

When restoring in another node while the service continues in the operating node, this procedure must not be performed.

Execute the following commands on all the nodes.

```
# rm -f /var/opt/FJSVsdx/log/.sdxnetmirror_disable.db
# rm -f /var/opt/FJSVsdx/log/.sdxnetmirror_timestamp
```

4-2) Restore the slice status of netmirror volume.

This procedure is performed when restoring after stopping all the nodes.

When restoring in another node while the service continues in the operating node, this procedure must not be performed.

Execute the following command on any node for all the netmirror volume.

Check the status of the copy source slice for the synchronization copying when both nodes are stopped before restoring.

Restore the status of the slice using sdxfix command when the copy source slice for the synchronization copying is INVALID.

For the -d option of sdxfix command, specify the copy source disk for the synchronization copying.

Example) When the class name is class0001, the volume name is volume0002, and the disk name is disk0001.

```
# sdxfix -v -c class0001 -v volume0002 -d disk0001 -x NoRdchk
```

4-3) Undo the automatic startup setting of RMS.

If the automatic startup setting of RMS is changed in step 1) or "7.4.1 Settings Before Backing Up," set the automatic startup of RMS.

```
# hvsetenv HV_RCSTART 1
# hvsetenv HV_RCSTART
1
```

4-4) Start up RMS.

5) Undo the configuration parameter settings of GDS.

Comment out the line of SDX_DB_FAIL_NUM=0 in the /etc/opt/FJSVsdx/sdx.cf configuration file.

```
# vim /etc/opt/FJSVsdx/sdx.cf
...
# SDX_DB_FAIL_NUM=0
```

7.5 Backing Up and Restoring System Disk through Alternative Boot Environment [EFI]

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.


7.5.1 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 sdb and sdd 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 7.28 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 sdf and sdh 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 7.29 Proxy Group Configuration





Figure 7.30 Joining Proxy Groups



7.5.2 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 7.31 Backup





Figure 7.32 Backup Schedule



7.5.3 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 7.33 Restore

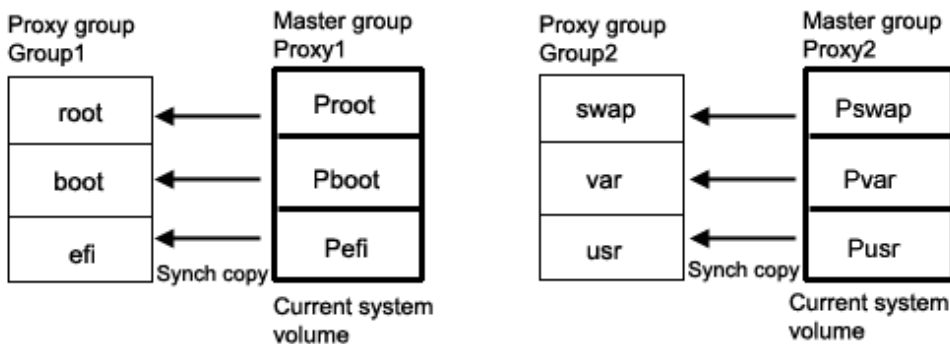
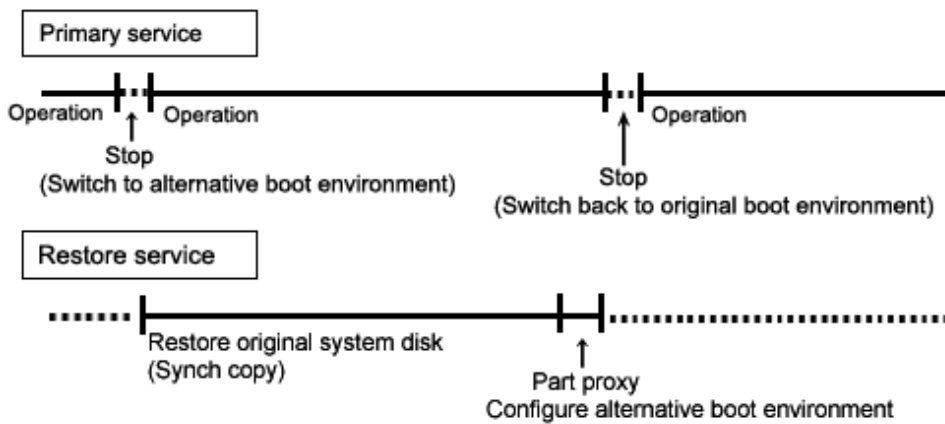


Figure 7.34 Restore Schedule



7.5.4 Summary of Procedure

Figure 7.35 Outline of the Configuration Procedure




Figure 7.36 Outline of the Backup Procedure

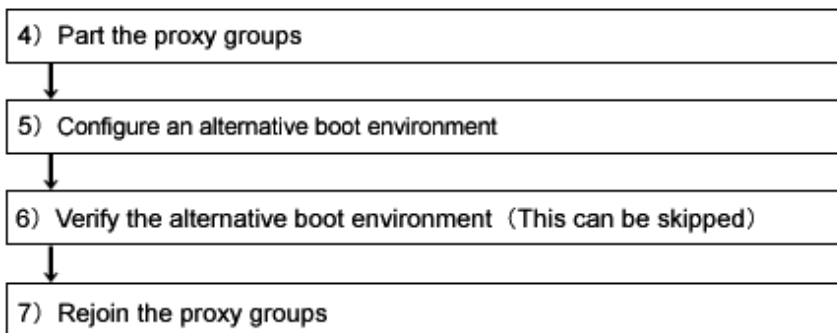



Figure 7.37 Outline of the Restore Procedure



7.5.5 Configuring 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 "[7.5.1 System Configuration](#)."



See

For details on GDS Management View, see "[6.5.2.1 Setting with GDS Management View](#)."

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  
sda=Root1:keep,sdb=Root2:undef,sdc=Var1:keep,sdd=Var2:undef
```

1-3) Connect system disks to groups respectively.

```
# sdxdisk -C -c System -g Group1 -d Root1,Root2 -v 1=root:on,2=boot:on,  
3=efi:on  
# sdxdisk -C -c System -g Group2 -d Var1,Var2 -v 1=swap:on,2=var:on,3=usr:on
```

1-4) Confirm that the mirror definition is complete.

```
# sdxroot -M -c System -d Root1,Var1
```

1-5) Reboot the system.

```
# shutdown -r now
```

1-6) Confirm that mirroring is complete.

Use the mount(8) 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 "[7.5.1 System Configuration](#)."

2-1) Register the disks with the root class.

```
# sdxdisk -M -c System -d sde=Proot1,sdf=Proot2,sdg=Pvar1,sdh=Pvar2
```



See

For details on GDS Management View, see "[8.2.1.2 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 "6.6.1.5 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 "7.5.1 System Configuration."

3-1) Join the proxy groups.

```
# sdxproxy Join -c System -m Group1 -p Proxy1 -a
root=Proot:on,boot=Pboot:on,efi=Pefi:on
# sdxproxy Join -c System -m Group2 -p Proxy2 -a
swap=Pswap:on,var=Pvar:on,usr=Pusr:on
```



See

For details on GDS Management View, see "6.10.1.1 Join."

3-2) 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   boot    ACTIVE
slice  System Group1  Root2   boot    ACTIVE
slice  System Group1  Root1   efi     ACTIVE
slice  System Group1  Root2   efi     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    usr     ACTIVE
slice  System Group2  Var2    usr     ACTIVE
slice  System Proxy1  Proot1  Proot   STOP
slice  System Proxy1  Proot2  Proot   STOP
slice  System Proxy1  Proot1  Pboot   STOP
slice  System Proxy1  Proot2  Pboot   STOP
slice  System Proxy1  Proot1  Pefi    COPY
slice  System Proxy1  Proot2  Pefi    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   Pusr    COPY
slice  System Proxy2  Pvar2   Pusr    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.

7.5.6 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.

- Boot the system in single user mode. (This can be skipped.)
- Exit all active application programs writing in the system disks. (This can be skipped.)
- Execute the sync(1) command to write file system data updated in memory but not yet written to the disks.

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) When the system was started in single user mode in a. of step 4-1), start the GDS services.

```
# /usr/bin/systemctl start fjsvsdx.service
```

4-3) Part the proxy groups.

```
# sdxproxy Part -c System -p Proxy1,Proxy2
```



See

When using GDS Management View, see "Part" in "7.8.1.2 Operation Procedure."

4-4) When the system was booted in single user mode in a. of step 4-1), reboot it in multi-user mode.

4-5) When application programs were exited in b. of step 4-1), launch the application programs.

5) Configuring an alternative boot environment

Enable the system to boot from 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,Pboot,Pefi,Pswap,Pvar,Pusr
# sdxattr -V -c System -v Proot -a mode=rw
# sdxattr -V -c System -v Pboot -a mode=rw
```

```
# sdxattr -V -c System -v Pefi -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 Pusr -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(8) command. If the file system type is xfs, use xfs_repair(8) instead of fsck(8).

```
# sdxvolume -N -c System -v Proot,Pboot,Pefi,Pswap,Pvar,Pusr
# fsck /dev/sfdsk/System/dsk/Proot
# fsck /dev/sfdsk/System/dsk/Pboot
# fsck /dev/sfdsk/System/dsk/Pefi
# fsck /dev/sfdsk/System/dsk/Pvar
# fsck /dev/sfdsk/System/dsk/Pusr
```

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=Root1 Root2
alternative-boot-device=Proot1 Proot2
```

Be sure to keep a copy of the output boot device names in the current boot environment (current-boot-device values) and in the alternative boot environment (alternative-boot-device values).

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,Pboot,Pefi,Pswap,Pvar,Pusr
```



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.

From boot devices displayed in the EFI boot manager's boot option selection screen, select one of the devices in the alternative boot environment output in the message as shown in step 5-3).

```
Root1
Root2
Proot1 <- Select
Proot2
```

6-2) When the system is started in single user mode, restart it in multi-user mode.

If the system was started in single user mode in a. of step 4-1), the alternative boot environment is started in single user mode. In this case, after changing the settings to start the alternative boot environment in multi-user mode, stop the system. Then, execute step 6-1) again. For the settings to start in multi-user mode, see "[A.10 Changing system run level.](#)"

6-3) Confirm that it was booted normally.

Using the mount(8) 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.

6-4) Return to the original boot environment.

From boot devices displayed in the EFI boot manager's boot option selection screen, select one of the devices in the alternative boot environment output in the message as shown in step 5-3).

```
Root1 <- Select
Root2
Proot1
Proot2
```

6-5) 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,Pboot,Pefi,Pswap,Pvar,Pusr
```

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 4-4) or step 6).

7-1) Rejoin the proxy groups.

```
# sdxproxy Rejoin -c System -p Proxy1,Proxy2
```

See

When using GDS Management View, see "[Rejoin](#)" in "[7.8.1.2 Operation Procedure.](#)"

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	boot	ACTIVE
slice	System	Group1	Root2	boot	ACTIVE
slice	System	Group1	Root1	efi	ACTIVE
slice	System	Group1	Root2	efi	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	usr	ACTIVE
slice	System	Group2	Var2	usr	ACTIVE
slice	System	Proxy1	Proot1	Proot	STOP
slice	System	Proxy1	Proot2	Proot	STOP
slice	System	Proxy1	Proot1	Pboot	STOP
slice	System	Proxy1	Proot2	Pboot	STOP
slice	System	Proxy1	Proot1	Pefi	COPY
slice	System	Proxy1	Proot2	Pefi	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	Pusr	COPY
slice	System	Proxy2	Pvar2	Pusr	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).

7.5.7 Restoring

 **Note**

Disk in DISABLE status

When taking this procedure to restore the system disk, check before taking this procedure or check in the procedure 8-2) that no DISABLE disks exist.

If a disk in DISABLE status exists, follow "(1) Disk is in DISABLE status." in "D.1.2 Disk Status Abnormality" to restore the disk in DISABLE status first, and then take the following procedure.

However, when the condition to forcibly remove a class that is described in "(1) Disk is in DISABLE status." in "D.1.2 Disk Status Abnormality" is satisfied, the disk cannot be restored by the following procedure. In this case, follow "(4) System cannot be booted. (Failure of all boot disks)" in "D.1.5 System Disk Abnormality [EFI]" to restore the disk.

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 in the alternative boot environment.

From boot devices displayed in the EFI boot manager's boot option selection screen, select one of the devices in the alternative boot environment output in the message as shown in step 5-3).

```
Root1
Root2
Proot1 <- Select
Proot2
```

8-2) When the system is started in single user mode, restart it in multi-user mode.

If the system was started in single user mode in a. of step 4-1) at "7.5.6 Backing Up" and step 6) was not executed, the alternative boot environment is started in single user mode. In this case, after changing the settings to start the alternative boot environment in multi-user mode, stop the system. Then, execute step 6-1) again. For the settings to start in multi-user mode, see "A.10 Changing system run level."

8-3) Confirm that it was booted normally.

Using the mount(8) 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.

8-4) 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,boot,efi,swap,var,usr
# sdxvolume -R -c System -v root
# sdxvolume -R -c System -v boot
# sdxvolume -R -c System -v efi
# sdxvolume -R -c System -v swap
# sdxvolume -R -c System -v var
# sdxvolume -R -c System -v usr
# 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 "9.1 Breaking Proxy," "9.3.1.2 Removing Volume," "9.3.1.3 Removing Group," and "8.2.1.2 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
```



See

When using GDS Management View, see "9.1 Breaking Proxy."

9-2) Remove the master volumes.

```
# sdxvolume -F -c System -v root,boot,efi,swap,var,usr
# sdxvolume -R -c System -v root
# sdxvolume -R -c System -v boot
# sdxvolume -R -c System -v efi
# sdxvolume -R -c System -v swap
# sdxvolume -R -c System -v var
# sdxvolume -R -c System -v usr
```



See

When using GDS Management View, see "9.3.1.2 Removing Volume."

9-3) If an original system disk crashed, swap the failed disk.

The following is an example of swapping disk Root1 (physical disk sda).



See

When using GDS Management View, see "7.3.1 Setting with GDS Management View."

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 sda.

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,Pboot=boot:on,Pefi=efi:on
# sdxproxy Join -c System -m Proxy2 -p Group2 -a
Pswap=swap:on,Pvar=var:on,Pusr=usr:on
```



See

When using GDS Management View, see "6.10.1.1 Join."

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	boot	STOP
slice	System	Group1	Root2	boot	STOP
slice	System	Group1	Root1	efi	COPY
slice	System	Group1	Root2	efi	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	usr	COPY
slice	System	Group2	Var2	usr	COPY
slice	System	Proxy1	Proot1	Proot	ACTIVE
slice	System	Proxy1	Proot2	Proot	ACTIVE
slice	System	Proxy1	Proot1	Pboot	ACTIVE
slice	System	Proxy1	Proot2	Pboot	ACTIVE
slice	System	Proxy1	Proot1	Pefi	ACTIVE
slice	System	Proxy1	Proot2	Pefi	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	Pusr	ACTIVE
slice	System	Proxy2	Pvar2	Pusr	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.

9-6) Through similar steps as 4) and 5) in "7.5.6 Backing Up" and 8) in this sub-section, part the proxy groups, create an alternative boot environment, and switch to the alternative boot environment.

 **Information**

Canceling system disk mirroring in an alternative boot environment

To cancel system disk mirroring after switching to the alternative boot environment in step 8), perform the following procedure. The following describes the procedure concerning the configuration as below.



10) Breaking up the former boot environment

Break up the former boot environment according to need. Break the master and proxy relationships, remove the master volumes, and remove groups and disks from the master groups as follows.

```
# sdxproxy Break -c System -p Proxy1
# sdxproxy Break -c System -p Proxy2
# sdxvolume -F -c System -v root, boot, efi, swap,var,usr
# sdxvolume -R -c System -v root
# sdxvolume -R -c System -v boot
# sdxvolume -R -c System -v efi
# sdxvolume -R -c System -v swap
# sdxvolume -R -c System -v var
# sdxvolume -R -c System -v usr
# 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
```



See

When using GDS Management View, see "9.1 Breaking Proxy," "9.3.1.2 Removing Volume," "9.3.1.3 Removing Group," and "8.2.1.2 Class Configuration."

11) Unmirroring system disks in an alternative boot environment



See

When using GDS Management View, see "9.4.1 Unmirroring System Disk with GDS Management View."

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 -r now
```

11-5) Verify that the mirroring was canceled normally.

Using the mount(8) command or the sdxinfo command, verify that the system disk mirroring was canceled properly.

11-6) Cancel the system disk management.

```

# sdxvolume -F -c System -v Proot,Pboot,Pefi,Pswap,Pvar,Pusr
# sdxvolume -R -c System -v Proot
# sdxvolume -R -c System -v Pboot
# sdxvolume -R -c System -v Pefi
# sdxvolume -R -c System -v Pswap
# sdxvolume -R -c System -v Pvar
# sdxvolume -R -c System -v Pusr
# 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

```

7.6 Backup (Offline)

This section discusses how to back up the volume data by stopping the services in a system where GDS Snapshot has not been installed.

This procedure enables you to back up the data of the following volumes.

- Shared class volume
- Local class volume
- Volumes (e.g. /opt, /home) other than system volumes (/usr, /var, swap area) among volumes in the root class

The following is an example of backing up and restoring volume Volume1 in class Class1.

1) Stopping the services

1a) With a shared volume used in a cluster application

1a-1) Exit RMS.

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 example, data is backed up to a tape medium of tape device /dev/st0.

- When backing up data with the dd(1) command

```

# dd if=/dev/sfdsk/Class1/dsk/Volume1 of=/dev/st0 bs=32768

```



See

For details on the backup methods, see the manuals of file systems to be backed up and used commands.

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 RMS.

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 the ext4 file system

```
# mount -t ext4 /dev/sfdsk/Class1/dsk/Volume1 /mnt1
```

3b-2) Resume the services.

7.7 Backup (by Slice Detachment)

For mirror volumes, data can be backed up through snapshot by slice detachment during the services operation.

To secure consistency of data in a detached slice, the services must be stopped temporarily when detaching the slice.



See

- For details on snapshots by slice detachment, see "[1.3.8 Snapshots by Slice Detachment](#)."
- For securing consistency of snapshot data, see "[3.17 Ensuring Consistency of Snapshot Data](#)."



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. Additionally, this snapshot is impossible with mirror volumes in the root class.



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/dsk/volume_name. For details, see "[7.10.2.1 Without Using GDS Snapshot](#)."

7.7.1 Using GDS Management View

7.7.1.1 Operation Flow

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

Figure 7.38 Backup Operation (by Slice Detachment)



7.7.1.2 Operation Procedure

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

Information

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.

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].

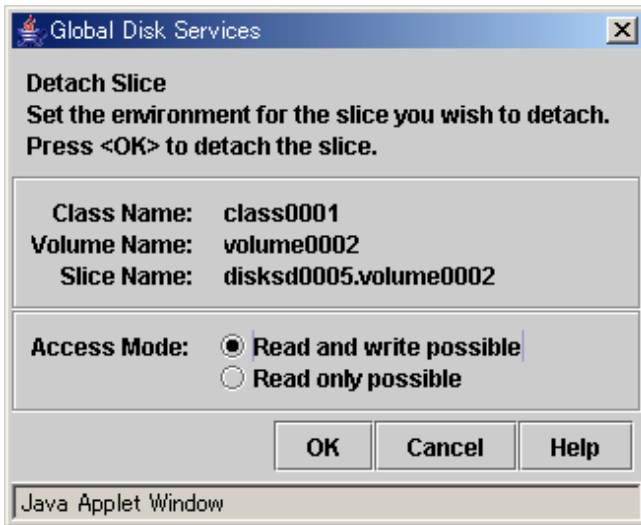
Figure 7.39 Detach Slice



4. Setting the environment of the detaching slice

We will set the environment of the detaching slice.

Figure 7.40 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 canceled.



Note


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 7.41 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 7.42 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 7.43 Attach Slice



The confirmation screen shown below appears.

To continue the process, click <Yes>. Clicking <No> cancels the slice attachment process.

Figure 7.44 Confirming "Attach Slice"



3. Notifying the completion of Attach Slice

Click <OK>, and a message notifying the completion of Attach Slice appears.

Figure 7.45 Notifying Attach Slice Completion



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.

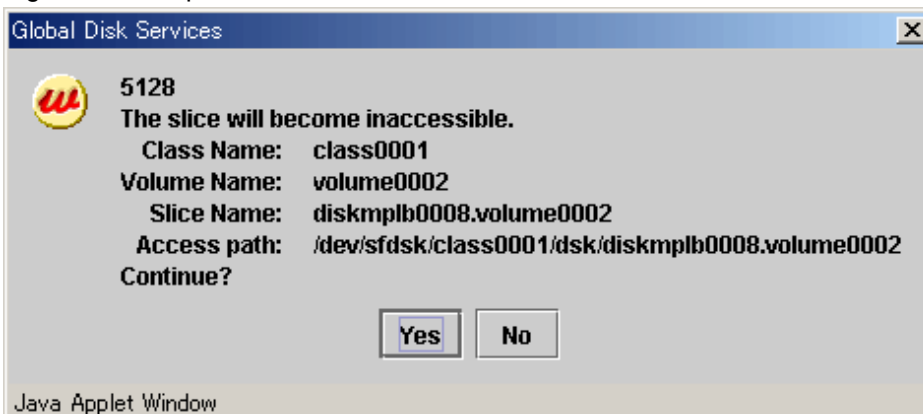
1. Selecting the slice to stop

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 7.46 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.

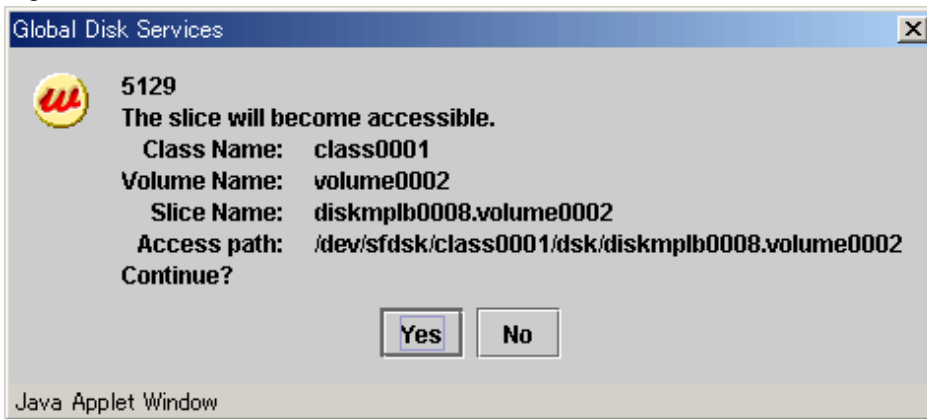
1. Selecting the slice to activate

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 7.47 Activate Slice



To continue the process, click <Yes>. Clicking <No> cancels the slice activation process.

7.7.2 Using Commands

This section describes the procedure for creating backup of a mirror volume making use of a snapshot by detaching a slice.

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
```

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 the ext4 file system

```
# mount -t ext4 /dev/sfdsk/Class1/dsk/Volume1 /mnt1
```

3b-2) Resume the services.

4) 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 the ext4 file system

```
# fsck -t ext4 -y /dev/sfdsk/Class1/dsk/Disk Volume1
```

5) Backing Up

Back up data in the detached slice. In the following example, data is backed up to a tape medium of tape device /dev/st0.



See

For details on the backup methods, see the manuals of file systems to be backed up and used commands.

- When backing up with the dd(1) command

```
# dd if=/dev/sfdsk/Class1/dsk/Disk1.Volume1 of=/dev/st0 bs=32768
```

6) Reattaching the slice

Reattach the detached slice to the mirror volume.

```
# sdxslice -R -c Class1 -d Disk1 -v Volume1
```

7.8 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 services must be stopped temporarily when creating a snapshot (when parting a proxy volume from a master volume).




See

- For details on snapshots by synchronization, see "1.5.1 Snapshot by Synchronization."

- For securing consistency of snapshot data, see "3.17 Ensuring Consistency of Snapshot Data."

Figure 7.48 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.8 Using Advanced Copy Function in Proxy Configuration](#)" and "[3.19 Using Dell EMC TimeFinder or Dell EMC SRDF in Proxy Configuration](#)."

 **Note**

When Using Dell EMC TimeFinder or Dell 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.

Note

Snapshot by Synchronization Use Conditions

See the following sections for points of concern.

- "[2.3.12 Proxy Configuration Preconditions](#)"
- "[2.3.13 Number of Proxy Volumes](#)"
- "[2.3.14 Proxy Volume Size](#)"
- "[2.3.15 Proxy Group Size](#)"

Information

Use Conditions for Copy Functions of Disk Units


See "[A.8 Using Advanced Copy Function in Proxy Configuration](#)" and "[3.19 Using Dell EMC TimeFinder or Dell EMC SRDF in Proxy Configuration](#)."

7.8.1 Using GDS Management View

7.8.1.1 Operation Flow

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

Figure 7.49 Backup Operation (by Synchronization)



7.8.1.2 Operation Procedure

This subsection describes the procedures for backing up volumes through the use of synchronization snapshot of GDS Snapshot. This method requires the following "Proxy Operations."

- Join
- Part
- Rejoin
- Break



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 "7.10.1.2 Operation Procedure."



"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 "3.17 Ensuring Consistency of Snapshot Data."

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 "6.10.1.1 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.




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 7.50 Part



3. Setting the environment for parting the proxy

The Part Proxy dialog box appears.

Figure 7.51 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.

Note

Instant Snapshot by OPC Use Conditions

See the following sections.

- ["A.8 Using Advanced Copy Function in Proxy Configuration"](#)
- ["3.18 Instant Snapshot by OPC"](#)

b. Just Resynchronization Mechanism

Set the mode of Just Resynchronization Mechanism (JRM) for proxies. The default value is "on."

See

For details on JRM for proxies, see ["A.4 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 7.52 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 7.53 Rejoin



3. Confirming

A confirmation dialog box appears asking you whether to rejoin the proxy.

Figure 7.54 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 7.55 Information Message of "Rejoin" Completion



Click <OK> to close the information message window.

Break

If no more backup is to be executed, break the relationship between the master and the proxy.

For "Break Proxy" procedures see "[9.1 Breaking Proxy](#)."

7.8.2 Using Commands

This subsection describes the procedures for backing up volumes through the use of synchronization snapshot of GDS Snapshot.

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
```

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
```

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 being used as a file system, unmount the file system. This example shows the procedure when the mount point is /DATA.

```
# cd /
# umount /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
```

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, mount the file system. This example shows the procedure when the ext4 file system on master volume Volume1 is mounted on mount point /DATA.

```
# mount -t ext4 /dev/sfdsk/Class1/dsk/Volume1 /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.



For details on the backup methods, see the manuals of file systems to be backed up and used commands.

6a) When backing up data with the dd(1) command

```
# dd if=/dev/sfdsk/Class1/dsk/Volume2 of=/dev/st0 bs=32768
```

6b) When backing up the ext4 file system with the tar(1) command

6b-1) Before mounting

Check and repair consistency of the ext4 file system on the proxy volume with the fsck(8) command. When the file system on the master volume was unmounted in step 3b-2), skip this step.

```
# fsck -t ext4 /dev/sfdsk/Class1/dsk/Volume2
```

6b-2) Mounting the snapshot

Mount the ext4 file system on proxy volume Volume2 on temporary mount point /DATA_backup.

```
# mkdir /DATA_backup
# mount -t ext4 /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/st0 with the tar(1) command.

```
# cd /DATA_backup
# tar cvf /dev/st0 .
```

6b-4) Unmounting the snapshot

Unmount the file system mounted in step 6b-2).

```
# cd /
# umount /DATA_backup
# rmdir /DATA_backup
```

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
```

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.

8) Breaking the proxy volume

When no more online backup is 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
```

7.9 Backup (by OPC)


If GDS Snapshot is installed in a system where a disk array unit with the OPC or QuickOPC 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 details on snapshots by OPC, see "[1.5.3 Instant Snapshot by OPC.](#)"
- For securing consistency of snapshot data, see "[3.17 Ensuring Consistency of Snapshot Data.](#)"

Figure 7.56 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 to the disk array unit and may affect services using master volumes.

Note

Instant Snapshot by OPC

For use conditions for instant snapshot by OPC, see "[3.18 Instant Snapshot by OPC.](#)" If proxy volumes are in mirroring configuration with the multiplicity of two and higher, the OPC function is unavailable for copying data from master volumes to the proxy volumes. For details, see "[A.8 Using Advanced Copy Function in Proxy Configuration.](#)"

Note

Instant Snapshot by OPC Use Conditions

See the following sections for points of concern.


- "2.3.12 Proxy Configuration Preconditions"
- "2.3.13 Number of Proxy Volumes"
- "2.3.14 Proxy Volume Size"
- "2.3.15 Proxy Group Size"
- "3.18 Instant Snapshot by OPC"
- A.8 Using Advanced Copy Function in Proxy Configuration"

7.9.1 Using GDS Management View

7.9.1.1 Operation Flow

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

Figure 7.57 Backup Operation (by OPC)



7.9.1.2 Operation Procedure

This subsection describes the procedures for backing up volumes through the use of OPC snapshot of GDS Snapshot. This method requires the following "Proxy Operations."

- Relate
- Update
- Break

Note

When using QuickOPC function

- If a QuickOPC session exists between the master volume and the proxy volume, QuickOPC function is used for updating the proxy. The presence or absence of a QuickOPC session can be checked by using the `sdxinfo -S -e` long command. When "QOPC" is displayed in the FUNC field, QuickOPC session exists.
 - Use the `sdxproxy` command for the first backup operation by using QuickOPC function. If you use the GDS Management View for the first backup operation, OPC function is used for the operation instead of QuickOPC function.
-

See

- For details on the `sdxproxy` command, see "[B.2.1 sdxproxy - Proxy Object Operations.](#)"
-

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 "[3.17 Ensuring Consistency of Snapshot Data.](#)"

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 "[6.10.1.2 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.

Note


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 7.58 Update



3. Setting the environment for updating the proxy

The Update Proxy screen appears.

Figure 7.59 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 7.60 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 7.61 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 "[9.1 Breaking Proxy](#)."

7.9.2 Using Commands

This subsection describes the procedures for backing up volumes through the use of OPC snapshot of GDS Snapshot.

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
```

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
```

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 being used as a file system, unmount the file system. This example shows the procedure when the mount point is /DATA.

```
# cd /  
# umount /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.

- For using OPC function

```
# sdxproxy Update -c Class1 -p Volume2 -e instant
```



When a QuickOPC session exists

When a QuickOPC session exists between master volume and proxy volume, the above command can copy only the differences between master and proxy. If you want to copy the whole master volume, stop the QuickOPC session. Then, specify -e OPC option and execute the sdxproxy Update command.

```
# sdxproxy Cancel -c Class1 -p Volume2  
# sdxproxy Update -c Class1 -p Volume2 -e instant,OPC
```

- For using QuickOPC function

```
# sdxproxy Update -c Class1 -p Volume2 -e instant,QOPC
```

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.

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, mount the file system. This example shows the procedure when the ext4 file system on master volume Volume1 is mounted on mount point /DATA.

```
# mount -t ext4 /dev/sfdsk/Class1/dsk/Volume1 /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
```

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.



For details on the backup methods, see the manuals of file systems to be backed up and used commands.

8a) When backing up data with the dd(1) command

```
# dd if=/dev/sfdsk/Class1/dsk/Volume2 of=/dev/st0 bs=32768
```

8b) When backing up the ext4 file system with the tar(1) command

8b-1) Before mounting

Check and repair consistency of the ext4 file system on the proxy volume with the fsck(8) command. When the file system on the master volume was unmounted in step 3b-2), skip this step.

```
# fsck -t ext4 /dev/sfdsk/Class1/dsk/Volume2
```

8b-2) Mounting the snapshot

Mount the ext4 file system on proxy volume Volume2 on temporary mount point /DATA_backup.

```
# mkdir /DATA_backup
# mount -t ext4 /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/st0 with the tar(1) command.

```
# cd /DATA_backup
# tar cvf /dev/st0 .
```


8b-4) Unmounting the snapshot

Unmount the file system mounted in step 8b-2).

```
# cd /  
# umount /DATA_backup  
# rmdir /DATA_backup
```

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
```

10) Re-executing online backup

Re-execute step 3) through 9) for performing online backup again.

11) 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
```


7.10 Restoring

7.10.1 Using GDS Management View

7.10.1.1 Operation Flow

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

Figure 7.62 Restore Operation



 See

- To restore with proxy volumes, see "[1.5.4 Instant Restore](#)."

7.10.1.2 Operation Procedure

This subsection describes the procedures for restoring volumes through use of GDS Snapshot proxy volumes. This method requires the following "Proxy Operations."

- Restore



System Volume Restoration [EFI]

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 "[7.4 Backing Up and Restoring System Disk](#)" or "[7.10.1.2 Operation Procedure](#)."

Stop services using a volume to be restored (a master volume), stop the master volume, and then perform the following procedures.

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.



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.



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 7.63 Restore



3. Setting the environment for restoring the master

The Restore Master screen appears.

Figure 7.64 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."



For the OPC function use conditions, see the following sections.

- "3.18 Instant Snapshot by OPC"
- "A.8 Using Advanced Copy Function in Proxy Configuration"



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."



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 7.65 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 7.66 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 7.67 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 7.68 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 "[9.1 Breaking Proxy](#)."

7.10.2 Using Commands

7.10.2.1 Without Using GDS Snapshot

This subsection explains how to restore the volume data in a system where GDS Snapshot has not been installed.

This procedure enables you to restore the data of the following volumes.

- Shared class volume
- Local class volume
- Volumes (e.g. /opt, /home) other than system volumes (/usr, /var, swap area) among volumes in the root class

The following is an example of backing up and restoring volume Volume1 in class Class1.

1) Stopping the services

1a) With a shared volume used by a cluster application.

1a-1) Exit RMS.

1a-2) Activate the volume on a node on which restoration 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) Restoring

Restore the volume data. The following shows an example of restoring data from a tape medium of tape device /dev/st0.



See

For details on the restore methods, see the manuals of file systems to be restored and used commands.

- When restoring data with the dd(1) command

```
# dd if=/dev/st0 of=/dev/sfdsk/Class1/dsk/Volume1 bs=32768
```

3) Resuming the services

3a) With a shared volume used by a cluster application

3a-1) Inactivate the volume on the node where restoration was conducted.

```
# sdxvolume -F -c Class1 -v Volume1
```

3a-2) Launch RMS.

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 the ext4 file system

```
# mount -t ext4 /dev/sfdsk/Class1/dsk/Volume1 /mnt1
```

3b-2) Resume the services.

7.10.2.2 Using GDS Snapshot

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).

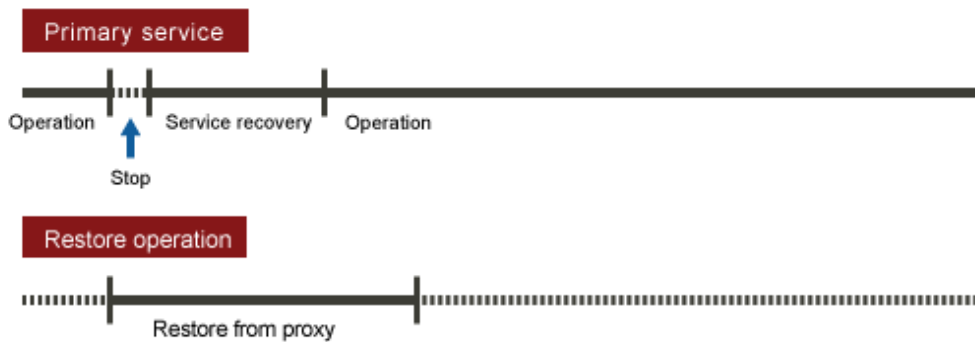
The master volume must be stopped temporarily for restoration, but you may start the master volume and make it accessible immediately after the restore is initiated without waiting until the copying is complete.

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 "7.10.2.1 Without Using GDS Snapshot."

Figure 7.69 Schedule for Instant Restore



Note

Instant Restore with Disk Unit's Copy Functions

To restore, only OPC is available, but (R)EC, Dell EMC TimeFinder or Dell EMC SRDF is unavailable. For details, see "A.8 Using Advanced Copy Function in 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
```

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
```

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.

3b) When the OPC function is available

Restore data from proxy volume Volume2 back to master volume Volume1.

- For using OPC function

```
# sdxproxy Restore -c Class1 -p Volume2 -e instant
```



When a QuickOPC session exists

When a QuickOPC session exists between master volume and proxy volume, the above command can copy only the differences between master and proxy. If you want to copy the whole proxy volume, stop the QuickOPC session. Then, specify -e OPC option and execute the sdxproxy Restore command.


```
# sdxproxy Cancel -c Class1 -p Volume2
# sdxproxy Restore -c Class1 -p Volume2 -e instant,OPC
```

- For using QuickOPC function

```
# sdxproxy Restore -c Class1 -p Volume2 -e instant,QOPC
```

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.



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.

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 restoration 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 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 this 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 a file system, mount the file system. In this example, the mount point is /DATA.

- For the ext4 file system

```
# mount -t ext4 /dev/sfdsk/Class1/dsk/Volume1 /DATA
```

4b-3) Resume the services 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 event 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 "7.8.2 Using Commands."

Note

When Using (R)EC, Dell EMC TimeFinder, Dell 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.

7.11 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, ETERNUS Disk storage system provides the Advanced Copy function and Dell EMC storage units 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 `sda` and `sdb` are registered with a shared class named `Class1` and mirrored, and a mirror volume named `Volume1` is used for services.

7.11.1 Configuring Environment

Note

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. For details on the resource registration, refer to "[6.3 Shared Disk Resource Registration](#)."

1) Creating an application volume

Create an application mirror volume onto application disks sda and sdb. The following settings are necessary on Node1 and Node2 in the primary domain.

1-1) Registering disks

Register disks sda and sdb 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  
sda=Disk1,sdb=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 sda and sdb, as copy sources, respectively with disks sdc and sdd, as copy destinations, with copy functions of the disk units.

See

For the synchronization methods, see the manuals of disk units for information regarding copy functions.

Note

Backup disks must be equivalent in size application disks to be backed up.

7.11.2 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."

```
# sdxinfol -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)



Information

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 "[3.17 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 "[D.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.



See

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 sdc and sdd from application disks sda and sdb.



See

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 sdc and sdd. 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 sda and sdb to sdc and sdd, which are the physical disk names of the backup disks, respectively.

```
# sdxconfig Convert -e replace -c Class1 -p sda=sdc,sdb=sdd \  
-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 sdc and sdd 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 sdc and sdd 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/st0.



For details on the backup methods, see the manuals of file systems to be backed up and used commands.

6a) When backing up data with the dd(1) command

```
# dd if=/dev/sfdsk/Class2/dsk/Volume1 of=/dev/st0 bs=32768
```

6b) When backing up the ext4 file system with the tar(1) command

6b-1) Check and repair consistency of the ext4 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 -t ext4 /dev/sfdsk/Class2/dsk/Volume1
```

6b-2) Mount the ext4 file system on backup volume Volume1 on /mnt1, a temporary mount point, in read only mode.

```
# mkdir /mnt1
# mount -t ext4 -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/st0 .
```

6b-4) Unmount the file system mounted in step 6b-2).

```
# cd /
# umount /mnt1
# rmdir /mnt1
```

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 sda and sdb, as copy sources, respectively with backup disks sdc and sdd, 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.
.....

7.11.3 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 sdc and sdd, as copy sources, respectively back to application disks sda and sdb, 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 sdc and sdd from application disks sda and sdb.



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 "7.11.2 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 -e chkps
# shutdown -r now
```

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 "[7.11.2 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.

```
# sdxvolume -N -c Class1 -v Volume1
```

17) Resuming the services

When the file system on application volume Volume1 was unmounted in step 8) of "[7.11.2 Backing Up](#)," unmount it again.

Start the applications using Volume1.

7.12 Backing Up and Restoring through External Server

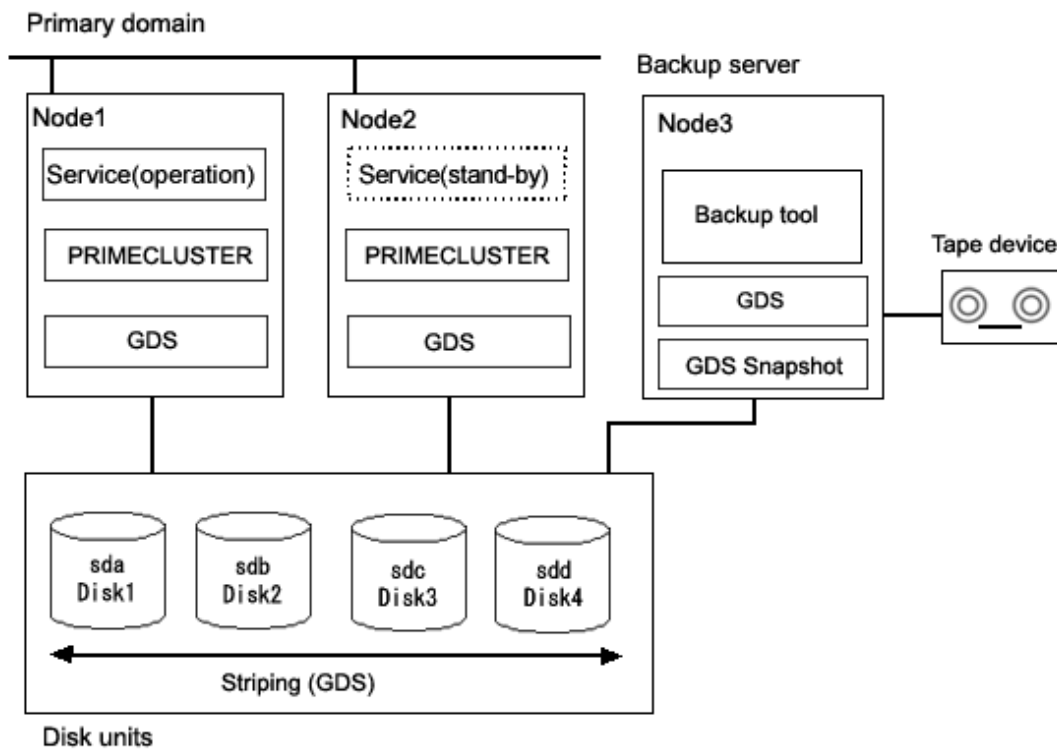
7.12.1 Backing Up and Restoring Logical Volume with No Replication

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 "[7.12.2 Backing Up and Restoring through Snapshot by Slice Detachment](#)" is recommended.

7.12.1.1 System Configuration

Figure 7.70 System Configuration



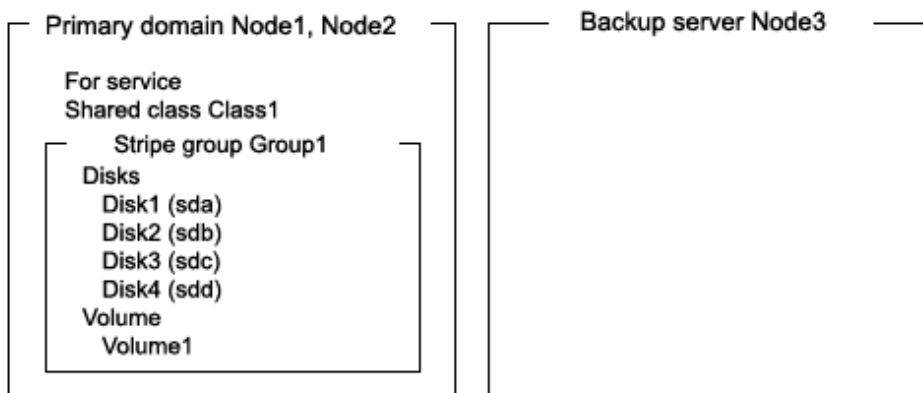
GDS: Global Disk Services
 GDS Snapshot: Global Disk Services Snapshot

Note

Physical Device Name

Different physical device names (such as sda) may be assigned to the identical physical disk in the primary domain and the backup server.


Figure 7.71 Object Configuration in Normal Operation



7.12.1.2 Summary of Backup

Backups can be created while the services are stopped and the application volume is not in use.

Figure 7.72 Backup



GDS: Global Disk Services

Figure 7.73 Object Configuration When Backing Up

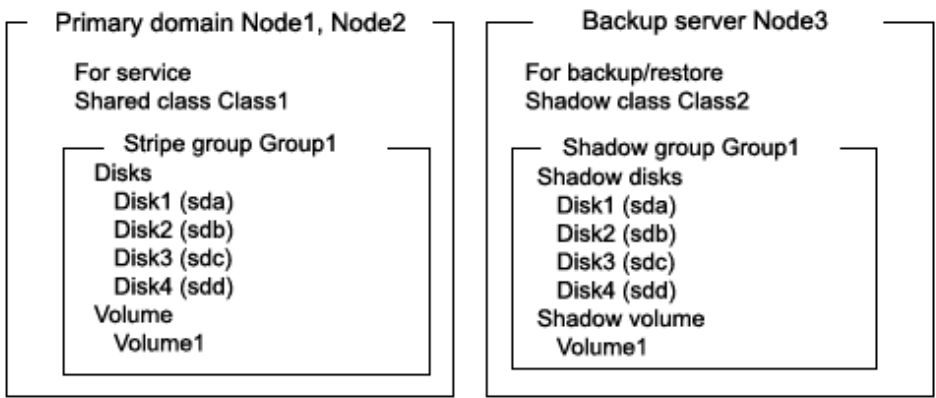



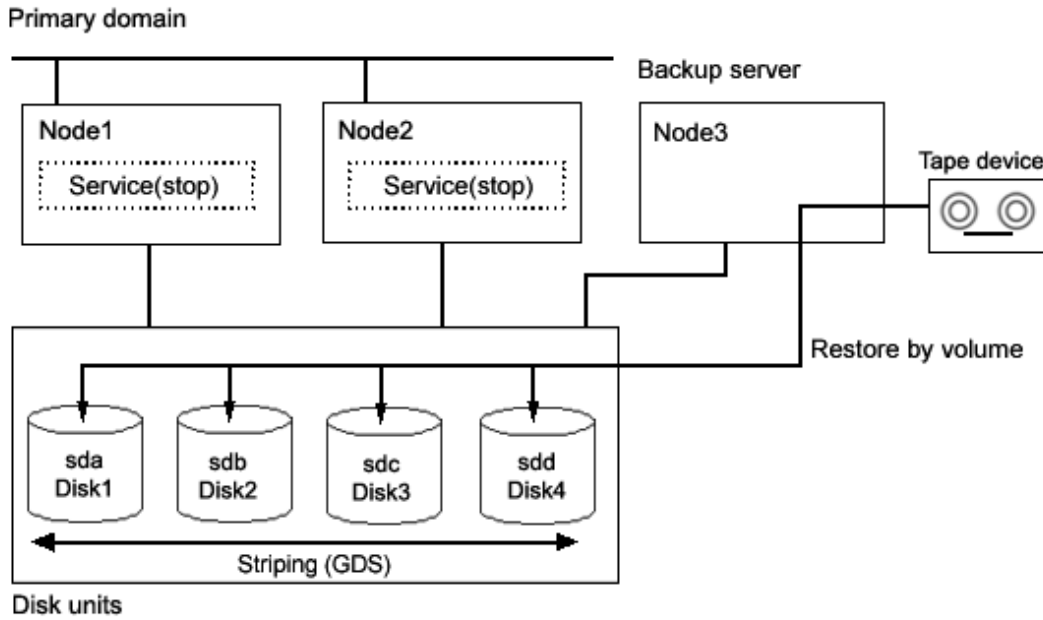
Figure 7.74 Backup Schedule



7.12.1.3 Summary of Restore

If volume data is damaged, it can be restored from tape. Data can be restored while the services are stopped and the application volume is not in use.

Figure 7.75 Restore



GDS: Global Disk Services

Figure 7.76 Object Configuration When Restoring

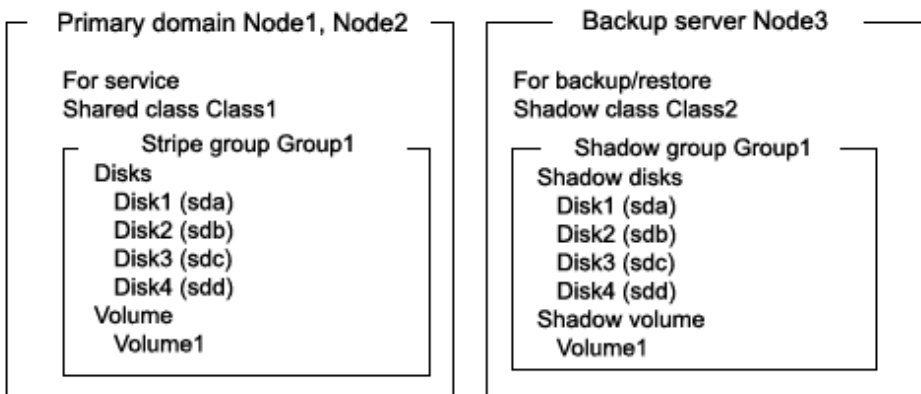



Figure 7.77 Restore Schedule



7.12.1.4 Summary of Procedure


Figure 7.78 Outline of the Configuration Procedure

1) Create an application volume	(Primary domain)
---------------------------------	------------------

Figure 7.79 Outline of the Backup Procedure

2) Stop the service	(Primary domain)
↓	
3) Stop the application volume	(Primary domain)
↓	
4) View the configuration of the application volume	(Primary domain)
↓	
5) Create a shadow volume for backup	(Backup server)
↓	
6) Backup to tape	(Backup server)
↓	
7) Remove the shadow volume	(Backup server)
↓	
8) Resume the service	(Primary domain)

Figure 7.80 Outline of the Restore Procedure



7.12.1.5 Configuring Environment



Note

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. For details on the resource registration, see "6.3 Shared Disk Resource Registration."

1) Creating an application volume

Create a stripe volume used for the services on disks sda, sdb, sdc, and sdd. The following settings are necessary on Node1 or Node2 in the primary domain.

1-1) Registering disks

Register disks sda, sdb, sdc, and sdd 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
sda=Disk1,sdb=Disk2,cdc=Disk2,sdd=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
```

7.12.1.6 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. Check the underlined parts.

```
# 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     sda        8380800    Node1:Node2    ENABLE
disk     Disk2    stripe    Class1     Group1     sdb        8380800    Node1:Node2    ENABLE
disk     Disk3    stripe    Class1     Group1     sdc        8380800    Node1:Node2    ENABLE
disk     Disk4    stripe    Class1     Group1     sdd        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.

Check the underlined parts.

```
# sdxinfo -G -c Class1 -o Group1 -e long
OBJ      NAME      CLASS     DISKS              BLKS      FREEBLKS    SPARE  MASTER  TYPE      WIDTH  ACTDISK
-----
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 backup server Node3.

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 sda, sdb, sdc, and sdd with shadow class Class2, and name them Disk1, Disk2, Disk3, and Disk4 respectively.

```
# sdxshadowdisk -M -c Class2 -d sda=Disk1,sdb=Disk2,sdc=Disk3,sdd=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. Check the underlined parts.

```
# sdxinfo -c Class2
OBJ      NAME      TYPE      SCOPE      SPARE
-----
class    Class2    local     Node3      0

OBJ      NAME      TYPE      CLASS      GROUP      DEVNAM      DEVBLKS      DEVCONNECT      STATUS
-----
disk     Disk1    stripe    Class2     Group1     sda          8380800      Node3           ENABLE
disk     Disk2    stripe    Class2     Group1     sdb          8380800      Node3           ENABLE
disk     Disk3    stripe    Class2     Group1     sdc          8380800      Node3           ENABLE
disk     Disk4    stripe    Class2     Group1     sdd          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. Check the underlined parts.

```
# sdxinfo -G -c Class2 -o Group1 -e long
OBJ      NAME      CLASS      DISKS          BLKS      FREEBLKS      SPARE      MASTER      TYPE      WIDTH      ACTDISK
-----
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, back up data in shadow volume Volume1 to a tape medium of tape device /dev/st0.



For details on the backup method, see the manuals of file systems to be backed up and used commands.



In a GFS Shared File System

Back up through the method as described in step 6a).

6a) When backing up data with the dd(1) command

```
# dd if=/dev/sfdsk/Class2/dsk/Volume1 of=/dev/st0 bs=32768
```

6b) When backing up the ext4 file system with the tar(1) command

6b-1) Mount the ext4 file system on shadow volume Volume1 on /mnt1, a temporary mount point, in read only mode.

```
# mkdir /mnt1
# mount -t ext4 -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/st0 .
```

6b-3) Unmount the file system mounted in step 6b-1).

```
# cd /
# umount /mnt1
# rmdir /mnt1
```

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) Resuming the services

When the file system on application volume Volume1 was unmounted in step 2), mount it again.

Start the applications using Volume1.

7.12.1.7 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 "[D.1.3 Volume Status Abnormality](#)". Check the underlined parts.

```
# 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     sda          8380800      Node1:Node2     ENABLE
disk    Disk2    stripe   Class1     Group1     sdb          8380800      Node1:Node2     ENABLE
disk    Disk3    stripe   Class1     Group1     sdc          8380800      Node1:Node2     ENABLE
disk    Disk4    stripe   Class1     Group1     sdd          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. Check the underlined parts.

```
# sdxinfo -G -c Class1 -o Group1 -e long
```

OBJ	NAME	CLASS	DISKS	BLKS	FREEBLKS	SPARE	MASTER	TYPE	WIDTH	ACTDISK
group	Group1	Class1	Disk1:Disk2:Disk3:Disk4	32964608	31850496	*	*	stripe	<u>256</u>	*

12) Creating a shadow volume for restoration

On backup server Node3, create a volume for restoration (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 restoration 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).

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).

12-1) Registering shadow disks

Register disks sda, sdb, sdc, and sdd with shadow class Class2, and name them Disk1, Disk2, Disk3, and Disk4 respectively.

```
# sdxshadowdisk -M -c Class2 -d sda=Disk1,sdb=Disk2,sdc=Disk3,sdd=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 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 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. Check the underlined parts.

```
# sdxinfo -c Class2
OBJ      NAME      TYPE      SCOPE      SPARE
-----
class    Class2    local     Node3      0

OBJ      NAME      TYPE      CLASS      GROUP      DEVNAM      DEVBLKS      DEVCONNECT      STATUS
-----
disk     Disk1    stripe    Class2     Group1     sda         8380800      Node3           ENABLE
disk     Disk2    stripe    Class2     Group1     sdb         8380800      Node3           ENABLE
disk     Disk3    stripe    Class2     Group1     sdc         8380800      Node3           ENABLE
disk     Disk4    stripe    Class2     Group1     sdd         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  ACTDISK
-----
group Group1 Class2  Disk1:Disk2:Disk3:Disk4  32964608  31850496  *      *       stripe  256    *
```

13) 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/st0.



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 with the dd(1) command

```
# dd if=/dev/st0 of=/dev/sfdsk/Class2/dsk/Volume1 bs=32768
```

13b) When restoring the ext4 file system with the tar(1) command

13b-1) Create the ext4 file system on Volume1.

```
# mkfs -t ext4 /dev/sfdsk/Class2/dsk/Volume1
```

13b-2) Mount the ext4 file system on shadow volume Volume1 on /mnt1, a temporary mount point.

```
# mkdir /mnt1
# mount -t ext4 /dev/sfdsk/Class2/dsk/Volume1 /mnt1
```

13b-3) Restore data held in the file system from tape.

```
# cd /mnt1
# tar xvf /dev/st0
```

13b-4) Unmount the file system mounted in step 13b-3).

```
# 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.

15-1) Activating the application volume

Activate application volume Volume1.

```
# sdxvolume -N -c Class1 -v Volume1
```

15-2) Resuming the services

When the file system on application volume Volume1 was unmounted in step 9), mount it again.

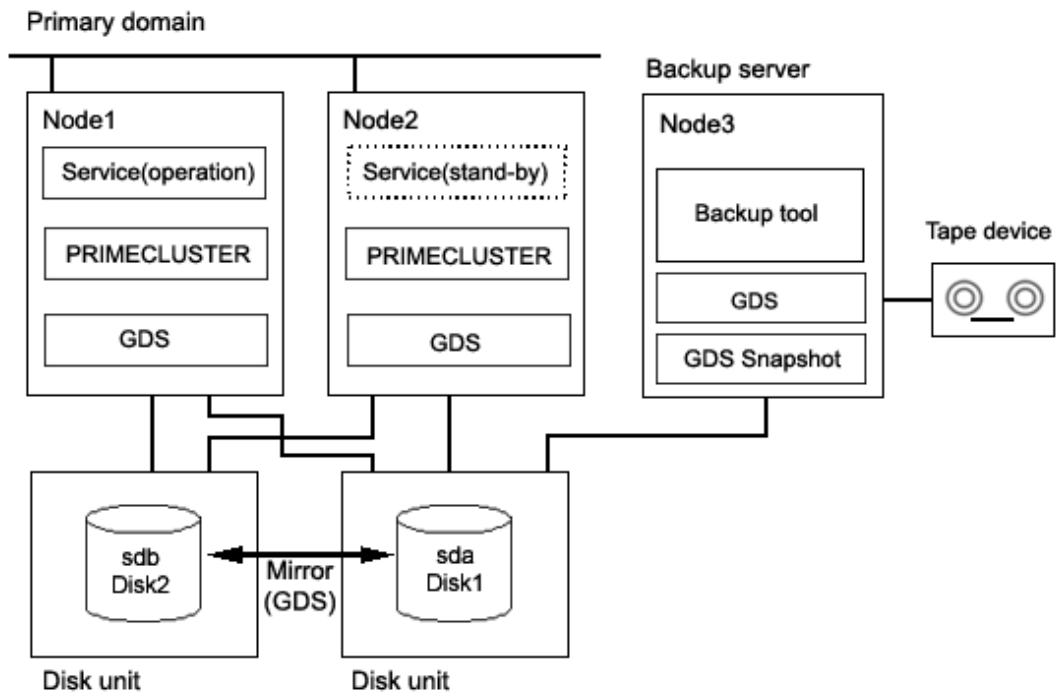
Start the applications using Volume1.

7.12.2 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.

7.12.2.1 System Configuration

Figure 7.81 System Configuration




GDS: Global Disk Services
 GDS Snapshot: Global Disk Services Snapshot

 **Note**

Physical Device Name

Different physical device names (such as sda) may be assigned to the identical physical disk in the primary domain and the backup server.

Figure 7.82 Object Configuration in Normal Operation



7.12.2.2 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 "3.17 Ensuring Consistency of Snapshot Data."

Figure 7.83 Backup




Figure 7.84 Object Configuration When Backing Up





Figure 7.85 Backup Schedule

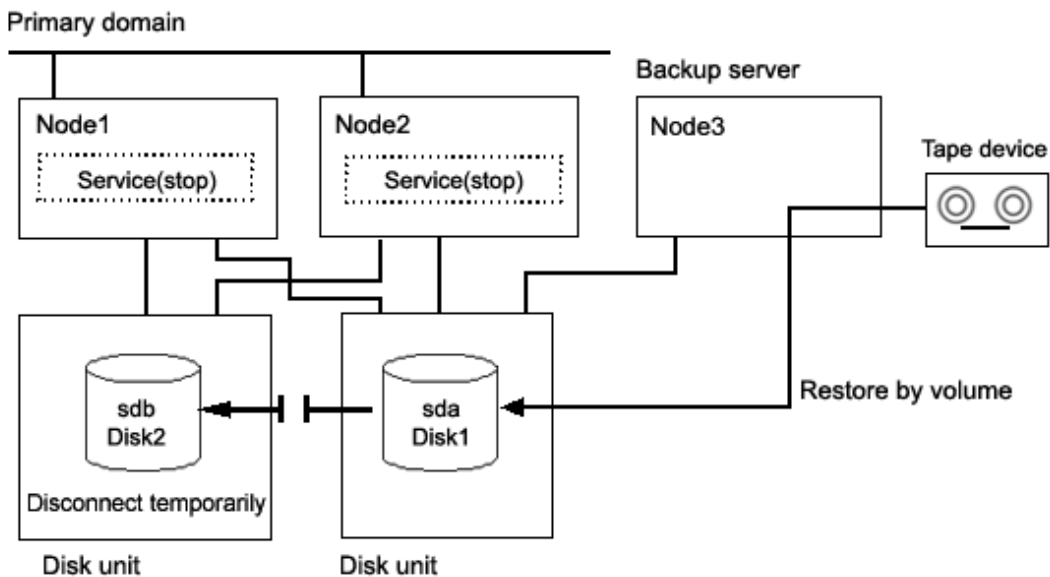


7.12.2.3 Summary of Restore

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 7.86 Restore



Information

In this configuration access cannot be gained from backup server Node3 to disk sdb. Therefore, after data is restored from tape back to sda while sdb is detached temporarily, resynchronization copying from sda to sdb must be performed by reattaching sdb. When access can be

gained from Node3 to both sda and sdb, it is not required that sdb be detached temporarily since data can be restored from tape back to both sda and sdb. For details on this restore method, see "7.12.1 Backing Up and Restoring Logical Volume with No Replication."

Figure 7.87 Object Configuration When Restoring





Figure 7.88 Restore Schedule



7.12.2.4 Summary of Procedure

Figure 7.89 Outline of the Configuration Procedure

1) Create an application volume	(Primary domain)
---------------------------------	------------------

Figure 7.90 Outline of the Backup Procedure

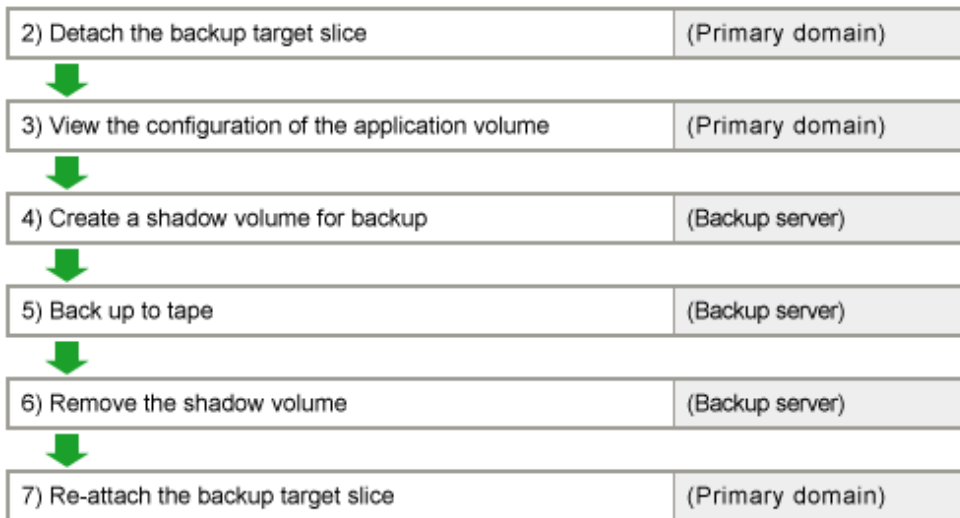



Figure 7.91 Outline of the Restore Procedure



7.12.2.5 Configuring Environment



Note

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. For details on the resource registration, see "6.3 Shared Disk Resource Registration."

1) Creating an application volume

Create a mirror volume used for the services on disks sda and sdb. The following settings are necessary on Node1 or Node2 in the primary domain.

1-1) Registering disks

Register disks sda and sdb 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
sda=Disk1,sdb=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
```

7.12.2.6 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 "3.17 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 Slice

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. Check the underlined parts.

```
# 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     sda         8380800      Node1:Node2     ENABLE
disk     Disk2     mirror    Class1     Group1     sdb         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 sda 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 sda with shadow class Class2, and name it Disk1.

```
# sdxshadowdisk -M -c Class2 -d sda=Disk1
```

Point

- The disk name must correspond to the disk name assigned to disk sda 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.

Check the underlined parts.

```
# 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  sda      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/st0.



For details on the backup method, see the manuals of file systems to be backed up and used command.

5a) When backing up data with the dd(1) command

```
# dd if=/dev/sfdsk/Class2/dsk/Volume1 of=/dev/st0 bs=32768
```

5b) When backing up the ext4 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 ext4 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 -t ext4 -y /dev/sfdsk/Class2/dsk/Volume1
```

5b-3) Mount the ext4 file system on shadow volume Volume1 on /mnt1, a temporary mount point, in the read only mode.

```
# mkdir /mnt1
# mount -t ext4 -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/st0 .
```

5b-5) Unmount the file system mounted in step 5b-3).

```
# cd /
# umount /mnt1
# rmdir /mnt1
```

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
```

7.12.2.7 Restoring



Information

If access can be gained from backup server Node3 to all of the disks constituting Volume1 (sda and sdb) data can be restored from tape back to both sda and sdb 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 (sda and sdb) 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 "[D.1.3 Volume Status Abnormality](#)" and "[D.1.1 Slice Status Abnormality](#)."

Check the underlined parts.

```
# 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 sda      8380800 Node1:Node2  ENABLE
disk Disk2  mirror Class1 Group1 sdb      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 (sda and sdb) you must confirm that all of the slices of Volume1 are in STOP status.

12) Creating a shadow volume for restoration

On backup server Node3, create a volume for restoration (shadow volume) on disk sda. The following settings are necessary on backup server Node3. A shadow volume for restoration 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 (sda and sdb) a shadow volume for restoration must be created in the same configuration as Volume1. Under these circumstances, those shadow volumes for restoration and backup are not common.

12-1) Registering a shadow disk

Register disk sda with shadow class Class2, and name it Disk1.

```
# sdxshadowdisk -M -c Class2 -d sda=Disk1
```

Point

- The disk name must correspond to the disk name assigned to sda 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 (sda and sdb) all of those disks (sda and sdb) 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 (sda and sdb) all of those disks (sda and sdb) 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.

Check the underlined parts.

```
# 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   sda      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/st0.



See

For details on the restore method, see the manuals of file systems to be restored and used commands.

13a) When restoring data with the dd(1) command

```
# dd if=/dev/st0 of=/dev/sfdsk/Class2/dsk/Volume1 bs=32768
```

13b) When restoring the ext4 file system with the tar(1) command

13b-1) Create the ext4 file system to shadow volume Volume1.

```
# mkfs -t ext4 /dev/sfdsk/Class2/dsk/Volume1
```

13b-2) Mount the ext4 file system on shadow volume Volume1 on /mnt1, a temporary mount point.

```
# mkdir /mnt1
# mount -t ext4 /dev/sfdsk/Class2/dsk/Volume1 /mnt1
```

13b-3) Restore data held in the file system from tape.

```
# cd /mnt1
# tar xvf /dev/st0
```

13b-4) Unmount the file system mounted in step 13b-3).

```
# 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 (sda and sdb) all of the disks registered with shadow class Class2 in step 12) (sda and sdb) 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), 15-3), 15-4) (confirming that the synchronization copying is complete), and 15-2).
.....

15-1) Activating the application volume

Activate application volume Volume1.

```
# sdxvolume -N -c Class1 -v Volume1
```

15-2) Resuming the services

When the file system on application volume Volume1 was unmounted in step 8), mount it again.

Start the applications using Volume1.

15-3) 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.

15-4) 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
```

7.12.3 Backing Up and Restoring Using Snapshots from 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.


7.12.3.1 System Configuration

Figure 7.92 System Configuration



GDS: Global Disk Services
 GDS Snapshot: Global Disk Services Snapshot

Figure 7.93 Object Configuration in Normal Operation



7.12.3.2 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 "[3.17 Ensuring Consistency of Snapshot Data](#)."

Figure 7.94 Backup

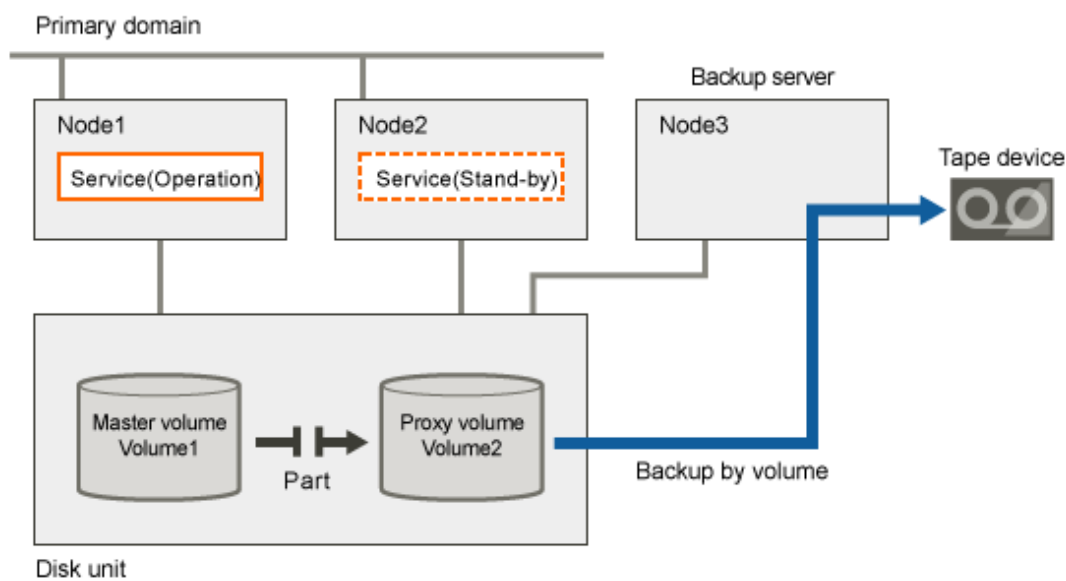


Figure 7.95 Object Configuration When Backing Up





Figure 7.96 Backup Schedule



7.12.3.3 Summary of Restore from 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 7.97 Restore from a Proxy Volume




Figure 7.98 Object Configuration When Restoring from a Proxy Volume


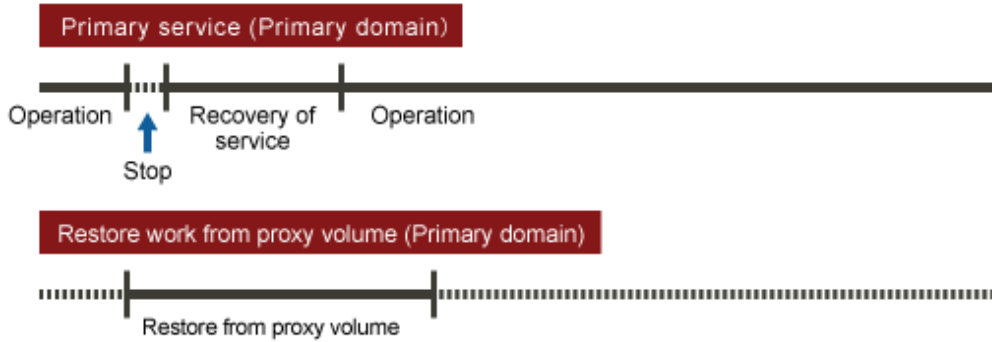


Figure 7.99 Schedule to Restore a Proxy Volume




7.12.3.4 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 7.100 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 7.101 Object Configuration When Restoring from Tape





Figure 7.102 Schedule to Restore from Tape



7.12.3.5 Summary of Procedure

Figure 7.103 Outline of the Configuration Procedure




Figure 7.104 Outline of the Backup Procedure

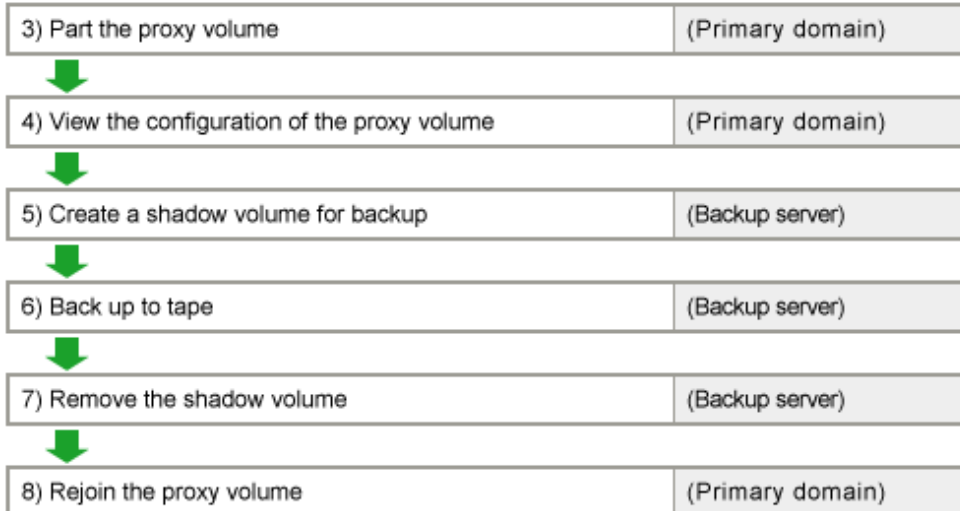


Figure 7.105 Outline of the Procedure for Restoring from a Proxy Volume





Figure 7.106 Outline of the Procedure for Restoring from Tape



7.12.3.6 Configuring Environment



Note

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. For details on the resource registration, see "[6.3 Shared Disk Resource Registration](#)."

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 sda and sdb 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
sda=Disk1,sdb=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 sdc and creates mirror volume Volume2.

```
# sdxdisk -M -c Class1 -d sdc=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.

7.12.3.7 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 "[3.17 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 "[D.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
```

Note

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
```

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 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/st0`.

See

For details on the backup method, see the manuals of file systems to be backed up and used commands.

6a) When backing up data with the `dd(1)` command

```
# dd if=/dev/sfdsk/Class2/dsk/Volume2 of=/dev/st0 bs=32768
```

6b) When backing up the ext4 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 ext4 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 -t ext4 -y /dev/sfdsk/Class2/dsk/Volume2
```

6b-3) Mount the ext4 file system on shadow volume Volume2 on `/mnt1`, a temporary mount point, in the read only mode.

```
# mkdir /mnt1
# mount -t ext4 -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/st0 .
```

6b-5) Unmount the file system mounted in step 6b-3).

```
# cd /
# umount /mnt1
# rmdir /mnt1
```

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
```

7.12.3.8 Restoring from 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) Resuming the services

When the file system on master volume Volume1 was unmounted in step 9), mount it again.

Start the applications using Volume1.

12-3) 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 "7.12.3.7 Backing Up."

7.12.3.9 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 to 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 "[D.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 restoration

On backup server Node3, create a volume for restoration (shadow volume).

17-1) Creating a shadow volume

```
# sdxshadowdisk -M -c Class2 -d sda=Disk1,sdb=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 restoration 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/st0.

See

For details on the restore method, see the manuals of file systems to be restored and used commands.

18a) When restoring data with the dd(1) command

```
# dd if=/dev/st0 of=/dev/sfdsk/Class2/dsk/Volume1 bs=32768
```

18b) When restoring the ext4 file system with the tar(1) command

18b-1) Create the ext4 file system to shadow volume Volume1.

```
# mkfs -t ext4 /dev/sfdsk/Class2/dsk/Volume1
```

18b-2) Mount the ext4 file system on shadow volume Volume1 on /mnt1, a temporary mount point.

```
# mkdir /mnt1
# mount -t ext4 /dev/sfdsk/Class2/dsk/Volume1 /mnt1
```

18b-3) Restore data held in the file system from tape.

```
# cd /mnt1
# tar xvf /dev/st0
```

18b-4) Unmount the file system mounted in step 18b-3).

```
# 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 "[7.12.3.8 Restoring from Proxy Volume](#)."

20-1) Activating the master volume

Activate master volume Volume1.

```
# sdxvolume -N -c Class1 -v Volume1
```

20-2) Resuming the services

When the file system on Volume1 was unmounted in step 13), mount it again.

Start the applications using Volume1.

7.12.4 Backing Up and Restoring by 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 a Dell EMC storage unit as a disk array unit and Dell EMC TimeFinder as a copy function.

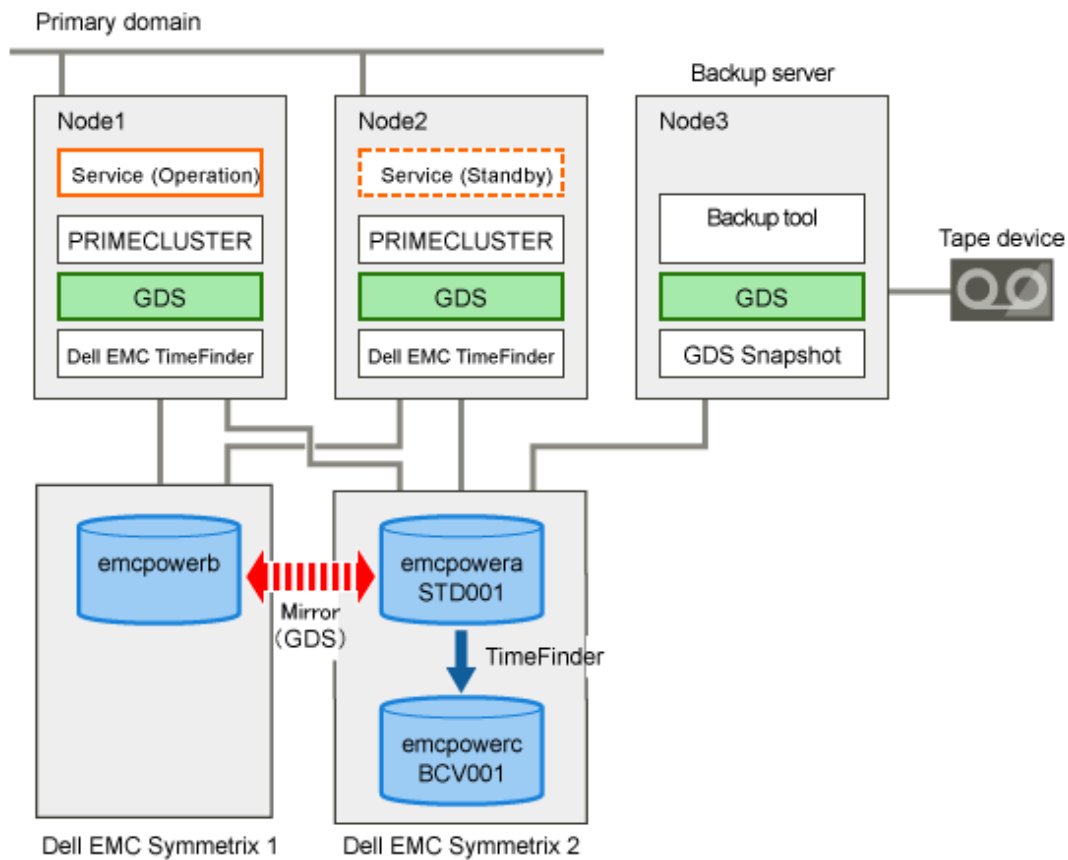
When restoration is performed with TimeFinder, configuration information within the private slice is also restored. For this reason, simultaneous backup and restore for all disks within the class is required. Additionally, the object configuration and status must match at backup and at restore, and so as in restore, it is necessary to disconnect the disks from mirroring before backup and reconnect after.

See

For using Dell EMC storage units, refer to notes described in "[3.3 Dell EMC Storage Units](#)" and "[6.1.3 To Use Dell EMC Storage Units](#)."

7.12.4.1 System Configuration

Figure 7.107 System Configuration



GDS: Global Disk Services

GDS Snapshot: Global Disk Services Snapshot

Information

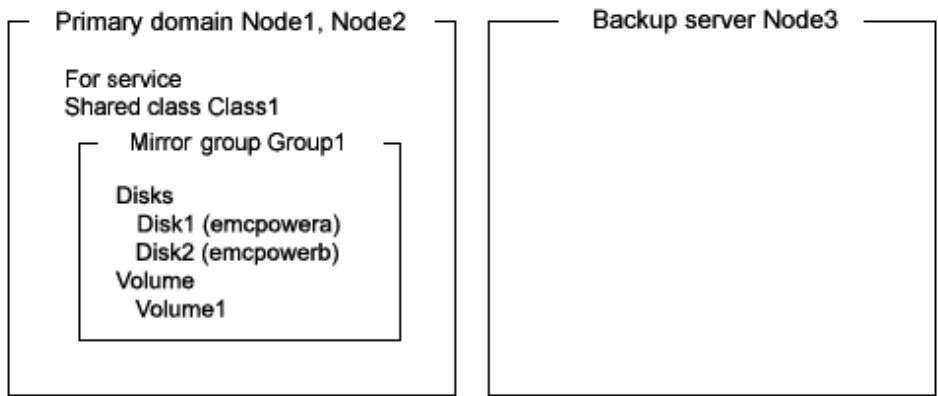
- A configuration that makes a node in the primary domain (e.g. Node2) work as a backup server is also available.

Note

Physical Device Name

Different physical device names (such as emcpowera) may be assigned to the identical physical disk in the primary domain and the backup server.

Figure 7.108 Object Configuration in Normal Operation



7.12.4.2 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 "3.17 Ensuring Consistency of Snapshot Data."

Figure 7.109 Backup

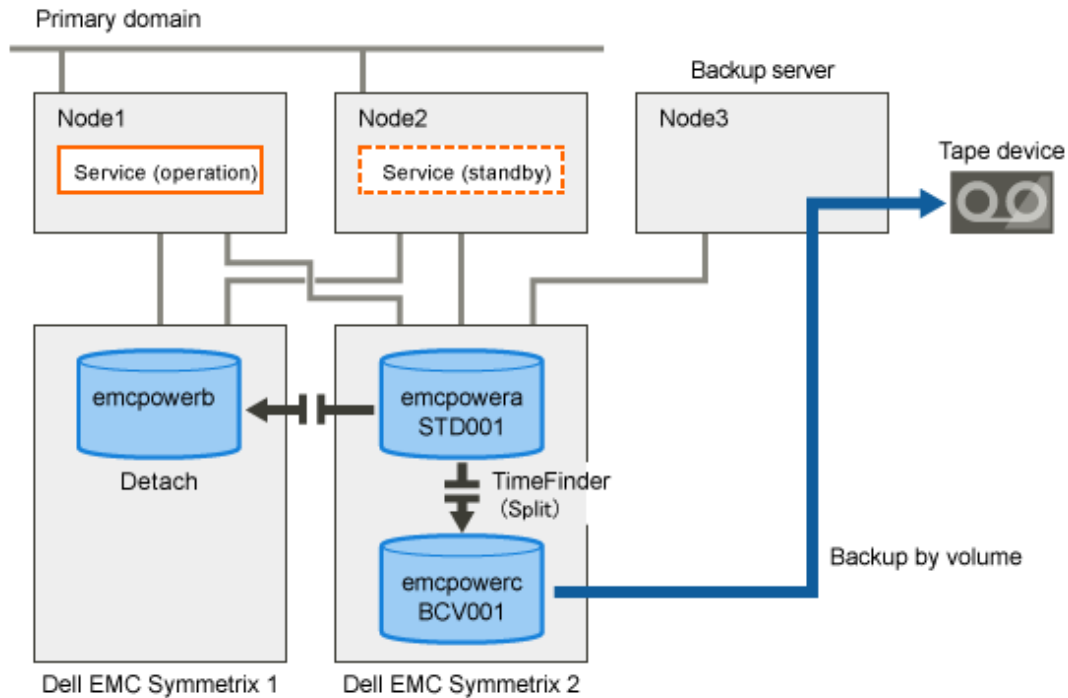


Figure 7.110 Object Configuration When Backing Up





Figure 7.111 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.

7.12.4.3 Summary of Restore from 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 to 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 7.112 Restore from BCV




Figure 7.113 Object Configuration When Backing Up from BCV


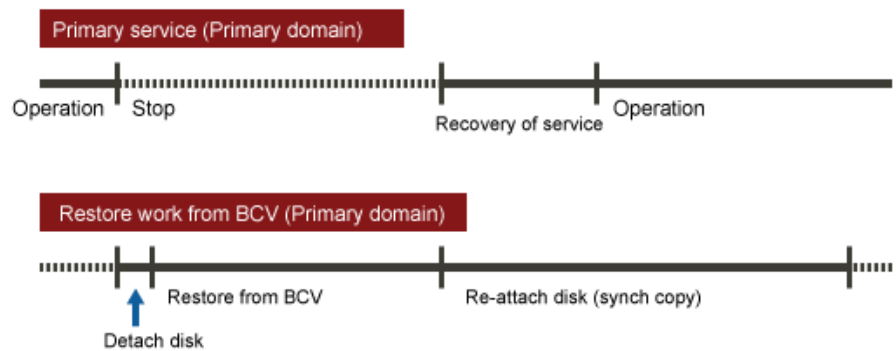


Figure 7.114 Schedule to Restore from BCV

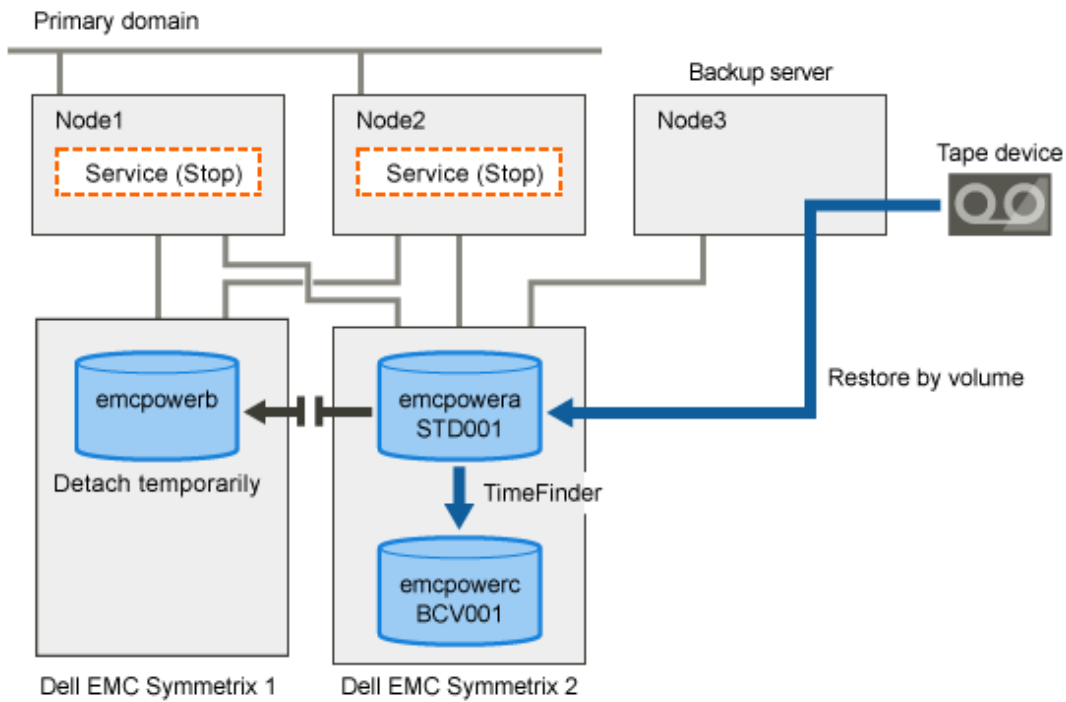


7.12.4.4 Summary of Restore from Tape

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 7.115 Restore from Tape



Information

In this configuration, access cannot be gained from backup server Node3 to disk emcpowerb. Therefore, after data held in emcpowera is restored from tape while emcpowerb is detached temporarily, synchronization copying from emcpowera to emcpowerb must be performed by reattaching emcpowerb. When access can be gained from Node3 to both emcpowera and emcpowerb, it is not required that emcpowerb be detached temporarily since data can be restored from tape back to both emcpowera and emcpowerb.

Figure 7.116 Object Configuration When Restoring from Tape


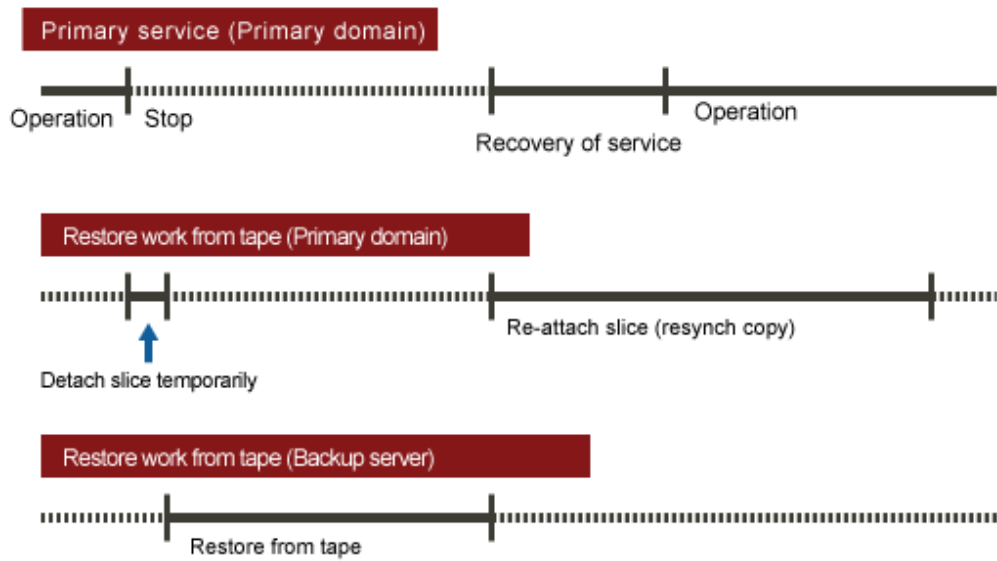


Figure 7.117 Schedule to Restore from Tape



7.12.4.5 Summary of Procedure

Figure 7.118 Outline of the Configuration Procedure



Figure 7.119 Outline of the Backup Procedure




Figure 7.120 Outline of the Procedure for Restoring from a BCV


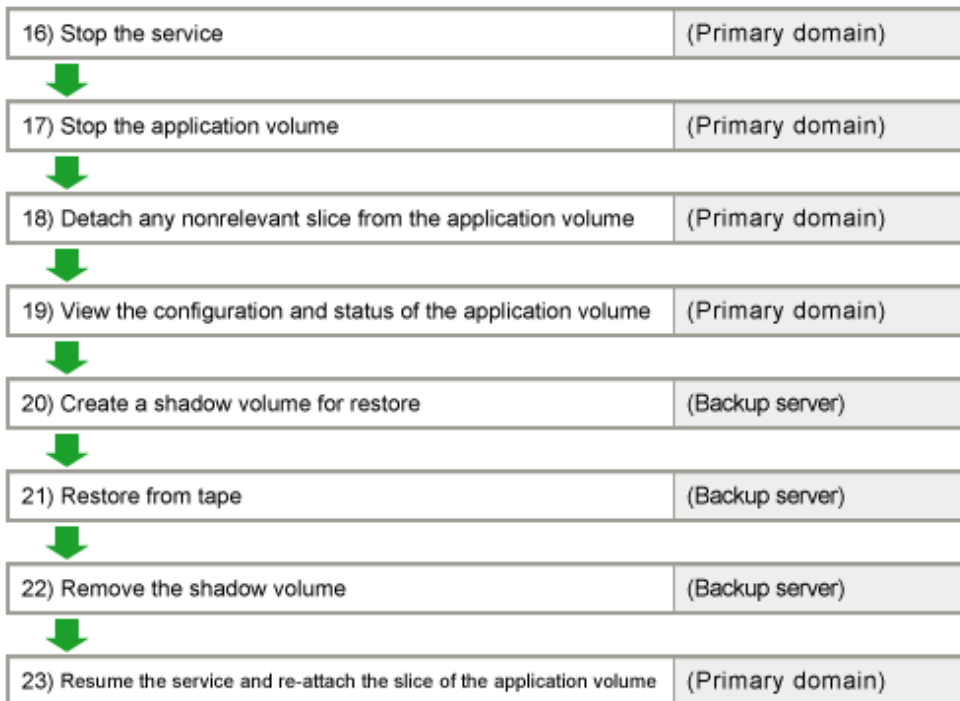


Figure 7.121 Outline of the Procedure for Restoring from Tape



7.12.4.6 Configuring Environment



Note

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. For details on the resource registration, see "6.3 Shared Disk Resource Registration."

1) Creating an application volume

Create a mirror volume used for the services operation on disks (standard devices) emcpowera and emcpowerb. The following settings are necessary on Node1 or Node2 in the primary domain.

1-1) Registering disks

Register disks (standard devices) emcpowera and emcpowerb 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  
emcpowera=Disk1,emcpowerb=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 emcpowera to BCV device emcpowerc 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 emcpowera with device group DevGroup, and name it STD001 as a logical device.

```
# symld -g DevGroup add pd /dev/emcpowera STD001
```

2-3) Relating a BCV device

Relate BCV device emcpowerc to device group DevGroup, and name it BCV001 as a logical device.

```
# symbcv -g DevGroup associate pd /dev/emcpowerc 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
```

7.12.4.7 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 emcpowerc from standard device emcpowera. 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 "3.17 Ensuring Consistency of Snapshot Data."

4-1) Viewing the status of the application volume

Check the slice on standard device emcpowera (Disk1) that is the copy source of BCV device emcpowerc 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 "D.1.1 Slice Status Abnormality."

4-2) Viewing the condition of the BCV pair

Confirm that standard device STD001 (emcpowera) and BCV device BCV001 (emcpowerc) 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. Check the underlined parts.

```
# 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     emcpowera     8380800      Node1:Node2     ENABLE
disk     Disk2     mirror    Class1     Group1     emcpowerb     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 emcpowerc. 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) emcpowerc with shadow class Class2, and name it Disk1.


```
# sdxshadowdisk -M -c Class2 -d emcpowerc=Disk1
```

Point

- The disk name must correspond to the disk name assigned in step 1-1) to standard device emcpowerc that is the copy source of BCV device emcpowerc. 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.

Check the underlined parts.

```
# 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  emcpowerc  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/st0.



See

For details on the backup method, see the manuals of file systems to be backed up and used commands.

8a) When backing up data with the dd(1) command

```
# dd if=/dev/sfdsk/Class2/dsk/Volume1 of=/dev/st0 bs=32768
```

8b) When backing up the ext4 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 ext4 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 -t ext4 /dev/sfdsk/Class2/dsk/Volume1
```

8b-3) Mount the ext4 file system on shadow volume Volume1 on /mnt1, a temporary mount point, in the read only mode.

```
# mkdir /mnt1
# mount -t ext4 -o ro /dev/sfdsk/Class2/dsk/Volume1 /mnt1
```

8b-4) Back up data held in the file system to tape.

```
# cd /mnt1
# tar cvf /dev/st0 .
```

8b-5) Unmount the file system mounted in step 8b-3).

```
# cd /
# umount /mnt1
# rmdir /mnt1
```

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).

7.12.4.8 Restoring from 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 ld 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), 15-3), 15-4) (confirming that the synchronization copying is complete), and 15-2).

15-1) Activating the application volume

Activate application volume Volume1.

```
# sdxvolume -N -c Class1 -v Volume1
```

15-2) Resuming the services

When the file system on application volume Volume1 was unmounted in step 11), mount it again.

Start the applications using Volume1.

15-3) 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.

15-4) 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
```

7.12.4.9 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 "[D.1.3 Volume Status Abnormality](#)" or "[D.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  emcpowera  8380800  Node1:Node2  ENABLE
disk  Disk2  mirror Class1  Group1  emcpowerb  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 restoration

On backup server Node3, create a volume for restoration (shadow volume) in disk emcpowera. The following settings are necessary on backup server Node3.

20-1) Registering a shadow disk

Register disk emcpowera with shadow class Class3, and name it Disk1.

```
# sdxshadowdisk -M -c Class3 -d emcpowera=Disk1
```

 **Point**

- The disk name must correspond to the disk name assigned to emcpowera 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 (emcpowera and emcpowerb), you must register all of the disks constituting Volume1 (emcpowera and emcpowerb) 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 (emcpowera and emcpowerb), you must connect all of the disks constituting Volume1 (emcpowera and emcpowerb) 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  DEVNAM  DEVBLKS  DEVCONNECT  STATUS
-----
disk  Disk1  mirror Class3  Group1  emcpowera  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/st0



See

For details on the restore method, see the manuals of file systems to be restored and used commands.

21a) When restoring data with the dd(1) command

```
# dd if=/dev/st0 of=/dev/sfdsk/Class3/dsk/Volume1 bs=32768
```

21b) When restoring the ext4 file system with the tar(1) command

21b-1) Create the ext4 file system to shadow volume Volume1.

```
# mkfs -t ext4 /dev/sfdsk/Class3/dsk/Volume1
```

21b-2) Mount the ext4 file system on shadow volume Volume1 on /mnt2, a temporary mount point.

```
# mkdir /mnt2
# mount -t ext4 /dev/sfdsk/Class3/dsk/Volume1 /mnt2
```

21b-3) Restore data held in the file system from tape.

```
# cd /mnt2
# tar xvf /dev/st0
```

21b-4) Unmount the file system mounted in step 21b-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 (emcpowera and emcpowerb), you must remove all of the disks registered with shadow class Class3 in step 20) (emcpowera and emcpowerb).

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), 23-3), 21-4) (confirming that the synchronization copying is complete), and 23-2).

23-1) Activating the application volume

Activate application volume Volume1.

```
# sdxvolume -N -c Class1 -v Volume1
```

23-2) Resuming the services

When the file system on application volume Volume1 was unmounted in step 16), mount it again.

Start the applications using Volume1.

23-3) Reattaching the slice of 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.

23-4) 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
```

7.13 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.



See

For details on how to use `sdxconfig` command, see "[B.1.13 sdxconfig - Object Configuration Operations](#)" and "[A.5 Backing Up and Restoring Object Configuration \(sdxconfig\)](#)".

7.13.1 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.

7.13.2 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 "7.13.1 Backing Up" in the event of loss of the Class1 object configuration caused by a problem of some kind.

1) Stopping and restriction of the automatic startup of RMS

When restoring the object configuration of the shared class registered in the cluster application, stop RMS and then restrict the automatic startup of RMS.

For details about stopping and restriction of the automatic startup of RMS, see "PRIMECLUSTER Reliant Monitor Services (RMS) with Wizard Tools Configuration and Administration Guide."

2) 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 "7.13.1 Backing Up."

3) 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 "7.13.1 Backing Up" In this example, the path to a destination file is "/var/tmp/Class1.conf".

4) Restoring the object configuration of the class

Execute the following command on the node where the configuration file was placed in step 3) 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 -r now
```



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 sda to sdb

```
# sdxconfig Convert -e replace -c Class1 -p sda=sdb -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 sdb.

```
# sdxconfig Convert -e replace -c Class1 -d Disk1=sdb -i /var/tmp/Class1.conf -o \  
/var/tmp/Class1.conf -e update
```

5) 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.

5-1) Stop the volume in the class.

```
# sdxvolume -F -c Class1
```

5-2) Change the class type and expand the class scope.

```
# sdxattr -C -c Class1 -a type=shared,scope=node1:node2
```

6) Settings for using shared classes in a cluster application

When the object configuration of the shared class Class1 registered as a Gds resource in the cluster application is restored, execute the following command on any one node within the class scope of Class1.

```
# /opt/SMAW/SMAWRrms/bin/hvgdsetup -a Class1  
...  
Do you want to continue with these processes ? [yes/no] y
```

7) Startup and canceling the restriction of the automatic startup of RMS

If the automatic startup setting of RMS is changed in step 1), restore the settings.

Also, if RMS is stopped in step 1), start RMS.

For the settings on how to startup and auto startup of RMS, see "PRIMECLUSTER Reliant Monitor Services (RMS) with Wizard Tools Configuration and Administration Guide."

7.14 Volume Expansion

7.14.1 Volume Expansion of Local or Shared Class

This section explains how to expand volumes of the shared class or the local class by using the sdxvolume -S command.

When the class name is Class1, the volume name is Volume1, and the expanded volume sizes are 204800 blocks, execute the following command on any one node within the class scope.

```
# sdxvolume -S -c Class1 -v Volume1 -s 204800
```



See

For details on the sdxvolum command, see "[B.1.4 sdxvolume - Volume Operations.](#)"

7.14.2 Volume Expansion Using Commands [EFI]

This section explains the procedural flow for expansion of /, /usr, and /var file systems with the snapshot functions of GDS Snapshot that can be performed while operations are running.

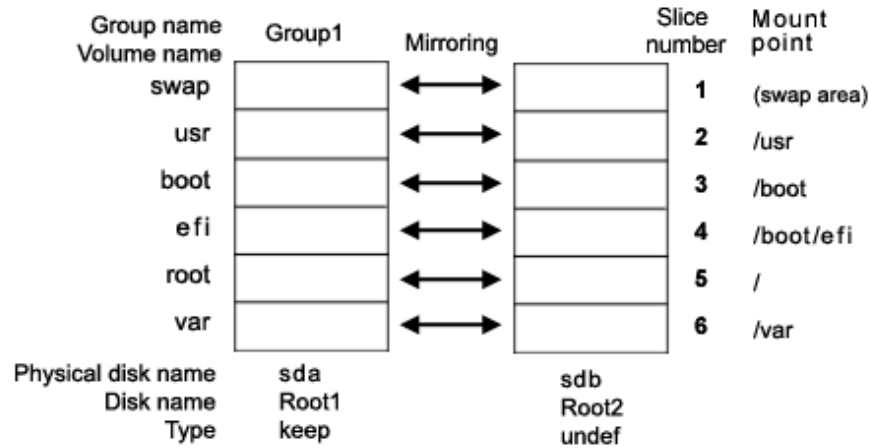
This procedure can be executed when the system volume is a file system whose capacity can be expanded. For details on the availability of file system expansion, see the manual of the file system.

[Procedure]



Assuming the configuration as below, the procedures for expanding /var file system space are explained.

As shown in the following figure, the / file system must be transferred to another volume no matter whether or not it is expanded.

before expansion



After expansion



 **Note**

The cylinder sizes of keep disks Root1 and Root3 must match.

Information

System disk mirroring is not a requirement. The configuration without `sdb` and `sdd` 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 the system disk

1-1) Stop the running application programs.

In order to ensure mirroring definition, all the running application programs must be stopped. When higher safety is required, create system disk backups.

1-2) Register the system disk to the root class.

```
# sdxdisk -M -c System -a type=root -d sda=Root1:keep,sdb=Root2:undef
```

1-3) Connect the system disk to a group.

```
# sdxdisk -C -c System -g Group1 -d Root1,Root2 -v 1=swap:on,2=usr:on,3=boot:on,4=efi:on,5=root:on,6=var:on
```

1-4) Check that mirroring definition is completed.

```
# sdxroot -M -c System -d Root1
```

1-5) Reboot the system.

```
# shutdown -r now
```

1-6) Check that mirroring is in effect.

Using the `mount(8)` command and the `sdxinfo` command, make sure that the system disk has been mirrored properly.

2) Creating proxy volumes

Create proxy volumes for the `/` file system and the expanded `/var` file system. At this point, one keep disk only should be connected to the group that will include the proxy volumes. This example shows the procedure for creating proxy volumes in the following configuration.

Group name	Group2	Slice	Mount
Volume name		number	point
root2		1	(N/A)
var2		2	(N/A)
(Free space)		(N/A)	(N/A)

Physical disk name	sdc
Disk name	Root3
Type	keep

Point

The disk must have sufficient free disk space following the last block of the proxy volume for the `/var` file system to be expanded.

2-1) Check 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  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 boot      System  Group1  off   on     2965680  3096899   131220  ACTIVE
volume efi       System  Group1  off   on     3096900  3228119   131220  ACTIVE
volume root      System  Group1  off   on     3228120  3403079   174960  ACTIVE
volume var       System  Group1  off   on     3403080  3638519   235440  ACTIVE
```

For the volume sizes, check the BLOCKS field displayed with the sdxinfo -V command. In this example, the root size is 174960 blocks and the var size is 235440 blocks.

2-2) Register the disk with the root class.

```
# sdxdisk -M -c System -d sdc=Root3:keep,sdd=Root4:undef
```



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 with a group.

```
# sdxdisk -C -c System -g Group2 -d Root3
```

2-4) Create volumes.

The volume sizes should be 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) Check the created volume sizes.

Make sure that the sizes of the volumes created in step 2-4) match the sizes 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  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 boot      System  Group1  off   on     2965680  3096899   131220  ACTIVE
volume efi       System  Group1  off   on     3096900  3228119   131220  ACTIVE
volume root      System  Group1  off   on     3228120  3403079   174960  ACTIVE
volume var       System  Group1  off   on     3403080  3638519   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

Join the created proxy volumes with the volumes for / and /var (master volumes) to copy data in / and /var to the proxy volumes. This example shows the procedure for joining the proxy volumes in the following configuration.



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) Check that synchronization copying is completed.

```
# 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  boot    ACTIVE
slice  System  Group1  Root2  boot    ACTIVE
slice  System  Group1  Root1  efi     ACTIVE
slice  System  Group1  Root2  efi     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
```

If the synchronization copying is in progress, "COPY" will be displayed in the STATUS field for the proxy volume slice. If "STOP" is displayed, synchronization copying is completed.

4) Parting the proxy volumes

With successful synchronization copying, the master volumes and the proxy volumes become equivalent. By parting those master volumes and proxy volumes, snapshots of the master volumes can be created on the proxy volumes.

4-1) Secure consistency of the file systems.

In order to secure consistency of snapshot file systems, it is necessary to prevent file system update. However, some file systems such as /, /usr, and /var are required for system operation and cannot be cancel the mounting statuses during system operation. There, use the following methods to minimize writing to the system disk and writing not updated on the system disk.

- a. Boot the system in single user mode (optional)
- b. Stop application programs having write access to the system disks (optional)
- c. Execute the sync(1) command to write file system data updated on the memory but not written yet.

Even if a., b., and c. are all put in action, file system update will not be completely prevented. Therefore, snapshot file systems may have inconsistency as in the case with after system panic occurrence.

When a., b., c. are all put in action, the snapshot file systems will be like those after system panic occurrence in single user mode.

When only c. is put in action skipping a. and b., the snapshot file systems will be like those after system panic occurrence during system operation.

In both situations, inconsistency may occur in the file systems, and it is necessary to check and repair consistency as instructed in step 5-1).

4-2) Part the proxy volumes.

```
# sdxproxy Part -c System -p root2,var2
```

4-3) When the system was rebooted in single user mode as instructed in step 4-1) a., reboot it in multi-user mode.

4-4) When the application programs were stopped as instructed in step 4-1) b., start them.

5) Configuring the alternative boot environment

Configure the environment to boot from the proxy volumes.

5-1) Check and repair the file systems on the proxy volumes.

The file systems on the proxy volumes may have inconsistency, and it is necessary to check and repair them with the fsck(8) command.

```
# fsck /dev/sfdsk/System/dsk/root2
# fsck /dev/sfdsk/System/dsk/var2
```

5-2) Configure the alternative boot environment.

```
# sdxproxy Root -c System -p root2,var2
```

With successful alternative boot environment configuration, the following message will be output.

```
SDX:sdxproxy: INFO: completed definitions of alternative boot environment:
current-boot-device=Root1 Root2
alternative-boot-device=Root3
```

Be sure to take a note of the output boot device names for the current boot environment (current-boot-device values) and those for the alternative boot environment (alternative-boot-device values).

6) Expanding snapshots

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.

This example shows the procedure for expanding the size of the snapshot volume for /var to 706320 blocks.

```
# sdxvolume -S -c System -v var2 -s 706320
```

6-3) Expand the snapshot file system size.



For details on how to expand the size of file system, see the manual of the file system.

7) Mirroring snapshots

Add a disk to the group that includes the snapshot volumes to mirror the snapshot volumes.

```
# sdxdisk -C -c System -g Group2 -d Root4
```

8) Switching to the alternative boot environment

By switching the environment to the alternative boot environment, replace the /var file system with the expanded volume.

8-1) Boot the system from the alternative boot environment.

Among boot devices shown in the EFI boot manager's boot option selection screen, select one from the boot devices for the alternative boot environment indicated by the message in step 5-2).

```
Root1  
Root2  
Root3 <- Select
```

8-2) Check that the boot was successful.

Using the mount(8) command the sdxinfo command, check that the system was booted normally in the alternative boot environment and that no errors exist in the GDS object statuses. As needed, check also that the file systems for the alternative boot environment are correct and that applications can run properly.



If the system was not booted normally, restore the original boot environment. To restore the original boot environment, among boot devices shown in the EFI boot manager's option selection screen, select one from the boot devices for the current boot environment indicated by the message in step 5-2).

```
Root1 <- Select  
Root2  
Root3
```

9) Deleting unnecessary volumes

After checking that the system was booted normally in the alternative boot environment, delete the volumes for the / file system and the original /var file system that build the previous boot environment.

```
# sdxvolume -F -c System -v root,var  
# sdxvolume -R -c System -v root  
# sdxvolume -R -c System -v var
```

7.15 I/O Response Time Assurance [IOmonitor]

Here the important points about operation and maintenance of the I/O Response Time Assurance Function provided as an option product "GDS I/O Monitor Option" for GDS will be explained.

- For the local class or shared class in a cluster environment, when the network load of cluster interconnect is under a heavy load (1 sec or more during transmission) or when CF (Cluster Foundation) has detected a disconnected heart beat, the I/O may return after exceeding the value of the I/O response time for 100 seconds or less.
- The setting of I/O Response Time Assurance cannot be backed up or restored by the backup (sdxcfg Backup) and restore (sdxcfg Restore) of the object configuration.

When backing up the object configuration of the class that the I/O Response Time Assurance configured, back up also the setting information (execution result of sdxcfg -I command) of I/O Response Time Assurance.

After restoring object configuration, check the execution result of sdxcfg -I command that was backed up and set the I/O Response Time Assurance by using sdxcfg command.

- To perform the following operations on a class that has system panic mode enabled, disable the system panic mode and then perform the following operations:
 - Canceling system disk mirroring (when the system panic mode is enabled for root class)
 - Deleting class (sdxcfg -R command or sdxcfg Remove command)

For procedure on how to change the system panic mode, see "[8.3.2 Changing the System Panic Mode.](#)"

7.16 Mirroring Among Servers

This chapter explains operation and maintenance when using the mirroring among servers.

7.16.1 I/O Error Due to Node Stop or Network Error

If the node stop or the network error occurs, the disk on another node cannot be accessed. In this case, the I/O error occurs on the disk of another node, and the slice where the I/O error occurs becomes INVALID.

After the node is restarted or the network is restored, the resynchronization copying is executed automatically. The slice of the operating node becomes the copy source and this resynchronization copying copies only the part where the data difference exists between slices.

However, the resynchronization copying which is executed after the following events copies the entire slice.

- When an I/O error occurs in the memory area of difference information inside the private slice
- When the resynchronization copying in process is canceled due to shutting the node, switching the cluster application, the sdxcfg command, etc.



When an error and restoration occurs repeatedly in the network or in the node, or when I/O error and network recovery occurs consecutively due to off and on failures, the synchronization or resynchronization copy may not be executed automatically.



I/O error on the disk of another node

When the I/O error occurs on the disk of another node that is connected to the netmirror group where the netmirror volume exists, the driver warning messages 22000 to 22010 are not output.

7.16.2 Stopping or Restarting the Node

- When one of the slices that configure the netmirror volume is in ACTIVE state and the other slice is in other than ACTIVE state, do not stop or restart the node that is connected to the disk where the ACTIVE slice exists. If the node is stopped or restarted, the ACTIVE slice does not exist in the netmirror volume. In this case, the service will stop.

To stop or restart the node, restore the status of the slice on the other node.

See below for how to restore the slice status.

- If the slice of the other node is in NOUSE state, execute the `sdxswap -I` command or use [Operation]:[Restore Physical Disk] of GDS Management View to restore the disk of the other node.
- If the slice of the other node is in INVALID state, refer "(1) The netmirror slice that configures the netmirror volume is in INVALID state." in "D.1.6 Errors in the Mirroring Among Servers" to restore the slice status.
- Even if all the slices that configure the netmirror volume are in ACTIVE or STOP state, the other slice becomes INVALID state when both nodes are restarted. This state is restored when the cluster application is started.

7.16.3 Operation when the LEFTCLUSTER Node Exists

If the cluster interconnect fails or the node panics, following events occur.

- One of the nodes becomes LEFTCLUSTER state.
- The I/O error occurs in the slices that configure the netmirror volume.

During this time, the following will be the state until the LEFTCLUSTER state is canceled.

- I/O from the application to the netmirror volume can be continued.

However, during synchronization or resynchronization copying, I/O to the area where copying process is running is suspended and the copying process does not progress.

- The slice where the I/O error has occurred will be in the following state:
 - It does not change to INVALID state but remains in ACTIVE, STOP or COPY state.
 - It becomes pre-detached state.

Pre-detached state is the state in which the message 44005 (WARNING: *object.volume*: detached *status* slice by an I/O error, class=*class*) is not output on any node after the warning message 22023 (WARNING: *sfdsk*: slice is pre-detached by an I/O error: ...) is output.

The pre-detached state cannot be checked in the `sdxinfo` command or the GDS Management View.

- I/O from the application to the net mirror volume is not issued to the slice where the I/O error occurred, but is issued only to another slice.

Therefore, the latest data is not stored in the slice where the I/O error occurred.

When the LEFTCLUSTER state is canceled, the slice in which the I/O error occurred is detached and becomes the INVALID state. Also, during synchronization or resynchronization copying, the suspended I/O returns after the copy destination slice is detached.

The LEFTCLUSTER state may be automatically canceled or not.

- When the LEFTCLUSTER state is automatically canceled

When the LEFTCLUSTER state is automatically canceled by the shutdown facility (SF) of PRIMECLUSTER, the maximum value of the time from the occurrence of an abnormality until the slice is detached is as follows.

[Maximum value of the time until the slice is detached (default)]

- PRIMERGY: 120 seconds
- PRIMEQUEST 3000 Series Business Model: 100 seconds
- PRIMEQUEST 3000 Series Enterprise Model: 140 seconds
- PRIMEQUEST 2000 Series: 100 seconds

- KVM guest: 160 seconds
- VMware guest: 110 seconds
- FJcloud-O environment: 270 seconds
- AWS environment: 50 seconds
- Azure environment: 110 seconds
- NIFCLOUD environment: 70 seconds
- RHOSP environment: 270 seconds

If [Timeout value of PRIMECLUSTER] below is changed from the default value, the maximum value of the time from the occurrence of an abnormality until the slice is detached can be calculated according to [Maximum value of the time until the slice is detached].

[Timeout value of PRIMECLUSTER]

- Timeout detection time of heartbeat of CF (CLUSTER_TIMEOUT)
- Timeout value of shutdown facility

[Maximum value of the time until the slice is detached]

<Maximum value of the time until the slice is detached (default) >
 + <The value of CLUSTER_TIMEOUT after removing the default value>
 + <Number of nodes> x <Timeout value of shutdown facility after removing the default value>

"Timeout value of shutdown facility" is the total value of timeout values of all the shutdown agents.

- When the LEFTCLUSTER state is not automatically canceled

The PRIMECLUSTER Shutdown Facility (SF) fails to shut down the node and the slice where the I/O error occurred is not detached without automatically canceling the LEFTCLUSTER state.

In this case, the LEFTCLUSTER state must be cleared manually.

For how to clear the LEFTCLUSTER state, see "[7.16.4 Restoration from the LEFTCLUSTER State](#)."

7.16.4 Restoration from the LEFTCLUSTER State

For procedures when restoring from the LEFTCLUSTER state, see "PRIMECLUSTER Installation and Administration Guide." During this time, do not restart the operating node.

When unavoidably restarting the operating node, such as when the operation state of the cluster system is on mutual standby, the synchronization copying and volume startup after restarting the node may be suppressed due to data protection. For restoration procedures, see "c." in "[7.16.8 Restoration when the Latest Disk Cannot Be Selected Automatically](#)."

7.16.5 Lock Volume of the Netmirror Volume

When the disk on one node cannot be accessed due to the node stop or the network error, data difference occurs between the disks in the netmirror group. If both nodes are started in this state, the disk containing the latest data is automatically selected when the cluster application is started, and this disk can be accessed.

After both nodes are started, the following will be the state until the disk containing the latest data is automatically selected when the cluster application is started.

- To protect the data, the start lock is set in the netmirror volume.
- Start of the netmirror volume or the synchronization copying is inhibited. (To start the netmirror volume, you need to start the cluster application.)
- The slice included in the disk containing the latest data may be in an INVALID state. (This state is only temporary.)



Note

Lock volume of the netmirror volume

Perform the lock volume operation only as the restoration procedure below. If a user performs the lock volume operation not as the restoration procedure, GDS may not work correctly.

7.16.6 Operation on Only One Node

To start the cluster application on only one node when another node cannot be started due to some error, for example, take the following steps to start the cluster application forcibly.

[Procedure]

1. If the netmirror slice is in INVALID state on the disk of the node on which the cluster application is forcibly started, restore the slice state.
Execute the following command on the node. For -d option, specify the disk name of the disk that is connected with the node among all the disks belonging to the netmirror group.

```
# sdxfix -V -c class_name -v netmirror_volume_name -d disk_name -e force -x NoRdchk
```

2. Start the cluster application forcibly.
For how to forcibly start the cluster application, see "PRIMECLUSTER Installation and Administration Guide."

In this case, the slice on the forcibly started node becomes ACTIVE while the slice of another node becomes INVALID. After that, the synchronization copying of the netmirror volume is inhibited.

7.16.7 Restoration to 2-node Operation after Operation on Only One Node

Take the following procedure to restore 2- node operation after forcibly starting the cluster application on only one node based on "7.16.6 Operation on Only One Node."



Information

2- node operation can be restored also by restarting both nodes. However, if GDS cannot automatically select the latest disk, the startup of cluster application after restarting will fail. In this case, follow a. in "7.16.8 Restoration when the Latest Disk Cannot Be Selected Automatically" and restore the disk.

Restoration Procedure

1. Start the stopped node.
2. Cancel the lock volume of the netmirror volume.

Take this procedure on both nodes.

```
# sdxattr -V -c class_name -v volume_name -a lock=off
```

3. Restore the iSCSI device and then perform the synchronization copying.

Take this procedure on the node where the cluster application is forcibly started.

For -d option of the sdxswap command, within the disks in the netmirror group, specify the disk of the node where the cluster application is not forcibly started.

```
# /etc/opt/FJSVsdx/bin/sdxiscsi_ctl -F -e init
# sdxswap -O -c class_name -d disk_name
# sdxswap -I -c class_name -d disk_name
```

7.16.8 Restoration when the Latest Disk Cannot Be Selected Automatically

For the following cases, GDS cannot determine the disk containing the latest data. In this case, a user must select the copy source disk for the synchronization copying to take the restoration procedure.

- a. Both nodes are stopped after the operation with an INVALID status slice, and the cluster application is forcibly started in the other node with the INVALID status slice.
- b. Both nodes are stopped after the standby node was stopped, and the cluster application is forcibly started with the node that was on standby.
- c. After the communication between the nodes is disabled due to the cluster interconnect failure or slowdown of the server, the operational node is stopped and switched.

In the above cases, starting the netmirror volume and the synchronization copying are inhibited.

The restoration procedure for the above a and b is as follows.

a. Both nodes are stopped after the operation with an INVALID status slice, and the cluster application is forcibly started in the other node with the INVALID status slice.

And b. both nodes are stopped after the standby node was stopped, and the cluster application is forcibly started with the node that was on standby.

Restoration Procedure

- 1. Start the stopped node.
- 2. Stop RMS on both nodes.
- 3. Restore the iSCSI device.

Take this procedure on the node where the cluster application is forcibly started last time.

```
# /etc/opt/FJSVsdx/bin/sdxiscsi_ctl -F -e init
```

- 4. The disk of the stopped node must be recognized on the operating node. For -d option of the sdxfix command, specify the disk of the stopped node.

Take this procedure on the operating node.

```
# sdxfix -D -c class_name -d disk_name -x Devlabel
```

-x Devlabel is the option to specify the restoration of the device information. Use this option only for this restoration procedure.

- 5. Decide the copy source slice for the synchronization copying.



- When selecting the slice of a node on which a cluster application has last started as a copy source for the synchronization copying, see "[Information: Verifying a node on which a cluster application has last started.](#)"
- When verifying the contents of each slice and then deciding the copy source slice for the synchronization copying, see "[Information: Verifying the contents of the slices that compose the netmirror volume.](#)"

- 6. If the copy source slice for the synchronization copying is in INVALID, restore the slice state. For -d option of the sdxfix command, specify the copy source disk for the synchronization copying.

Take this procedure on either of both nodes.

```
# sdxfix -V -c class_name -v volume_name -d disk_name -x NoRdchk
```

-x NoRdchk is the option to skip the read check of the slice. Use this option only for this restoration procedure.

- 7. Cancel the lock volume of the netmirror volume.

Take this procedure on both nodes.

```
# sdxattr -V -c class_name -v volume_name -a lock=off
```

8. Delete the iSCSI device information.

Take this procedure on both nodes.

```
# rm /var/opt/FJSVsdx/log/.sdxnetmirror_disable.db
```

9. Restore the status of the copy destination disk for the synchronization copying.

Take this procedure on either of both nodes.

For -d option of the sdxswap command, within the disks in the netmirror group, specify the copy destination disk for the synchronization copying.

```
# sdxswap -O -c class_name -d disk_name
# sdxswap -I -c class_name -d disk_name
```

10. Start RMS on both nodes.

In this case, the synchronization copying of the netmirror volume is performed.

c. After the communication between the nodes is disabled due to the cluster interconnect failure or slow down of the server, the operational node is stopped and switched.

Restoration Procedure

1. Decide the copy source slice for the synchronization copying.



.....
- When selecting the slice of a node on which a cluster application has last started as a copy source for the synchronization copying, see "[Information: Verifying a node on which a cluster application has last started.](#)"
- When verifying the contents of each slice and then deciding the copy source slice for the synchronization copying, see "[Information: Verifying the contents of the slices that compose the netmirror volume.](#)"
.....

2. If the copy source slice for the synchronization copying is not in INVALID, restore the iSCSI device information. For -d option of the sdxnetdisk command, specify the copy source disk for the synchronization copying.
Take this procedure on either of both nodes.

```
# /etc/opt/FJSVsdx/bin/sdxnetdisk -F -e devinfo -c class_name -v volume_name -d disk_name
```

3. If the copy source slice for the synchronization copying is in INVALID, take the following steps.

3-1) Stop RMS on both nodes.

3-2) Restore the slice state.

For -d option of the sdxfix command, specify the copy source disk for the synchronization copying.

Take this procedure on either of both nodes.

```
# sdxfix -V -c class_name -v volume_name -d disk_name -x NoRdchk
```

-x NoRdchk is the option to skip the read check of the slice. Use this option only for this restoration procedure.

4. Perform the synchronization copying.
Take this procedure on either of both nodes.

```
# sdxcopy -B -c class_name -v volume_name
```

5. Restart the service if necessary.

Information: Verifying a node on which a cluster application has last started

Execute the following command on both nodes, and compare the output results of both nodes.


```
# /etc/opt/FJSVsdx/bin/sdxiscsi_ctl -L -e time -c class_name
```

[Execution example]

Node 1

```
# /etc/opt/FJSVsdx/bin/sdxiscsi_ctl -L -e time -c class1
Wed Oct 26 19:13:16 2016
```

Node 2

```
# /etc/opt/FJSVsdx/bin/sdxiscsi_ctl -L -e time -c class1
Tue Oct 11 11:34:56 2016
```

The output displays the date and time the Gds resource of the class specified in the argument last started on the node where the command was executed.

In the above example, the last started cluster application was on Node 1.

Information: Verifying the contents of the slices that compose the netmirror volume

Take the following steps to change the state of the slice whose contents you want to verify to STOP, and to change the state of the other slice to INVALID. After that, verify the contents of the volume and you can verify the contents of the slice.

- 1) Stop RMS on both nodes.
- 2) If the slice whose contents you want to verify is in INVALID, change the slice status.

For -d option of the sdxnetdisk command, specify the disk of the slice in INVALID.

Take this procedure on either of both nodes.

```
# /etc/opt/FJSVsdx/bin/sdxnetdisk -F -e slice -c class_name -v volume_name -d disk_name
```

After executing this command, the slice whose contents you want to verify becomes STOP, and the other slice becomes INVALID.

- 3) Start the volume.

Take this procedure on either of both nodes.

```
# sdxvolume -N -c class_name -v volume_name -e unlock,mode=ro
```

- 4) Verify the contents of the volume.
- 5) Stop the volume.

Perform this step on the node on which 3) of this procedure has been performed.

```
# sdxvolume -F -c class_name -v volume_name
```

7.16.9 Stopping both Nodes

Before stopping both nodes, execute the hvshut -a command to stop RMS.

If both nodes are stopped without stopping RMS when the cluster application is activated, if the standby node is stopped first, the following error may occur when the next time the cluster application is started.

[Event]

When only the first stopped node is started, the cluster application may not be forcibly started.

The cluster application must be forcibly started to start the cluster application on only one node.

If the standby node is stopped first, writing to the disk on the node is interrupted. At this time, data or file system may be inconsistent. In this case, the cluster application may fail to be forcibly started.

[Restoration procedure]

Start the last stopped node and start RMS to restore the cluster application to be started.

If the last stopped node cannot be started, take the following steps on the started node.

1. Cancel the start lock of the netmirror volume.

```
# sdxattr -V -c class_name -v volume_name -a lock=off
```

2. Start the netmirror volume.

```
# sdxvolume -N -c class_name -v volume_name
```

3. Restore data or file system.
Restore the file system, re-create the file system, or restore the data.

4. Stop the netmirror volume.

```
# sdxvolume -F -c class_name -v volume_name
```

5. Start RMS.



Information

You can stop both nodes without stopping RMS when the cluster application is activated if the following operations are performed after stopping both nodes.

- Starting both nodes again
In this case, the cluster application can be started correctly because the correct data on the last stopped node can be used. At this time, the synchronization copying from the disk on the last stopped node to the disk on the first stopped node is automatically performed. In this case, inconsistency in the data or the file system on the first stopped node is automatically cleared.
- Starting only the last stopped node
The cluster application can be forcibly started.
For how to restore the 2-node operation, see "[7.16.7 Restoration to 2-node Operation after Operation on Only One Node.](#)"

7.16.10 Restarting the Node during Synchronization Copying

When the synchronization copying of the netmirror volume is performed, do not restart the node on which the copy source disk exists.

If the node on which the copy source disk exists is restarted, a slice that contains the appropriate data in the netmirror volume no longer exists. In this case, I/O to the netmirror volume cannot be performed, and operation will stop.



See

For details, see "(2) [The cluster application becomes Faulted or Inconsistent after the node is restarted.](#)" in "[D.1.6 Errors in the Mirroring Among Servers.](#)"

7.16.11 Switchover of the Cluster Application when the Network is Disconnected

To switch over the cluster application when the slice of the operating node is ACTIVE and the slice of the standby node is in any status except STOP (when the slice is detached) among the slices that configure the netmirror volume, check the status of the network that is used for the mirroring among servers.

If the cluster application is switched over when the slice of the standby node is detached as well as the network is disconnected, operation will stop. This is because on the new operating node, the slice on the disk of the own node is detached and cannot be used, and the disk on another node cannot be used due to the disconnected network. At the same time, offline processing of the cluster application may fail on the new operating node, and the node may panic.

7.16.12 Configuration Change when the Node is Stopped

Following operations can be performed even if the stopped node exists in the configuration other than the mirroring among servers. However, in the configuration of mirroring among servers, these operations cannot be performed when the stopped node exists.

In the configuration of mirroring among servers, perform these operations when both nodes are active.

- Connect the disk (sdxdisk -C)
- Disconnect the disk (sdxdisk -D)
- Remove the group (sdxgroup -R)
- Create the volume (sdxvolume -M)
- Remove the volume (sdxvolume -R)
- Change the class attribute (sdxattr -C, except scope extension)
- Change the disk name (sdxattr -D -a name)
- Change the group attribute (sdxattr -G)
- Change the volume attribute (sdxattr -V, except the change of lock mode)
- Swap the physical disk (sdxswap -I)
- Restore the physical disk (sdxswap -O)
- Restore the volume (sdxfix -V) (except when -e force option is specified)

7.17 Maintenance in Cloud Environment

When maintaining the system in the cloud environment, refer to the following manual.



See

.....
See "PRIMECLUSTER Installation and Administration Guide 4.6 Cloud Services."
.....

7.18 Removing Errata Kernel

If an errata kernel has been removed in an environment where the system disk is mirrored, to update the grub.cfg file, execute the following command before restarting the system.

```
# grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg
```

If the grub2-mkconfig command is not executed, the system may not be started. For the resolution of this error, see "(6) System cannot be booted. (Removing errata kernel)" in "D.1.5 System Disk Abnormality [EFI]."

 See

.....
For details on the grub2-mkconfig command, see the OS manual.
.....

 Note

.....
Do not remove the running errata kernel when the system disk mirroring was set. If the errata kernel is removed, no boot entry that can be started may exist on the boot option selection screen of the EFI boot manager.
.....

Chapter 8 Configuration Change

This section explains how to change the configuration of GDS, GDS Snapshot, and GDS I/O Monitor Option.

8.1 Changing Web-Based Admin View Settings

When changing one of the following settings after installing Web-Based Admin View, see "PRIMECLUSTER 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

8.2 Changing Object Configuration

This section explains how to change the object configuration of GDS classes, groups, and volumes.



Recreating shared class registered in a cluster application

When recreating a shared class registered as a Gds resource in a cluster application, execute the following command on any one node within the class scope.

```
# /opt/SMAW/SMAWRrms/bin/hvgdsetup -a class_name
...
Do you want to continue with these processes ? [yes/no] y
```

This command is set up in order to start up or stop the volume of shared class along with the change in the state of the cluster application.

8.2.1 Using GDS Management View

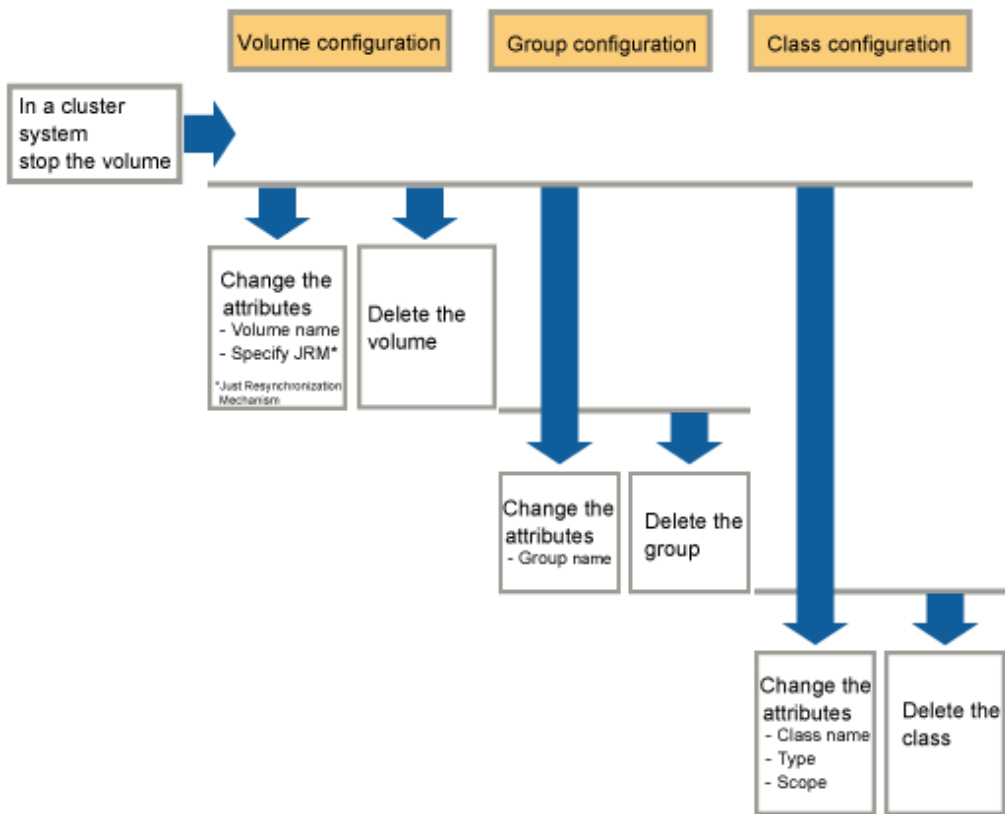
Class configurations, group configurations, and volume configurations can be changed through two types of operation: reconfiguration and attribute change.

This section explains the changing procedures for each configuration.

8.2.1.1 Operation Flow

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

Figure 8.1 Configuration Changing Operation



8.2.1.2 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 8.2 Class Configuration



2. Selecting class to change

Select the class you want to change from Class Configuration screen [Class Name].

3. Changing configuration

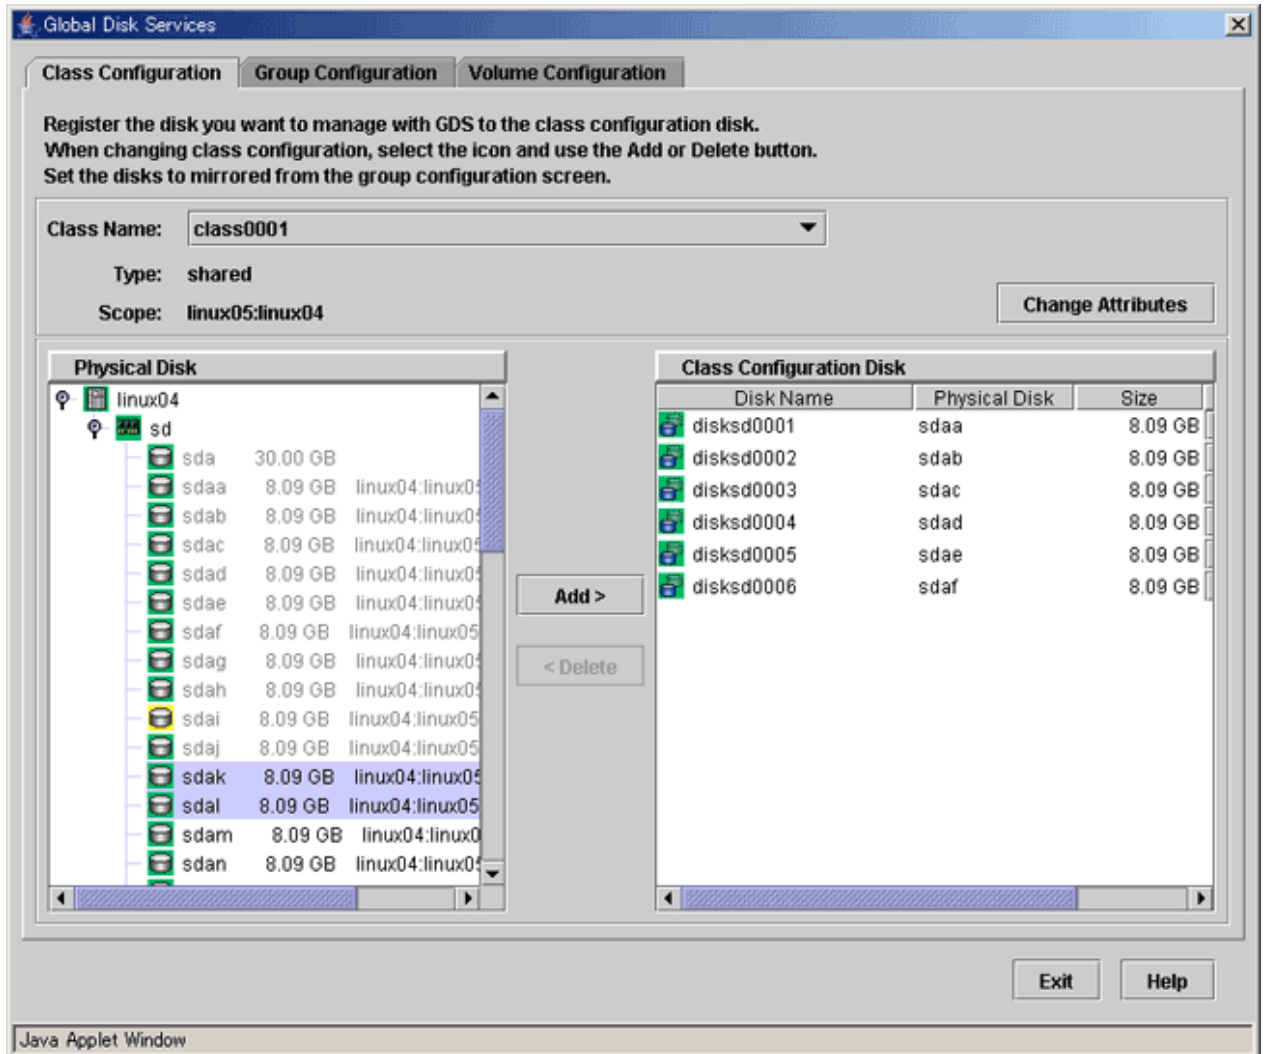
Follow the procedures below to register a physical disk (create a disk), or to remove a disk.

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 8.3 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.

- 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.
- Do not change the class name.

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 8.4 Class Attributes Definition Screen



2. Changing attributes

a. For a single node

You cannot change the class attributes.

b. For a cluster

You can change the following attributes:

- Type
- Scope

3. Implementing changes

Click <OK> to implement the attributes change, or <Cancel> to cancel.

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 8.5 Disk Attribute Definition Screen



3. Changing attributes

You can change the following attributes.

- Disk Name
- Disk Type



For details on assigning a disk name, see "[2.3.1 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 "[2.3.1 Object Name.](#)"
.....

8.2.1.3 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 8.6 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 8.7 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.

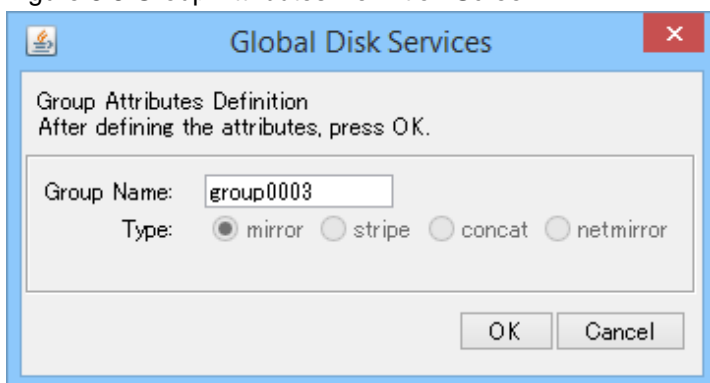
a. Operation menu in the Main screen

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 8.8 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 "2.3.1 Object Name."

8.2.1.4 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.

- If the volume name is described in the /etc/fstab file of the selected node in the GDS Configuration Tree Field, the volume name of that volume cannot be changed. When changing the volume name, change first the volume name described in the /etc/fstab file.


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 8.9 Volume Attributes Definition Screen



3. Changing attributes

You can change the following attributes:

- Volume Name
- JRM (on/off)
- Physical Slice (on/off)



For information on assigning a volume name, see "2.3.1 Object Name."

4. Applying changes

Click <OK> to apply changes in the attributes change, or <Cancel> to cancel.



Changes in Special File Path Name by Changing Volume Name

Changing the volume name will also change the special file path name used to access the volume, so you must update the settings of software (example: the /etc/fstab.pcl file of PRIMECLUSTER) which accesses to the special file of volume.

8.2.2 Using Commands

When changing the object configuration by using the commands, use the following commands.

- Changing class configuration
 - Registering a disk to a class: `sdxdisk -M`
 - Deleting a disk from a class: `sdxdisk -R`
 - Changing class attributes: `sdxattr -C`
 - Changing disk attributes: `sdxattr -D`
- Changing group configuration
 - Connecting disks to a group: `sdxdisk -C`
 - Disconnecting disks from a group: `sdxdisk -D`
 - Connecting a lower level group to a group: `sdxgroup -C`
 - Disconnecting a lower level group from a group: `sdxgroup -D`
 - Changing group attributes: `sdxattr -G`
- Changing volume configuration
 - Changing volume attributes: `sdxattr -V`



See

For details on each command, see "[Appendix B Command Reference](#)."

8.3 Changing the Settings of I/O Response Time Assurance [IOMonitor]

8.3.1 Changing the I/O response time

Here the method of changing the maximum time before the I/O responds will be explained.

1. Check the settings of the I/O response time assurance.

Example) If checking the I/O Response Time Assurance settings for the class Class.

```
# sdxinfo -I -c Class
CLASS  MODE  TIME  PANIC
-----
Class  on   77   on
```

2. Change the I/O response time.

Example) Change the I/O response time of the class Class to 90 seconds.

```
# sdxattr -C -c Class -a iotimeout=on:90
```

3. Check that the settings of the I/O Response Time Assurance Function has been changed.

Example) If checking the I/O Response Time Assurance settings for the class Class.

```
# sdxinfo -I -c Class
CLASS  MODE  TIME  PANIC
-----
Class  on   90   on
```

If the value set in procedure 2 is displayed in the MODE field, the TIME field and the PANIC field of the specified class, the change in the settings of the I/O Time Response Assurance is complete.

- Change the settings of the Hang-up Monitoring Function if a value other than the default value was set for the I/O response time. Set the conclusion time with the same value of I/O response time, and set the monitoring interval twice the value of conclusion time. For details, see "[6.12.2 Hang-up Monitoring Function.](#)"

8.3.2 Changing the System Panic Mode

Here the method of changing the settings for System Panic Mode will be explained.

- Check the settings of the I/O Response Time Assurance.

Example) If checking the I/O Response Time Assurance settings for the class Class.

```
# sdxinfo -I -c Class
CLASS  MODE  TIME  PANIC
-----
Class  on   77   on
```

- Change the settings for the System Panic Mode.

Example) Change the System Panic Mode to off for the class Class.

```
# sdxattr -C -c Class -a iotimeout=on:77:off
```

- Check that the settings of the I/O Response Time Assurance Function has been changed.

Example) If checking the I/O Response Time Assurance settings for the class Class.

```
# sdxinfo -I -c Class
CLASS  MODE  TIME  PANIC
-----
Class  on   77   off
```

If the value set in procedure 2 is displayed in the MODE field, the TIME field and the PANIC field of the specified class, the change in the settings of the I/O Time Response Assurance is complete.

8.4 Deleting the Settings of the I/O Response Time Assurance [IOmonitor]

Here the method of deleting the settings of the I/O Response Time Assurance will be explained.

- Delete the settings of the I/O Response Time Assurance.

Example) If deleting the I/O Response Time Assurance settings for the class Class.

```
# sdxattr -C -c Class -a iotimeout=off
```

- Check the settings of the I/O Response Time Assurance.

Example) If checking the I/O Response Time Assurance settings for the class Class.

```
# sdxinfo -I -c Class
CLASS  MODE  TIME  PANIC
-----
Class  off   *    *
```

If off is displayed in the MODE field, and an asterisk (*) in the TIME field and PANIC field of the specified class, the settings of the I/O Response Time Assurance has been deleted successfully.

8.5 Migration from Single Nodes to Cluster System

The following describes methods to change over to a cluster system from one or more single nodes where classes have already existed by installing the cluster control facility. The procedure varies depending on each class type.

- Root class

Unmirror the system disk, then install and set up the initial configuration of the cluster control facility. After that, set up the mirror of the system disk.

- Local class

Back up the volume data as necessary, and then delete the local class. Install and set up the initial configuration of the cluster control facility, and then create the classes and volumes again. Restore the backed up volume data as necessary.

Important Point 1

A root or local class created on a single node cannot be used directly in a cluster system. For local classes, 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 "[D.1.10 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.

8.6 Server Migration

This sub-section describes the method to migrate a server to a new server on the system where GDS is used.

Setting mistakes during migration can be prevented and servers can be migrated efficiently by backing up and restoring object configuration information by using the `sdxconfig` command.

Figure 8.10 Server migration



The migration procedure is as follows:

- 1) Back up and delete the GDS configuration on the migration source server.
 - 1-1) Back up object configuration information of local and shared classes. (sdxcfg Backup command)
 - 1-2) Stop the volumes of local and shared classes. (sdxvolume -F command)
 - 1-3) Convert shared classes to local classes. (sdxattr -C command)
 - 1-4) Remove object configuration information of local classes (including the classes converted in step 1-3)). (sdxcfg Remove command)
 - 1-5) When the system disk is being mirrored, cancel the system disk mirroring. (Unmirror System Disk of the Management View)
- 2) Remove the storage unit and the internal disks registered with local classes from the migration source server and connect them to the migration destination server.

Information

.....
If the internal disk cannot be connected to the migration destination server, back up the volume data of a local class before performing step 1) on the migration source server, and then restore it on the migration destination server after performing step 4).
.....

- 3) Install and set up the OS and software on the migration destination server.
- 4) Restore the GDS configuration on the migration destination server.
 - 4-1) When storage is a shared disk in the cluster system, register the disk with the PRIMECLUSTER resource database. (clautoconfig command)
 - 4-2) When mirroring the system disk, set the system disk mirroring. (System Disk Settings of the Management View)
 - 4-3) Convert the physical disk names described in the configuration files backed up in step 1) to physical disk names on the migration destination server. (sdxcfg Convert command)
 - 4-4) Restore object configuration information of local and shared classes. (sdxcfg Restore command)
 - 4-5) When storage is a shared disk in the cluster system, convert the restored class to a shared class. (sdxattr -C command)



- For the method of using the `sdxconfig` command, see "B.1.13 `sdxconfig` - Object Configuration Operations", "7.13 Backing Up and Restoring Object Configurations", and "A.5 Backing Up and Restoring Object Configuration (`sdxconfig`)."
- For the method of using the `clautoconfig` command, see "6.3 Shared Disk Resource Registration."

8.7 Storage Migration

This sub-section describes the method to migrate storage units in which volumes of local and shared classes are created on the system where GDS is used, to new storage units.


The following two methods are available for storage migration:

- Using the GDS Snapshot online disk migration.
- Using the copy function of storage units and the `sdxconfig` command.

8.7.1 Using GDS Snapshot Online Disk Migration

By using the online disk migration, storage units can be migrated to new units without stopping the services.

Figure 8.11 Online disk migration



Migrate volume data by following methods:

- Use REC for migration between ETERNUS Disk storage systems.
- Use Dell EMC SRDF for migration between Dell EMC Symmetrix.
- Use the soft copy function of GDS Snapshot for storage units other than above.

To migrate with this method, the following conditions must be met:

- GDS Snapshot has been installed.
- Storage of the migration source and migration destination can be connected simultaneously to a server.

- GDS volumes are single or mirror volumes.
- When using REC, GDS volumes are single volumes or volumes in a mirror group which is not hierarchized.
- When using Dell EMC SRDF, GDS volumes are mirror volumes in a mirror group which is not hierarchized.

 See

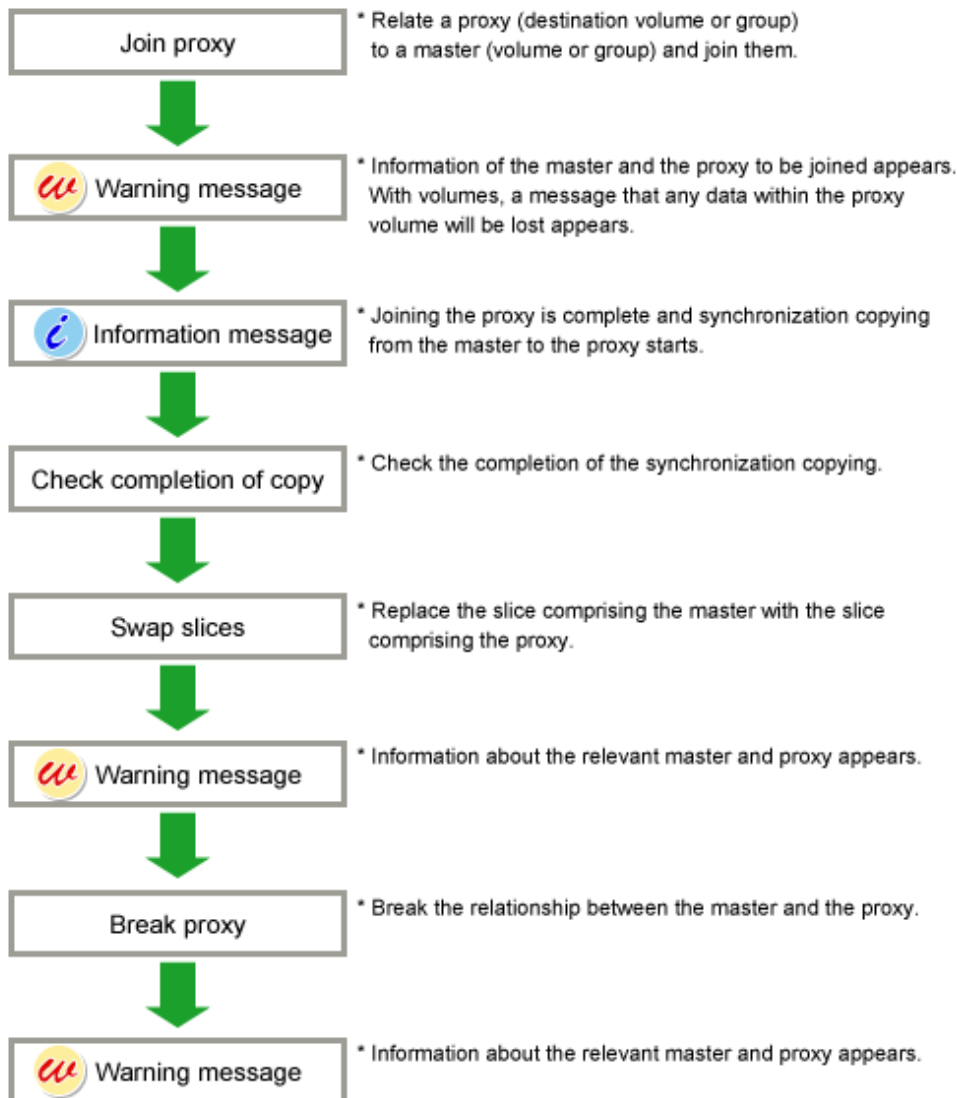
- For disk migration with proxy volumes, see "[1.5.5 Online Disk Migration](#)."

8.7.1.1 Using GDS Management View

8.7.1.1.1 Operation Flow

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

Figure 8.12 Disk Migration Operation



 Note

Disk Migration Precondition

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

8.7.1.1.2 Operation Procedure

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



Proxy Volume Use Conditions

See the following sections for points of concern.

- ["2.3.12 Proxy Configuration Preconditions"](#)
 - ["2.3.13 Number of Proxy Volumes"](#)
 - ["2.3.14 Proxy Volume Size"](#)
 - ["2.3.15 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 ["6.10.1.1 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.

Note

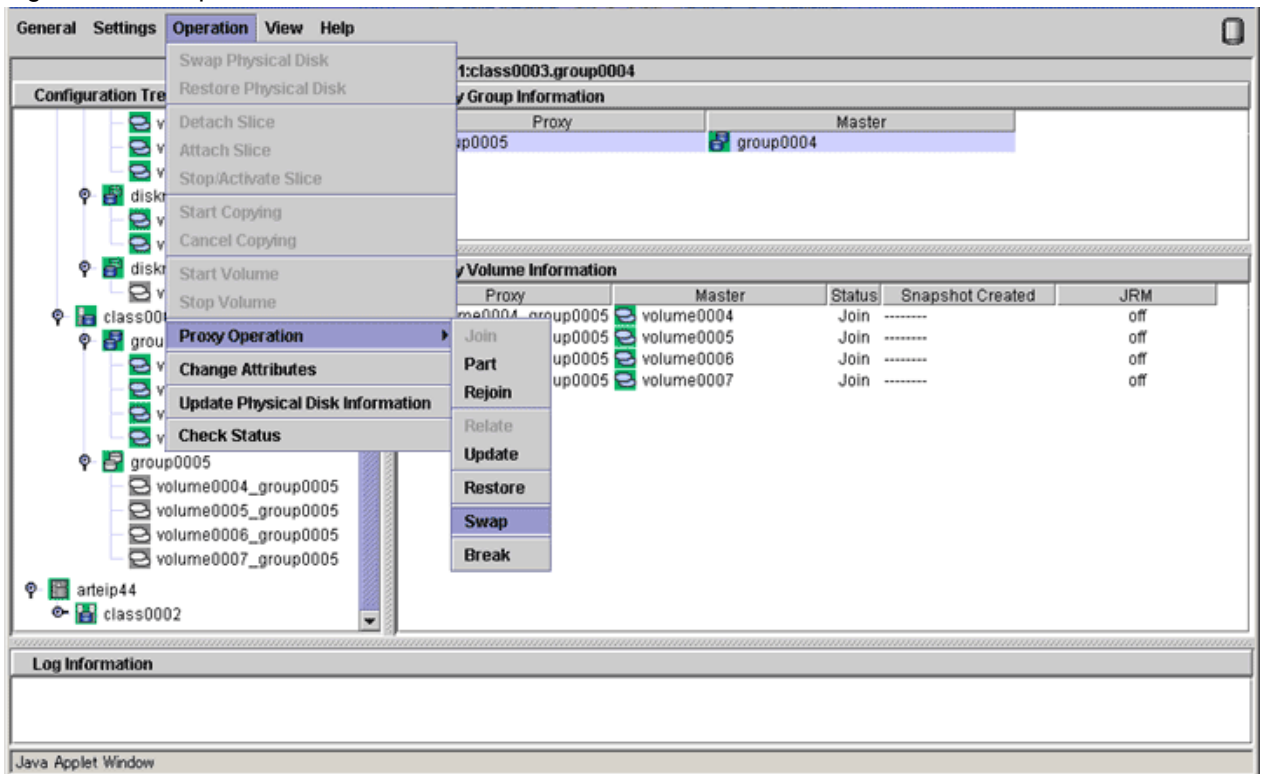
If There Is a Session by Disk Unit's Copy Function

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 "[B.1.6 sdxinfo - Display Object Configuration and Status Information](#)" and "[B.2.1 sdxproxy - Proxy Object Operations](#)."

2. Selecting the [Swap Slice] menu

Select [Operation]:[Proxy Operation]:[Swap Slice] in the Main screen.


Figure 8.13 Swap Slice



3. Confirming

A confirmation screen appears asking you whether to swap the slices.

Figure 8.14 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 8.15 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 "[9.1 Breaking Proxy](#)."

8.7.1.2 Using Commands

The migration procedure is as follows:

1) Create GDS objects on the migration destination storage.

1-1) Connect the migration destination storage to a server.

1-2) When the migration destination storage is a shared disk in the cluster system, register the migration destination disks with the PRIMECLUSTER resource database. (clautoconfig command)

1-3) Create volumes or groups on the migration destination storage.

- When the migration source are single volumes

Create single volumes of the same size as the migration source. (sdxdisk -M and sdxvolume -M commands)

- When the migration source are mirror volumes

Create mirror groups of the same size as the migration source. (sdxdisk -M and sdxdisk -C commands)

When using REC or Dell EMC SRDF, connect only one disk to each of the mirror groups.

2) Copy data to the migration destination storage from the migration source storage.

2-1) Join the master with the proxy as the migration source storage is master and the migration destination storage is proxy. (sdxproxy Join command)

- When the migration source is a single volume, join it per volume.

- When the migration source is a mirror volume, join it per group.

2-2) When copying by using REC or Dell EMC SRDF in step 2-1) and the master and the proxy become synchronized, cancel the session of REC or SRDF between the master and proxy. (sdxproxy Cancel command)

3) Swap storage of the master and proxy. (sdxproxy Swap command)

4) Break the proxy. (sdxproxy Break command)

5) When copying by using REC or Dell EMC SRDF in step 2-1) and mirroring the migration destination storage, add a disk to the mirror group of the migration destination. (sdxdisk -C command)

6) Delete GDS objects from the migration source storage.

6-1) Delete volumes, groups, and disks from the migration source storage. (sdxvolume -R, sdxgroup -R, and sdxdisk -R commands)

6-2) When the migration source storage is a shared disk in the cluster system, delete the migration source disk from the PRIMECLUSTER resource database. (cldehrsc command)

6-3) Remove the migration source storage from a server.

 See

- For the method of using the sdxproxy command, see "B.2.1 sdxproxy - Proxy Object Operations."
- For the method of using the clautoconfig and cldehrsc commands, see "6.3 Shared Disk Resource Registration."

8.7.2 Using Copy Function of Storage Units and sdxconfig Command

Setting mistakes can be prevented during storage migration and storage can be migrated efficiently by backing up and restoring object configuration information by using the sdxconfig command.

Migrate volume data by using the copy function per LUN (such as ACM CCM) of storage units in one of the following methods:

- Copy to one migration destination case from one migration source case.
- Copy to each migration destination case from each migration source case.

Figure 8.16 Copying to one migration destination case from one migration source case


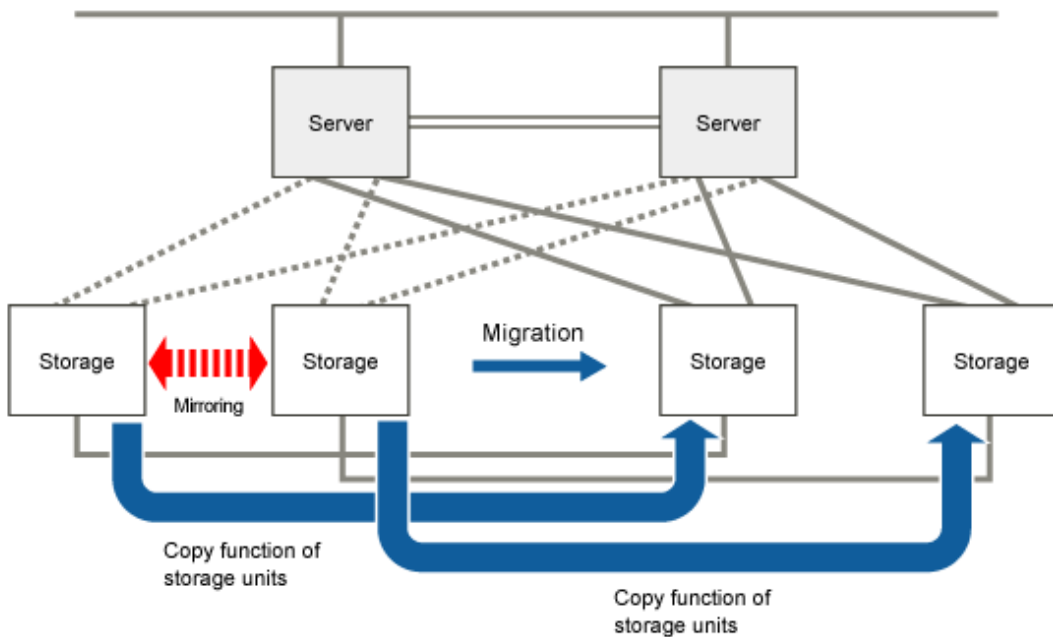


Figure 8.17 Copying to each migration destination case from each migration source case



Information

For a single volume or a mirror volume in a mirror group which is not hierarchized, it can be migrated by using the copy function per slice (such as REC) of storage units. In this case, create a volume of the same size as the migration source on the migration destination storage, and then copy data between slices configuring the migration source and destination volumes.

The migration procedure is as follows:

- 1) Copy data to the migration destination storage from the migration source storage.
 - 1-1) Copy data to the migration destination storage from the migration source storage per LUN by using the copy function of storage.
 - 1-2) When copying of storage become synchronized, stop services where storage is used. After that, once volumes are stopped, stop copying of storage.
- 2) Back up and delete the GDS configuration from the migration source storage.
 - 2-1) Back up object configuration information. (sdxconfig Backup command)
 - 2-2) Stop volumes. (sdxvolume -F command)
 - 2-3) For a shared disk in the cluster system, convert the migration source class to a local class. (sdxattr -C command)
 - 2-4) Remove the object configuration information. (sdxconfig Remove command)
 - 2-5) For a shared disk in the cluster system, delete the migration source disk from the PRIMECLUSTER resource database. (cldelrsc command)
- 3) Remove the migration source storage from a server and connect the migration destination storage to a server.
- 4) Restore the GDS configuration to the migration destination storage.

4-1) For a shared disk in the cluster system, register the migration destination disk with the PRIMECLUSTER resource database. (clautoconfig command)

4-2) Convert the configuration files backed up in step 2-1) for the migration destination storage as below. (sdxcfg Convert command)

- When copying to one migration destination case from one migration source case in step 1), remove disks from each mirror group so that only one disk will remain in each mirror group.
- Change the physical disk names to names of the migration destination.

4-3) Restore object configuration information. (sdxcfg Restore command)

When copying data from each migration source case to each migration destination case, specify the -e skipsync option at the time of executing the sdxcfg Restore command.

4-4) For a shared disk in the cluster system, convert the classes restored in step 4-3) to shared classes. (sdxattr -C command)

5) Restart the services to use storage.

6) When copying to one migration destination case from one migration source case in step 1) and mirroring the migration destination storage, add a disk to the mirror group of the migration destination. (sdxdisk -M and sdxdisk -C commands)



See

- For the method of using the sdxcfg command, see "[B.1.13 sdxcfg - Object Configuration Operations](#)", "[7.13 Backing Up and Restoring Object Configurations](#)", and "[A.5 Backing Up and Restoring Object Configuration \(sdxcfg\)](#)."
- For the method of using the clautoconfig and eldelrsc commands, see "[6.3 Shared Disk Resource Registration](#)."

8.8 Changing Driver Major Number

This section explains how to change the GDS driver major numbers.

8.8.1 Changing sfdsk Driver Major Number

This section explains how to change the sfdsk driver major numbers.

When the following message is output at the time of system startup, change the sfdsk driver major number to a number other than 487 according to this procedure.

```
ERROR: sfdsk: internal error, func=register_blkdev(sfdsk), errno=16
```

For details of this message, see the error message 26000 in "[C.2.5 Internal Error Messages \(26000-26001\)](#)."

For a cluster system, perform the following procedures on all nodes in the cluster.

1) Stopping volumes in local and shared classes

If a local or shared class exists, stop all volumes in the class. With a shared class registered with a cluster application, stop the cluster application to stop the volumes.

2) Determining a new major number

Choose a number that is not contained in /proc/devices nor /usr/include/linux/major.h. For a cluster system, the major number must be constant on all nodes in the cluster.

3) Changing the major number

To /etc/opt/FJSVsdx/modules/sfdsk.conf, add the "sfdsk_major=*New_Major_Number*," line.

```
#
# Copyright (c) 1998-2001 FUJITSU Ltd.
# All rights reserved.
#
#ident "@(#)sfdsk.conf 41.4 04/10/04 TDM"

name="sfdsk" parent="pseudo";
...
sfdsk_major= New_Major_Number;
```

4) Re-creating device files

4-1) Re-creating device files for control

Re-create the device files `_adm` and `_diag` used by GDS for control.

```
# cd /dev/sfdsk
# ls -l
crw-r--r-- 1 root root Old_Major_Number, 0 May 9 18:47 _adm
crw-r--r-- 1 root root Old_Major_Number, 1 May 9 18:47 _diag
drwxr-xr-x 4 root root 4096 May 13 13:00 Class_Name
...
# rm _adm _diag
# /bin/mknod _adm c New_Major_Number 0
# /bin/mknod _diag c New_Major_Number 1
```

4-2) Checking the re-created device files

Check whether the device files `_adm` and `_diag` used by GDS for control were created correctly.

```
# cd /dev/sfdsk
# ls -l
crw-r--r-- 1 root root New_Major_Number, 0 May 9 18:47 _adm
crw-r--r-- 1 root root New_Major_Number, 1 May 9 18:47 _diag
drwxr-xr-x 4 root root 4096 May 13 13:00 Class_Name
...
```

4-3) Deleting local volume and shared volume device files

If a local or shared class exists, delete volume device files.

```
# cd /dev/sfdsk/Class_Name/dsk
# ls -l
brw----- 1 root root Old_Major_Number, Minor_Number_1 May 13 13:00
Volume_Name_1
brw----- 1 root root Old_Major_Number, Minor_Number_2 May 13 13:00
Volume_Name_2
...
# rm Volume_Name_1 Volume_Name_2 ...
```

The deleted device files will automatically be re-generated when the system is rebooted.

5) Rebooting the system

Reboot the system. For a cluster system, reboot all nodes in the cluster together.

6) Checking the major number

Check whether the sfdsk driver major number was set to the number determined in step 2).

```
# grep sfdsk /proc/devices
...
New_Major_Number_sfdsk
```

7) Checking volume device files

If a local or shared class exists, check whether volume device files were re-generated correctly.

```
# cd /dev/sfdsk/Class_Name/dsk
# ls -l
brw----- 1 root    root      New_Major_Number, Minor_Number_1 May 13 13:00
Volume_Name_1
brw----- 1 root    root      New_Major_Number, Minor_Number_2 May 13 13:00
Volume_Name_2
...
```

8.8.2 Changing sfdsklog Driver Major Number

This section explains how to change the sfdsklog driver major numbers.

When the following message is output at the time of system startup, change the sfdsklog driver major number according to this procedure.

```
sfdsklog: internal error, func=register_chrdev(sfdsklog) fail(-16)
```

For details of this message, see the error message 26001 in "[C.2.5 Internal Error Messages \(26000-26001\)](#)."

1. Determine the major number of sfdsklog driver.

Choose a 300 or higher number which is not contained in /proc/devices.

2. Change the major number of sfdsklog driver.

Add the major number settings of sfdsklog driver (`sdxlog_major=New_Major_Number`) to `/etc/opt/FJVSdx/modules/sfdsklog.conf`.

Example) When changing a major number to 300

```
# All Rights Reserved, Copyright (c) FUJITSU LIMITED 1999
#ident "@(#) $Id: sfdsklog.conf,v 1.1.1.1 2010/10/01 06:06:02 juli Exp $"
name="sfdsklog" parent="pseudo";
sdxlog_major=300;    <- Added
```

3. Restart the system.

8.9 Upgrading Procedure for Dell EMC PowerPath

This section explains how to upgrade PowerPath in the system where the Dell EMC PowerPath device is registered to GDS.

All user operations must be stopped for taking this procedure.

Back up the system disks and the data of customers before taking the following steps.

[Procedure]

1) Displaying class information

Execute the following command on all the nodes to check the configuration information of created classes.

```
# sdxinfo -C
OBJ      NAME      TYPE      SCOPE      SPARE
-----
class    class00  shared   node1:node2  0
class    class01  shared   node1:node2  0
class    class02  local    node1        0
```

2) Stopping applications

Stop all the applications that are accessing the disks of Dell EMC storage units.

Also, execute the `umount(8)` command to unmount the file systems created on the disks of Dell EMC storage units.

3) Stopping RMS

Take the following steps to stop RMS.

3-1) Execute the following command on any one of the nodes where RMS is running to stop RMS.

```
# hvshut -a
```

3-2) Execute the following command on all the nodes to check that RMS is correctly stopped and the message "RMS is not running." is output.

```
# hvdisp -a
RMS is not running.
```

3-3) Set up to avoid the automatic startup of RMS until Dell EMC PowerPath is upgraded.

Execute the following command on all the nodes where RMS is used.

```
# hvsetenv HV_RCSTART 0
```

4) Backing up GDS configuration information

Take the following steps to back up the GDS configuration information.

Specify any directory other than `/tmp` or its subdirectories for `mydir` in the following procedure.

4-1) Back up class configuration information.

For all the classes where the disks of Dell EMC storage units are registered, execute the following command on any one of the nodes that is included in the class scope.

You can check the class scope in the `SCOPE` field output by the `sdxinfo -C` command executed in step 1).

```
# sdxconfig Backup -c classname -o /mydir/classname.bkup
```

4-2) Stop volumes.

For all the classes where the disks of Dell EMC storage units are registered, execute the following command on any one of the nodes that is included in the class scope.

- Shared classes:

```
# sdxvolume -F -c classname -e allnodes
```

- Local classes:

```
# sdxvolume -F -c classname
```

4-3) Change shared classes to local classes.

For all the shared classes where the disks of Dell EMC storage units are registered, execute the following command on any one of the nodes that is included in the class scope.

```
# sdxattr -C -c classname -a type=local
```

4-4) Remove classes.

For all the classes where the disks of Dell EMC storage units are registered, execute the following command on the same node as step 4-3).

```
# sdxconfig Remove -c classname
```

5) Switching to single-user mode

Take the following steps to switch the mode to single-user mode.

5-1) Check the current default target.



Note

The default target (multi-user.target in the example below) changes depending on an environment. Write down the current default target to change back the default target later.

(Example)

```
# systemctl get-default  
multi-user.target
```

5-2) Change the default target.

```
# systemctl set-default rescue.target
```

5-3) Restart the system in single-user mode.

```
# shutdown -r now
```

6) Upgrading PowerPath

PowerPath must be upgraded by Dell EMC engineers.

After PowerPath is upgraded, check that both the emcpower device and the native device name are not changed before and after the upgrade.

If these device names have been changed, contact Dell EMC engineers to change them back to the original names before the upgrade.

7) Switching to multi-user mode

7-1) Change the default target.

```
# systemctl set-default multi-user.target
```

7-2) Restart the system.

```
# shutdown -r now
```

8) Setting excluded device list

If the excluded device list is not created, create it.

Take this step on all the nodes where the disks of Dell EMC storage units are registered to the class.

For how to create the excluded device list, see "[6.1.1 Excluded Device List](#)" and "[6.1.3 To Use Dell EMC Storage Units.](#)"

9) Restoring GDS configuration information

Take the following steps to restore the GDS configuration information.

9-1) Restore the class configuration information.

For all the classes where the disks of Dell EMC storage units are registered, execute the following command on any one of the nodes that is included in the class scope.

You can check the class scope in the SCOPE field output by the `sdxinfo -C` command executed in step 1).

```
# sdxconfig Restore -c classname -i /mydir/classname.bkup
```

When the above command is executed, the classes will be restored as local classes on the node where the command was executed.

9-2) Restart the system.

Execute the following command on the same node as step 9-1).

```
# shutdown -r now
```

9-3) Change local classes to shared classes.

For all the shared classes where the disks of Dell EMC storage units are registered, change the local classes that were restored in step 9-1) to shared classes.

For the classes whose TYPE fields are listed as shared in the output of the `sdxinfo -C` command when it was executed in step 1), execute the following command on the same node as step 9-1).

```
# sdxattr -C -c classname -a type=shared,scope=nodename1:nodename2
```

9-4) Set up the classes to be used as RMS resources.

For all the shared classes where the disks of Dell EMC storage units are registered, execute the following command on any one of the nodes that is included in the class scope.

```
# /opt/SMAW/SMAWRrms/bin/hvgsdsetup -a classname
```

10) Checking restored classes

Execute the following command on all the nodes to check that all the restored classes are correctly displayed.

```
# sdxinfo -C
```

11) Starting RMS

Take the following steps to start RMS.

11-1) Execute the following command on any one of the nodes to start RMS.

```
# hvcm -a
```

11-2) Execute the following command on all the nodes to check that RMS has correctly started.

```
# hvdisp -a
```

11-3) To start RMS automatically during node startup, execute the following command on all the nodes where RMS is used.

```
# hvsetenv HV_RCSTART 1
```

12) Starting applications

Execute the `mount(8)` command to mount the file systems on the disks of Dell EMC storage units.

Also, start the applications that use the disks of Dell EMC storage units.

8.10 Changing IP Addresses Used for Mirroring among Servers

This section describes the procedures to change IP addresses used for mirroring among servers.

Change both nodes to single user mode, and then perform the procedures.

Because the console cannot be used in a cloud environment (except an FJcloud-Baremetal environment), do not select single user mode. In this case, perform step 1 to 3 to restrict the automatic startup of the PRIMECLUSTER service.

In an environment other than a cloud environment (except an FJcloud-Baremetal environment), follow the procedure from step 4.

1. Stop RMS.

If RMS is running, execute the following command on any one node to stop RMS.

```
# hvshut -a
```


Note

If RMS is stopped on all nodes during the synchronization copying of the GDS volume, the synchronization copying of the entire volume area is performed after the corrections are applied and all nodes are restarted.

If you do not want to perform the synchronization copying of the entire area of volume, stop RMS after the synchronization copying is completed.

To check the slice status of the GDS volume, execute the following command.

Execute the following command on any one node to check the value of the STATUS field of the command output.

The status of the copy destination slice is COPY during the synchronization copying, and after copying is complete, the status becomes ACTIVE or STOP.

```
# sdxinfo -S
```

2. Restrict the automatic startup of the PRIMECLUSTER service.

Restrict the automatic startup of the PRIMECLUSTER service by executing the following command on all nodes.

```
# /opt/FJSPVpclinst/bin/pclservice off
```

3. Restart the system.

Restart the system on all nodes.

```
# /sbin/shutdown -r now
```

4. Delete the registration of the target portal on the local node.

Example) when IP address for the mirroring among servers is 192.168.56.20

```
# iscsiadm -m discovery --op delete --portal 192.168.56.20
```

5. Start the network service.

For the cloud environment (except an FJcloud-Baremetal environment), this procedure is unnecessary because the network service has already been started.

- For RHEL7

```
# systemctl start network.service
```

The following message may be output but, no action is required because the operation is not affected.

```
Error getting authority: Error initializing authority: Could not connect:  
No such file or directory (g-io-error-quark, 1)
```

- For RHEL8 or later

```
# systemctl start NetworkManager.service
```

6. Copy the setting information file of the iSCSI target.

Example) when the file name for the copy destination is /var/tmp/targetfile.json

```
# cp /etc/target/saveconfig.json /var/tmp/targetfile.json
```

7. Change the IP address described in the copy destination file in step 6.

Example) When changing the IP address for the mirroring among servers to 192.168.56.21

```
# vim /var/tmp/targetfile.json
```

[Before changing]

```
{
  "fabric_modules": [],
  "storage_objects": [
    {
      ...
      "portals": [
        {
          "ip_address": "192.168.56.20",
          "iser": false,
          "port": 3260
        }
      ],
      ...
    }
  ]
}
```

[After changing]

```
{
  "fabric_modules": [],
  "storage_objects": [
    {
      ...
      "portals": [
        {
          "ip_address": "192.168.56.21",
          "iser": false,
          "port": 3260
        }
      ],
      ...
    }
  ]
}
```

8. Reflect the setting information file of the iSCSI target edited in step 7 to the target.

Example) when the file name is /var/tmp/targetfile.json

```
# targetctl restore /var/tmp/targetfile.json
```

The following message may be output but, no action is required because the operation is not affected.

Unable to load target_core_user

9. Confirm that the iSCSI target was set correctly.

```
# targetcli /iscsi/ ls
```

 **Point**

(*1) Confirm that the modified IP address was reflected.

[Output example]

```
o- iscsi ..... [Targets: 1]
  o- iqn.2003-01.org.linux-iscsi.nodel.x8664:sn.6e665c7c1be0 ..... [TPGs: 1]
    o- tpg1 ..... [no-gen-acls, no-auth]
      o- acls ..... [ACLs: 2]
        | o- iqn.1994-05.com.redhat:dee92ff9979d ..... [Mapped LUNs: 1]
        | | o- mapped_lun0 ..... [lun0 block/store1 (rw)]
        | o- iqn.1994-05.com.redhat:fa7eb9cf483c ..... [Mapped LUNs: 1]
        | | o- mapped_lun0 ..... [lun0 block/store1 (rw)]
      o- luns ..... [LUNs: 1]
        | o- lun0 ..... [block/store1 (/dev/disk/by-id/scsi-3500000e111c56610)]
      o- portals ..... [Portals: 2]
        o- 192.168.56.21:3260 ..... [OK] (*1)
```

10. Store the target information that was restored in step 8.

```
# targetctl save
```

11. Change the IP addresses described in the `/etc/opt/FJSVsdx/.sdxnetmirror_ipaddr` file.

[Before changing]

```
192.168.56.10
192.168.56.20
```

[After changing]

```
192.168.56.11
192.168.56.21
```

12. Establish iSCSI session.

For the procedure, refer to "[4.8.4 Establishing iSCSI Session.](#)"

Step 8 does not need to be performed.

13. Restart both nodes in multi-user mode.

In a cloud environment (except an FJcloud-Baremetal environment), take the following procedure to cancel restricting the automatic startup of the PRIMECLUSTER service.

1. Cancel restricting the automatic startup of the PRIMECLUSTER service.

Execute the following command on all nodes and cancel the PRIMECLUSTER service settings restricted in step 2.

```
# /opt/FJSVpclinst/bin/pclservice on
```

2. Restart the system.

Restart the system on all nodes.

```
# /sbin/shutdown -r now
```

8.11 Adding Disks Used for Mirroring among Servers

When adding disks used for mirroring among servers, configure iSCSI device according to this procedure.

When adding disks to both nodes, after performing this procedure on a node, perform the same procedure on the other node.

1. Add disks.
2. On the node where the disks are added, modify the target settings to add the information of the added disks.
 - 2-1. Execute `targetcli` to enter the interactive mode.

```
# targetcli
```

- 2-2. Register a disk or a partition of a local node where the mirroring among servers is performed.

```
/> /backstores/block/ create storage_name device_path
```



- Specify any character string for `storage_name`.
- For `device_path`, specify the path for the by-id file of a disk or a partition on a local node which is used for the mirroring among servers.

In an Azure environment, a NIFCLOUD environment, or an environment where the by-id file does not exist, do not specify the

path for the by-id file, but create a partition in the disk where the gpt disk label is set, and specify the path for the by-partuuid file for that partition.

.....

Execution example

```
/> /backstores/block/ create storage1 /dev/disk/by-id/scsi-3500000e111c56611
```

2-3. Confirm iqn name.

```
/> /iscsi/ ls
```

Execution example

```
/> /iscsi/ ls
o- iscsi ..... [Targets: 1]
  o- iqn.2003-01.org.linux-iscsi.node1.x8664:sn.2613f8620d98 ..... [TPGs: 1]
    o- tpg1 ..... [no-gen-acls, no-auth]
      o- acls ..... [ACLs: 0]
      o- luns ..... [LUNs: 0]
      o- portals ..... [Portals: 0]
```

This example shows that iqn name is "iqn.2003-01.org.linux-iscsi.node1.x8664:sn.2613f8620d98".

2-4. Connect the device registered in step 2-2.

```
/> /iscsi/iqn_name/tpgX/luns create /backstores/block/storage_name
```



-
- Specify the iqn name confirmed in step 2-3 to *iqn_name*.
 - Specify the directory name (such as tpg1) created in *iqn_name* directory to *tpgX*.
 - Specify the storage name registered in step 2-2 to *storage_name*.
-

Execution example

```
/> /iscsi/iqn.2003-01.org.linux-iscsi.node1.x8664:sn.2613f8620d98/tpg1/luns create \
/backstores/block/storage1
```

2-5. End the interactive mode of targetcli. This operation records the settings.

```
/> exit
```

3. Execute the following on both nodes to rescan iSCSI session.

```
# iscsiadm -m session --rescan
```

4. Confirm that LUN was added on both nodes.

```
# iscsiadm -m session -P 3
```

Execution example

```
# iscsiadm -m session -P 3
iSCSI Transport Class version xxxxxxxx
version xxxxxxxx
Target: iqn.2003-01.org.linux-iscsi.node2.x8664:sn.c585c8864f22
  Current Portal: 192.168.56.10:3260,1
  Persistent Portal: 192.168.56.10:3260,1
  *****
  Interface:
  *****
```

```

...
    iSCSI Connection State: LOGGED IN
    iSCSI Session State: LOGGED_IN
    Internal iscsid Session State: NO CHANGE
...
    *****
    Attached SCSI devices:
    *****
    Host Number: 12 State: running
    scsil2 Channel 00 Id 0 Lun: 0
        Attached scsi disk sdd State: running
...
Target: iqn.2003-01.org.linux-iscsi.node1.x8664:sn.2613f8620d98  (*1)
    Current Portal: 192.168.56.20:3260,1
    Persistent Portal: 192.168.56.20:3260,1
    *****
    Interface:
    *****
...
    iSCSI Connection State: LOGGED IN
    iSCSI Session State: LOGGED_IN
    Internal iscsid Session State: NO CHANGE
...
    *****
    Attached SCSI devices:
    *****
    Host Number: 13 State: running
    scsil3 Channel 00 Id 0 Lun: 0
...
    scsil3 Channel 00 Id 0 Lun: 2                (*2)
        Attached scsi disk sde State: running    (*3)
...

```

This example shows that sde (*3) for LUN2 of (*2) was added in the target "iqn.2003-01.org.linux-iscsi.node1.x8664:sn.2613f8620d98" of (*1).

Chapter 9 Removing Configuration

This section explains how to remove the configuration of GDS, GDS Snapshot, and GDS I/O Monitor Option.

9.1 Breaking Proxy

This subsection describes the procedures for breaking the relationships between masters and proxies in GDS Snapshot installed systems.

9.1.1 Using GDS Management View

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 9.1 Break



3. Confirming

A confirmation dialog box appears asking you whether to break the proxy.

Figure 9.2 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 9.3 Information Message of "Break" Completion



Click <OK> to close the information message window.

9.1.2 Using Commands

Break the relationships between masters and proxies by using the `sdxproxy Break` command.



See

.....
For details on the `sdxproxy Break` command, see "[B.2.1 sdxproxy - Proxy Object Operations.](#)"
.....

9.2 Removing File System

When removing a file system that is created in a volume, remove the file system configuration from the `/etc/fstab` file.



See

.....
For details on how to unset file systems on shared disk units used in a cluster system, see the "[PRIMECLUSTER Installation and Administration Guide.](#)"
.....

9.3 Removing Volumes, Groups, and Classes


9.3.1 Using GDS Management View

When file system is not created, you can start the unmirroring process from removing the volume.

9.3.1.1 Operation Flow

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

Figure 9.4 Configuration Changing Operation



9.3.1.2 Removing Volume

In this section, how to remove a volume is explained.

1. Stopping an operation

Stop the operation which is using the volume.

Unmount the file system on the volume when it has been mounted.

2. Removing the settings of a file system or an application

When a file system is created in the volume, delete the setting of the volume and the file system in the `/etc/fstab` file.

If the setting of a file system or an application other than `/etc/fstab` is performed, delete them.

3. Stopping the volume

When removing a shared class volume, stop the volume from `[Operation]:[Stop Volume]` on all nodes which belongs a class scope.

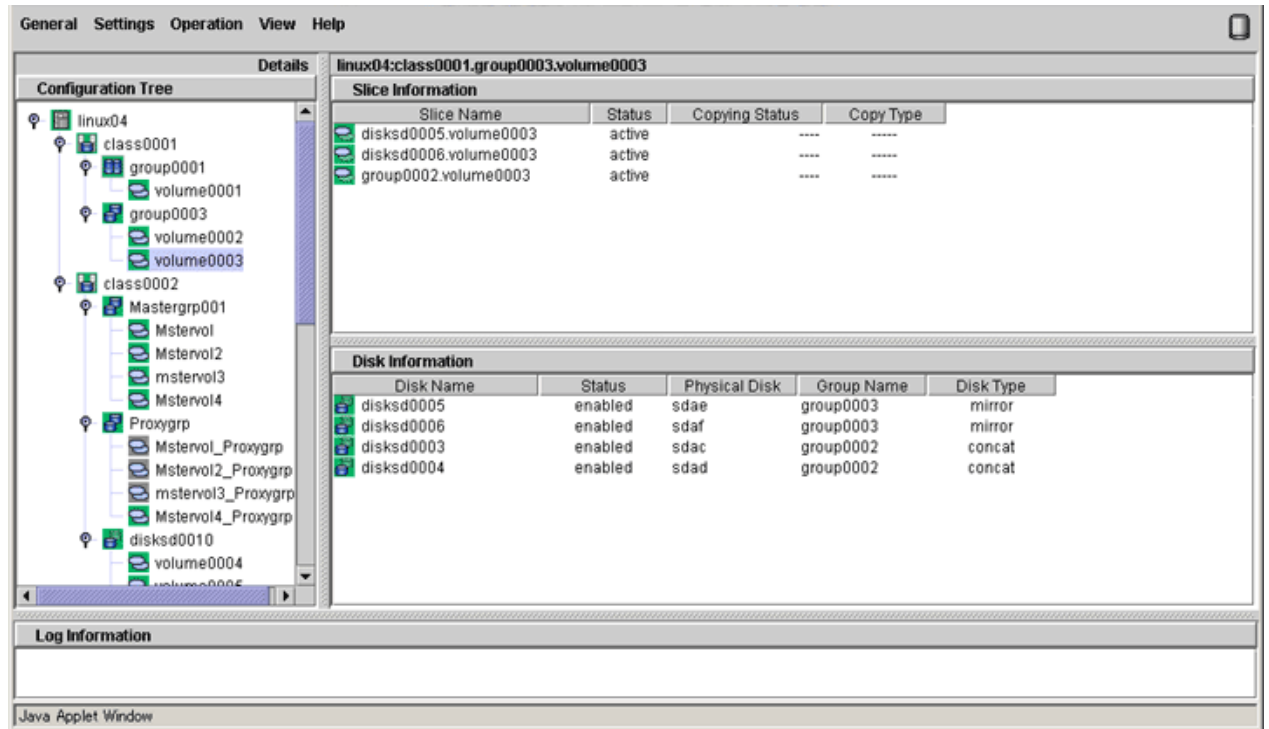
4. Confirming the volume status

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 9.5 Displaying Volume Information



5. 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 "[7.7.1.2 Operation Procedure](#)."

6. 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.




Note

Size displayed in the <Unused> field

The size displayed in the <Unused> field is the maximum size that can be created as a single volume.

Depending on the position of the volume to be removed, the maximum size and consequently the size displayed in the <Unused> field may remain unchanged.

Figure 9.6 Removing Volume



9.3.1.3 Removing 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 ["9.3.1.2 Removing 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.

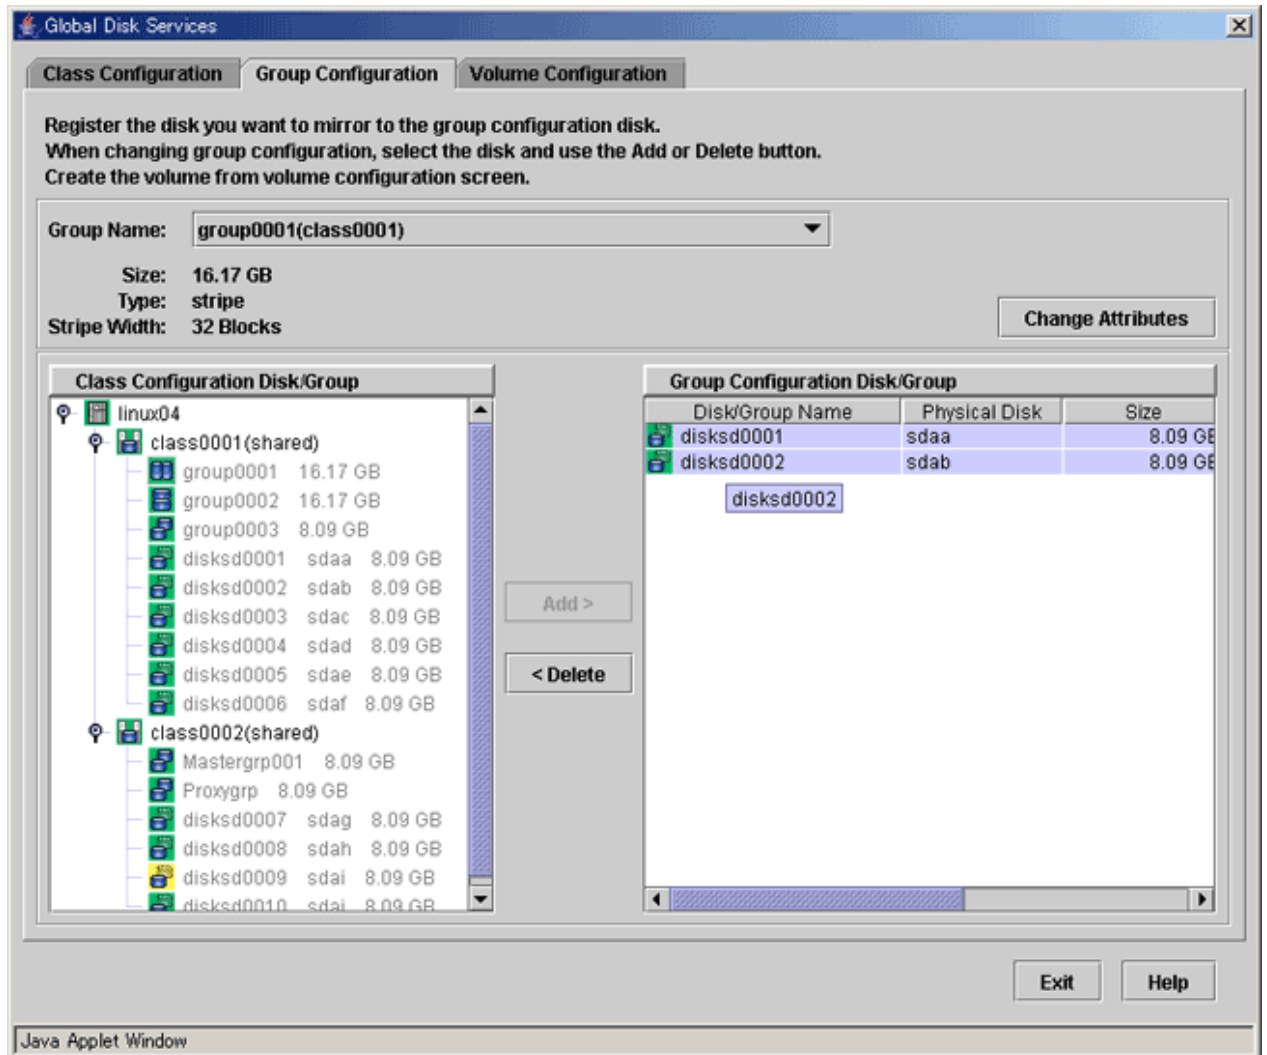
2. **Selecting group to remove**

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 9.7 Disconnecting Disk/Lower level group



9.3.1.4 Removing 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 "9.3.1.3 Removing 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.

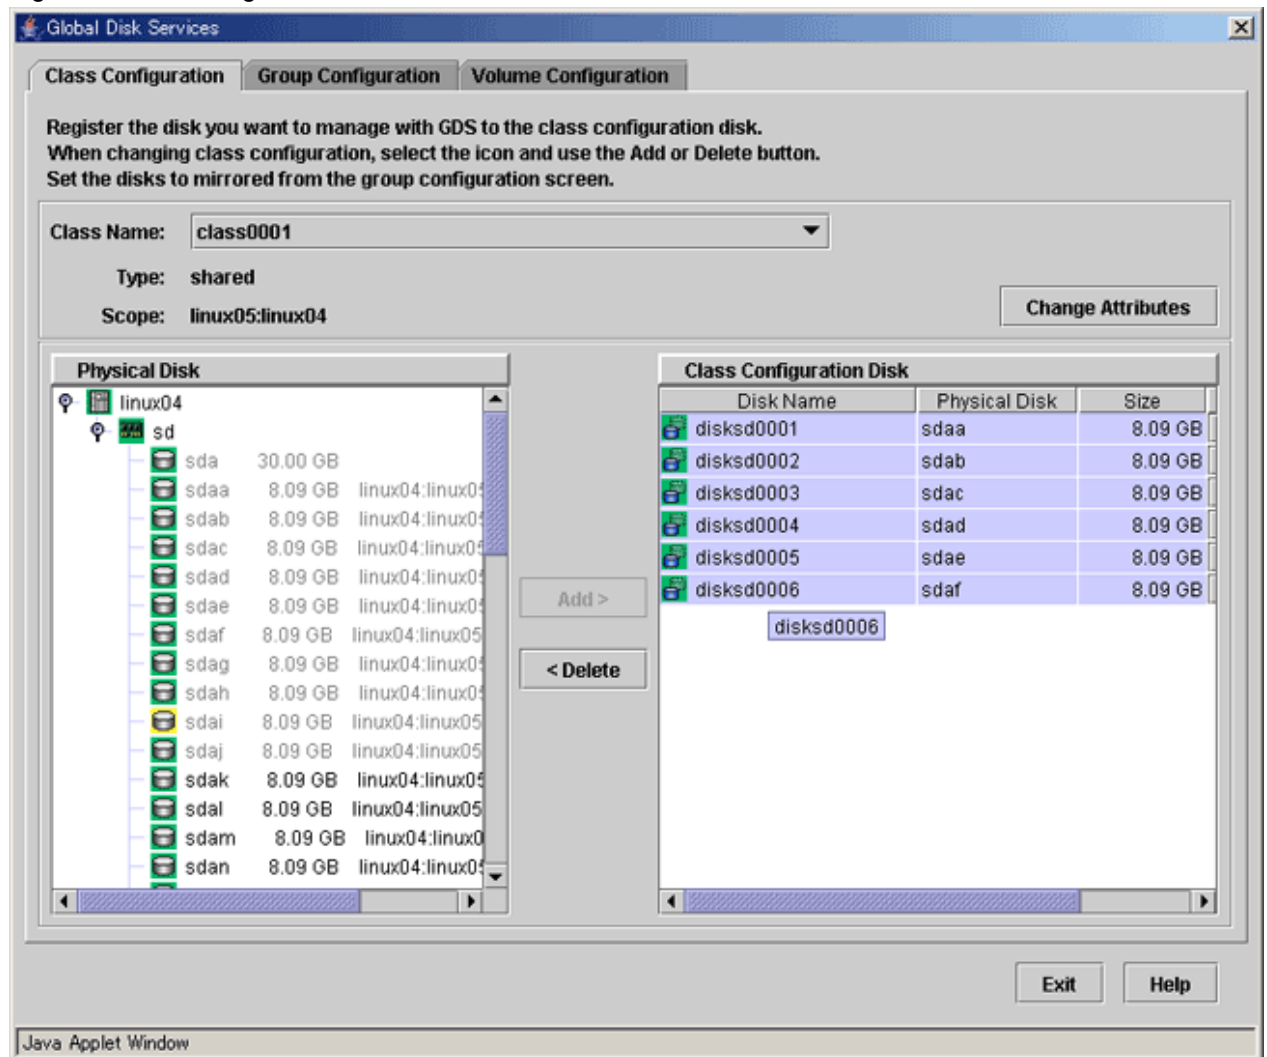
2. Selecting class to remove

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 9.8 Removing Disk



Note

In Cluster Systems

If a class resource is registered to the cluster application, delete the resource from the cluster application and then delete the class. For the method for deleting resources from cluster applications, refer to the "PRIMECLUSTER Installation and Administration Guide."

9.3.2 Using Commands

When removing volumes, groups, and classes by using the commands, use the following commands.

- Removing volumes: `sdxvolume -R`
- Removing groups: `sdxgroup -R`
- Removing classes: `sdxclass -R`



See

For details on each command, see "[Appendix B Command Reference](#)."



Note

[IOmonitor]

If you use the I/O Response Time Assurance Function and use the `sdxclass -R` command to delete a class that has been enabled to system panic mode, disable system panic mode before you execute the `sdxclass -R` command. For details on how to change the system panic mode, see "[8.3.2 Changing the System Panic Mode](#)."

9.4 System Disk Mirroring Cancellation [EFI]

This section explains how to unmirror the system disk.



Note

[IOmonitor]


If you use the I/O Response Time Assurance Function and system panic mode is enabled for the root class, disable system panic mode for the root class before you cancel the system disk mirroring. For details on how to change the system panic mode, see "[8.3.2 Changing the System Panic Mode](#)."

9.4.1 Unmirroring System Disk with GDS Management View

9.4.1.1 Unmirroring the System Disk [EFI]

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

Figure 9.9 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.

9.4.1.2 Operation Flow

In this section, how to unmirror a system disk is explained.

Note

When You Cannot Unmirror the System Disk

Under the following conditions, the system disks cannot be unmirrored. To avoid these conditions, restore the disk status or change the configuration first, and then unmirror the system disks.

- In the system disk group, a disk has been disconnected with [Swap Physical Disk].
- In the system disk group, a disk is in disabled status.
- All the disks that are connected with the system disk group contain a slice that is not active.
- Among the groups in the root class, any groups other than the system disk group have volumes.

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

Select [Settings]: [Unmirror System Disk] in the Main screen to display system disk mirroring configurations.

Figure 9.10 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.

Note

The disk displayed in the [Mirror Disk] field cannot be used as a system disk after unmirroring.

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 9.11 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 9.12 Unmirroring System Disk Completion Screen



Confirm that unmirroring of system disk is complete and click <OK>.

Figure 9.13 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

For the cluster system

After unmirroring the system disk on one node, restart the node in multi-user mode immediately. Do not restart multiple nodes at the same time.

When unmirroring the system disk on multiple nodes in the cluster system and the nodes are restarted at the same time, a object in the root class on some nodes may be left without removed. In this case, follow "[9.3 Removing Volumes, Groups, and Classes](#)" to remove the object in the root class.



Note

Definition of devices in the `/etc/fstab` file

After unmirroring the system disk, the devices of the system volumes (`/`, `/var`, `usr`, `boot`, `/boot/efi`, and swap area) are defined with the UUID format in the `/etc/fstab` file.

When you want to change the definition of devices, modify the `/etc/fstab` file by using the editor such as `vim(1)` after unmirroring the system disk.

9.4.2 Cancellation with Commands

Assuming the disk on which `/` (`root`), `/usr`, `/var`, `/boot`, and `/boot/efi` 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.



Note

- In this version, to mirror `/` (`root`), `/usr`, `/var`, `/boot`, `/boot/efi`, or swap areas, these areas should be located in the same disk.
- In this version, cancel the system disk mirroring from the GDS Management View, not by using the command. For details, see "[9.4.1 Unmirroring System Disk with GDS Management View](#)."

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`, `/var`, `/boot`, and `/boot/efi` is "sda", and the disk allocated as a swap area is "sdb."

```
# sdxdisk -M -c System -a type=root -d sda=Root1:keep,  
sdc=Root2:undef,sdb=Swap1:keep,sdd=Swap2:undef
```

3. Connect the system disks to a group.

```
# sdxdisk -C -c System -g Group1 -d Root1,Root2 -v 1=root:on,  
2=usr:on,3=var:on,4=home:on,5=boot:on,6=efi:on  
# sdxdisk -C -c System -g Group2 -d Swap1,Swap2 -v 1=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 not displayed with the `mount(8)` command may not be open.

4. Check that mirroring definition is completed.

```
# sdxroot -M -c System -d Root1,Swap1
```

5. Reboot the system.

```
# shutdown -r now
```

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 -r now
```

5. Check that the mirroring has been canceled.

Using the `mount` command and the `sdxinfo` command, make sure that the system disk mirroring has been canceled 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 boot
# sdxvolume -F -c System -v efi
# 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 boot
# sdxvolume -R -c System -v efi
# 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
```

9.5 OS Re-installation

On the system where GDS object such as class is created, even when you reconfigure the system from the reinstallation of the OS, delete the GDS settings as described in "9.3 Removing Volumes, Groups, and Classes" and "9.4 System Disk Mirroring Cancellation [EFI]."

If you re-install the operating system without deleting the GDS settings, execute the following restoration procedure.

Restoration procedure

1) If you re-install the operating system without canceling the system disk mirroring, delete the root class information.

If GDS is installed without taking this procedure, it may take about 10 minutes to start GDS. In this case, it takes time to start the cluster application in which the class resources are registered.

1-1) Delete the configuration data from the EFI environment variables.

```
# cd /sys/firmware/efi/vars/
# cat sysvol-db2-*/raw_var > del_var
# cat sysvol-db2-*/data
-> Make sure that nothing is displayed.
```

1-2) Delete SDX disk entries from the boot option of the EFI boot manager.

1-2-1) Check the current setting of the EFI boot manager.

```
# efibootmgr
BootCurrent: 0000
Timeout: 5 seconds
BootOrder: 0001,0002,0000
Boot0000* sda
Boot0001* RootDisk1
Boot0002* RootDisk2
```

In this example, the SDX disk entries are RootDisk1 of the number 0001 and RootDisk2 of the number 0002.

1-2-2) Remove the SDX disk entries from the boot option.

```
# efibootmgr -B -b 0001
BootCurrent: 0000
Timeout: 5 seconds
BootOrder: 0002,0000
Boot0000* sda
Boot0002* RootDisk2
# efibootmgr -B -b 0002
BootCurrent: 0000
Timeout: 5 seconds
BootOrder: 0000
Boot0000* sda
```

1-3) Remove the private slice of the disk that has been registered in the root class.

An example below shows the procedure when the physical disks sda and sdb have been registered in the root class, and the physical slices sda8 and sdb8 have been the private slices before the OS is reinstalled.

1-3-1) Execute the parted command.

```
# parted /dev/sda
```



See

For details of the parted command, see the manual of parted(8).

1-3-2) Display the partition table using the print subcommand.

```
(parted) print
Disk /dev/sdf: 34.4GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start   End     Size    File system  Name              Flags
 1      1049kB  630MB   629MB   fat32        EFI System Partition  boot, esp
 2       630MB  15.0GB  14.4GB   xfs
 3      15.0GB  17.0GB  2000MB   xfs
 4      17.0GB  19.0GB  2001MB   xfs
 5      19.0GB  21.0GB  2000MB   xfs
 6      21.0GB  23.0GB  2001MB   xfs
 7      23.0GB  25.0GB  2000MB   linux-swap(v1)      swap
 8      25.0GB  25.0GB  10.5MB           <== Private slice
```

1-3-3) Remove the private slice.

```
(parted) rm 8
```

1-3-4) Remove the private slice sdb8 as well.

Take step 1-3-2) and 1-3-3) for the physical disk sdb.

1-3-5) Finish the parted command.

```
(parted) quit
```

1-4) Restart the system.

2) The following errors may occur when you re-install the operating system and GDS without deleting the GDS setting and setup GDS.

```
SDX:sdxdisk: ERROR: physical_slice_name is private slice
SDX:sdxdisk: ERROR: physical_disk_name: configuration information exists in private slice
```

In this case, setup GDS after deleting private slices with the following procedure;

2a) When the error message is output for the physical disk that has been registered in the root class before the OS is re-installed.

Take step 1-3).

2b) When the error message is output for the physical disk that has been registered in the shared class or in the local class before the OS is re-installed.

As an example, the procedure when *physical_disk_name* is "sdb" is described below.

```
# dd if=/dev/zero of=/dev/sdb count=35
```

9.6 Deleting iSCSI Settings Used for Mirroring among Servers

The section describes how to delete iSCSI settings used for mirroring among servers.

9.6.1 Delete iSCSI devices

This section describes the procedure to delete the iSCSI devices used for the mirroring among servers.

Perform this procedure after deleting iSCSI devices to be deleted from a class of GDS and the resource database of PRIMECLUSTER.

For details, refer to the steps "Change the device names set in resources of the shared disk device" and "Delete resources of the shared disk device" in the procedure "Deleting a shared disk device" described in "PRIMECLUSTER Installation and Administration Guide."

1. Confirm LUN to be deleted on any node.

```
# iscsiadm -m session -P 3
```

[Execution example]

```
# iscsiadm -m session -P 3
iSCSI Transport Class version xxxxxxxx
version xxxxxxxx
Target: iqn.2003-01.org.linux-iscsi.node2.x8664:sn.c585c8864f22
  Current Portal: 192.168.56.10:3260,1
  Persistent Portal: 192.168.56.10:3260,1
  *****
  Interface:
  *****
...
  iSCSI Connection State: LOGGED IN
  iSCSI Session State: LOGGED_IN
  Internal iscsid Session State: NO CHANGE
...
  *****
  Attached SCSI devices:
  *****
  Host Number: 12 State: running
  scsi12 Channel 00 Id 0 Lun: 0
    Attached scsi disk sdd State: running
  scsi13 Channel 00 Id 0 Lun: 1
...
Target: iqn.2003-01.org.linux-iscsi.node1.x8664:sn.6e665c7c1be0          (*1)
  Current Portal: 192.168.56.20:3260,1
  Persistent Portal: 192.168.56.20:3260,1
  *****
  Interface:
  *****
...
  iSCSI Connection State: LOGGED IN
  iSCSI Session State: LOGGED_IN
  Internal iscsid Session State: NO CHANGE
...
  *****
  Attached SCSI devices:
  *****
  Host Number: 13 State: running
  scsi13 Channel 00 Id 0 Lun: 0                                     (*2)
    Attached scsi disk sde State: running                          (*3)
  scsi13 Channel 00 Id 0 Lun: 1
...

```

In the example, when deleting sde of (*3), LUN0 of (*2) is deleted in the target name "iqn.2003-01.org.linux-iscsi.node1.x8664:sn.6e665c7c1be0" of (*1).

2. Stop iSCSI devices to be deleted on both node.

```
# echo offline > /sys/block/iSCSI_device_name/device/state
```

[Execution example]

```
# echo offline > /sys/block/sde/device/state
```

3. Delete iSCSI devices to be deleted on both node.

```
# echo 1 > /sys/block/iSCSI_device_name/device/delete
```

[Execution example]

```
# echo 1 > /sys/block/sde/device/delete
```

4. Confirm the storage for LUN to be deleted on a node where LUN is to be deleted.

```
# targetcli ls
```

Point

(*1) Confirm that the target name is same as which was confirmed in step 1.

(*2) Confirm the storage name correspond to the LUN to be deleted.

[Execution example]

```
# targetcli ls
o- / ..... [ ... ]
  o- backstores ..... [ ... ]
    | o- block ..... [Storage Objects: 2]
    | | o- storage1 .. [/dev/disk/by-id/scsi-3500000e114784363 (8.0GiB) write-thru activated]
    | | o- storage2 .. [/dev/disk/by-id/scsi-3500000e1109cc830 (16.0GiB) write-thru activated]
    | o- fileio ..... [Storage Objects: 0]
    | o- pscsi ..... [Storage Objects: 0]
    | o- ramdisk ..... [Storage Objects: 0]
  o- iscsi ..... [Targets: 1]
    | o- iqn.2003-01.org.linux-iscsi.nodel.x8664:sn.6e665c7c1be0 ..... [TPGs: 1] (*1)
    |   o- tpg1 ..... [no-gen-acls, no-auth]
    |     o- acls ..... [ACLs: 2]
    |       | o- iqn.1994-05.com.redhat:dee92ff9979d ..... [Mapped LUNs: 2]
    |         | | o- mapped_lun0 ..... [lun0 block/storage1 (rw)]
    |         | | o- mapped_lun1 ..... [lun1 block/storage2 (rw)]
    |         | o- iqn.1994-05.com.redhat:fa7eb9cf483c ..... [Mapped LUNs: 2]
    |           | o- mapped_lun0 ..... [lun0 block/storage1 (rw)]
    |           | o- mapped_lun1 ..... [lun1 block/storage2 (rw)]
    |     o- luns ..... [LUNs: 2]
    |       | o- lun0 ..... [block/storage1 (/dev/disk/by-id/scsi-3500000e114784363)] (*2)
    |       | o- lun1 ..... [block/storage2 (/dev/disk/by-id/scsi-3500000e1109cc830)]
    |     o- portals ..... [Portals: 1]
    |       o- 192.168.125.123:3260 ..... [OK]
  o- loopback ..... [Targets: 0]
```

5. Delete the storage for LUN to be deleted on a node where LUN is to be deleted.

```
# targetcli /backstores/block/ delete storage_name
```

Point

Specify the storage name that was confirmed in step 4 to *storage_name*.

[Execution example]

```
# targetcli /backstores/block/ delete storage1
```

6. Confirm that LUN was deleted on a node where LUN is to be deleted.

```
# targetcli ls
```



Confirm that the storage and LUN information which were deleted were not output.

[Output example]

```
# targetcli ls
o- / ..... [..]
  o- backstores ..... [..]
    | o- block ..... [Storage Objects: 1]
    | | o- storage2 . [/dev/disk/by-id/scsi-3500000e1109cc830 (16.0GiB) write-thru activated]
    | o- fileio ..... [Storage Objects: 0]
    | o- pscsi ..... [Storage Objects: 0]
    | o- ramdisk ..... [Storage Objects: 0]
  o- iscsi ..... [Targets: 1]
    | o- iqn.2003-01.org.linux-iscsi.node1.x8664:sn.6e665c7c1be0 ..... [TPGs: 1]
    |   o- tpg1 ..... [no-gen-acls, no-auth]
    |     o- acls ..... [ACLs: 2]
    |       | o- iqn.1994-05.com.redhat:dee92ff9979d ..... [Mapped LUNs: 1]
    |         | | o- mapped_lun1 ..... [lun1 block/storage2 (rw)]
    |         | o- iqn.1994-05.com.redhat:fa7eb9cf483c ..... [Mapped LUNs: 1]
    |           | o- mapped_lun1 ..... [lun1 block/storage2 (rw)]
    |     o- luns ..... [LUNs: 1]
    |       | o- lun1 ..... [block/storage2 (/dev/disk/by-id/scsi-3500000e1109cc830)]
    |     o- portals ..... [Portals: 1]
    |       o- 192.168.125.123:3260 ..... [OK]
  o- loopback ..... [Targets: 0]
```

7. Save the target information on a node where LUN is to be deleted.

```
# targetctl save
```

9.6.2 Deleting iSCSI Initiator Settings

When deleting the configuration of the mirroring among servers, after deleting the objects of the mirroring among servers, iSCSI initiator settings must be deleted by the following procedure.

1. Confirm the iSCSI target name used for the mirroring among servers.

Confirm the target name by executing the following command on both nodes.

```
# targetcli ls iscsi 1
```

[Output example]

```
o- iscsi ..... [Targets: 1]
  o- iqn.2003-01.org.linux-iscsi.node1.x8664:sn.c053e5d35899 ..... [TPGs: 1]
```

This output example shows that "iqn.2003-01.org.linux-iscsi.node1.x8664:sn.c053e5d35899" is the target name.

2. Confirm that the iSCSI target information used for the mirroring among servers on any node.

```
# iscsiadm -m node
```

[Output example]

```
192.168.56.10:3260,1 target-0201
192.168.56.20:3260,1 target-0101
```

In this example, the target name and the portal name are as below.

Target name	Portal name
target-0201	192.168.56.10:3260,1
target-0101	192.168.56.20:3260,1

3. Log out from the iSCSI session used for the mirroring among servers on both nodes.

```
# iscsiadm -m node -T targetname --portal portalname --logout
```

Specify the target name and the portal name confirmed in step 2 to *targetname* and *portalname*.

In the example, execute the following 2 commands on each node.

[Execution Example]

```
# iscsiadm -m node -T target-0201 --portal 192.168.56.10:3260,1 --logout
# iscsiadm -m node -T target-0101 --portal 192.168.56.20:3260,1 --logout
```

4. Delete the iSCSI initiator information for the mirroring among servers on both nodes.

```
# iscsiadm -m node -o delete -T targetname --portal portalname
```

Specify the target name and the portal name confirmed in step 2 to *targetname* and *portalname*.

In this example, execute the following 2 commands on each node.

[Execution example]

```
# iscsiadm -m node -o delete -T target-0201 --portal 192.168.56.10:3260,1
# iscsiadm -m node -o delete -T target-0101 --portal 192.168.56.20:3260,1
```

9.6.3 Deleting iSCSI Target Settings

When deleting all the configurations of the mirroring among servers, after deleting the iSCSI initiator settings on all the node, iSCSI target settings must be deleted by the following procedure.



See

For the deletion procedure of iSCSI initiator settings, refer to "9.6.2 Deleting iSCSI Initiator Settings."

1. Delete the iSCSI target on both nodes.

Execute the following command on both nodes.

```
# targetcli /iscsi delete target_name
```

[Execution example]

```
# targetcli /iscsi delete iqn.2003-01.org.linux-iscsi.node1.x8664:sn.2613f8620d98
```

2. Delete the storage settings on both nodes.

Execute the following command on both nodes.

If there are multiple storages, execute this command for each storage.

```
# targetcli /backstores/block delete storage_name
```

[Execution example]

```
# targetcli /backstores/block delete storage1
```

3. Confirm that the setting was deleted on both nodes.

Execute the following command on both nodes.

```
# targetcli ls
```

[Output example]

```
o- / ..... [..]
  o- backstores ..... [..]
    | o- block ..... [Storage Objects: 0]
    | o- fileio ..... [Storage Objects: 0]
    | o- pscsi ..... [Storage Objects: 0]
    | o- ramdisk ..... [Storage Objects: 0]
  o- iscsi ..... [Targets: 0]
  o- loopback ..... [Targets: 0]
```

4. Save the target information on both nodes.

Execute the following command on both nodes.

```
# targetctl save
```

5. Delete the directory for setting related to the start and stop order for iSCSI target service.

Execute the following command on both nodes.

```
# rm -rf /etc/systemd/system/fjsvsdx.service.d/
```

6. Apply the start and stop order for iSCSI target service.

Execute the following command on both nodes.

```
# systemctl daemon-reload
```

7. Confirm the start and stop order for iSCSI target service.

Execute the following command on both nodes.

```
# systemctl show fjsvsdx.service | grep "After="
After=iscsi-shutdown.service iscsi.service ...
```



.....
Confirm that target.service is not output.
.....

The following procedure must not be performed on nodes where programs other than GDS use iSCSI target (LIO).

8. Set up to avoid the automatic startup of iSCSI target service when starting the system.

Execute the following command on both nodes.

```
# systemctl disable target.service
```


Appendix A General Notes

A.1 Restraining Access to Physical Special File

After registration of disks to a class other than the root class 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/sdXn          (for normal hard disk)
/dev/mapper/mpathXpn (for mpath devices of DM-MP)
/dev/emcpowerXn    (for emcpower devices)
/dev/vdXn          (for virtual disks on a KVM guest)
```

X indicates the device ID, and *n* is the slice number.

For example, if you execute the `dd` command to write to a disk using a physical special file, an error as below occurs.

```
# dd if=/dev/zero of=/dev/sda1
dd: /dev/sda1: open: Device busy
```

In the disk list output by the `parted(8)` command, disks that are registered to a class other than the root class are not displayed.

This access protection is canceled when a disk is removed from a class. It will also be canceled 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.

Access to physical special files is prevented on the following nodes:

- For disks that are registered in the shared class or the local class
 - 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
 - Access to the physical special files is prevented on the node where the local class is created.
- 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 in the root class
- 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

When mirroring among servers, access to the physical special files is prevented only for iSCSI devices. As the configuration that "[Figure 4.1 Disk setting for Performing Mirroring among Servers](#)" shows, access to the physical special file of the iSCSI devices (`sde` and `sdd`) will be prevented but access to the physical special file of the physical disk (`sdc`) will not be prevented. Do not access the physical special file of the physical disk.

A.2 Booting from a CD-ROM Device

If the system cannot be booted, for example, it may be required recovering the system booting from a CD-ROM device.

However, there are the following concerns when booting from a CD-ROM device.

- 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/sd[a-z]*[1-4]*`).

- 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 unless directed in this manual.

When mounting a file system by booting from a CD-ROM device in any procedure not described in this manual from necessity, it should be mounted as read-only.

A.3 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 resource registration and register shared disk units with the PRIMECLUSTER resource database.



For details on resource registration, see "[6.3 Shared Disk Resource Registration](#)."

- 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 [EFI] 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 resource registration in advance, and register all disks that belong to the local class with the resource database.
- In a cluster system with three or more nodes, if the physical scope of a shared disk does not match with the scope of a shared class, the shared disk cannot be registered to the shared class.
- 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.
- 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.4 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 "[6.6.1.6 Volume Configuration](#)" and "[B.1.4 sdxvolume - Volume Operations](#)."
- For details on the changing methods, see "[6.6.1.6 Volume Configuration](#)" and "[B.1.7 sdxattr - Change 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 "[7.1.1 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 "[B.1.6 sdxinfo - Display Object Configuration and Status Information](#)."



Under the following circumstances, normal resynchronization is performed even though the JRM mode of the volume is on:

- There is any slice in status other than ACTIVE or STOP, and a system panic occurs.

If the proxy is joined, this condition is also applied to the status of the proxy slice.

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 "[B.1.5 sdxslice - Slice Operations.](#)"
- or details on the changing methods, see "[B.1.7 sdxattr - Change 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 "[B.1.6 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 "[7.8.1.2 Operation Procedure](#)" and "[B.2.1 sdxproxy - Proxy Object Operations.](#)"
- For details on the changing methods, see "[B.1.7 sdxattr - Change 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 "[7.1.2 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 "[B.1.6 sdxinfo - Display Object Configuration and Status Information.](#)"



Note

JRM for Proxies of a Root Class [EFT]

In the root class, you cannot enable JRM for proxies. If you set "on" for JRM when the proxy volumes are detached from the master on the GDS Management View, JRM will be set "off." If the proxies are attached to the master again, the whole data of the master volume will be copied to the proxy volumes.

A.5 Backing Up and Restoring Object Configuration (sdxcfg)

- Do not use editors such as vim(1) and sed(1) to edit configuration tables created with the sdxcfg Backup command or those saved in configuration files. To edit configuration tables, use the sdxcfg Convert command.
- When the object of the class is removed using the sdxcfg Remove -e keepid command then it is restored using the sdxcfg Restore -e chkps command, restore the same configuration as the removed configuration.
- When the object configuration of the shared class registered as a Gds resource in the cluster application is restored using sdxcfg Restore command, execute the following command on any one node within the class scope.

```
# /opt/MAW/MAWRrms/bin/hvgdsetup -a class_name
...
Do you want to continue with these processes ? [yes/no] y
```

This command is set up in order to start up or stop the volume of shared class along with the change in the state of the cluster application.

- For Backup (sdxcfg Backup) and Restore (sdxcfg Restore) of the object configurations, the settings for I/O Response Time Assurance Function are not backed up or restored. If you are using the I/O Response Time Assurance Function, see "[7.15 I/O Response Time Assurance \[IOmonitor\]](#)" for the resolution.
- If you use the I/O Response Time Assurance Function and use the sdxcfg Remove command to delete a class that has been enabled to system panic mode, disable system panic mode before you execute the sdxcfg Remove command again. For details on how to change the system panic mode, see "[8.3.2 Changing the System Panic Mode.](#)"
- By using the sdxcfg command, you can use disks registered with the class and volumes in the class from other systems. However, when using disks that were registered in the class on a system using GDS 4.3A00 or earlier with a system that uses this version of GDS, backup and restoration for volume data and setup for disk labels may be required.

Below are the procedures. Usually, step 1, 5, and 7 are unnecessary. However, you have to perform these steps only when the above conditions are met.

1. Back up volume data in a class on a system before migration.
 2. Back up a class object configuration by using the sdxcfg Backup command on the system before migration.
 3. Remove the class object configuration by using the sdxcfg Remove command on the system before migration.
 4. Connect disk which was removed from the class in step 3 to the system after migration.
 5. Set up a MSDOS type disk label for the removed disk by using the parted(8) command on the system after migration. You can also skip this step and specify the -e mklable option in step 6.
 6. Restore the class object configuration by using the sdxcfg Restore command on the system after migration.
 7. Restore volume data in the class on the system after migration.
- When using multipath, register the native device to each node's Excluded Device List. For the Exclude Device List, see "[6.1 Creating Excluded Device List](#)"

A.6 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] of GDS Management View 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.

A.7 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, canceled, 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.8 Using Advanced Copy Function in Proxy Configuration

In a proxy configuration, by working in cooperation with the Advanced Copy functions, EC (including REC), or OPC (including QuickOPC) of ETERNUS Disk storage system, 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 master and the proxy belong to a root class. [EFI]
 - 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 "[7.10.1.2 Operation Procedure](#)."
 - Executing the `sdxproxy Restore` command. For details see "[B.2.1 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 "3.18 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 ETERNUS Disk storage system, 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.9 Shadow Volume



Note

Shadow volumes are not supported in this version.

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 might remain. 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 "[B.2.4 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(8) (excluding when using the -o ro option), fsck(8) or mkfs(8) 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.10 Changing system run level

For stopping and restarting the system and changing over the user mode, use the shutdown(8) command.

Operation	Command
Stopping the system	shutdown -h now
Restarting the system and changing over the user mode	shutdown -r now

The user mode after restarting the system by the shutdown -r now command is specified by the following procedure. For run levels other than the below, see the OS manual.

Note

Because the console cannot be used in a cloud environment (except an FJcloud-Baremetal environment), do not change the run levels.

As described in the table below, change the run level by using the systemctl command, and then restart the system.

User mode	Run level
Multi-user mode	multi-user.target
Single user mode	rescue.target

Example: When restarting the system in single user mode

```
# systemctl set-default rescue.target
# shutdown -r now
```

A.11 Root Class Operation [EFI]

For performing the root class operation with a command, you can only use the following commands:

- sdxinfo
- sdxattr -V -a mode={ro|rw}
- sdxattr -V -a jrm={on|off}
- sdxattr -C -a iotimeout={on [:time:]{on|off}] |off} [IOmonitor]

In a root class, you cannot use the following sub menus of the proxy operation menu on the GDS Management View:

- Relate
- Update
- Restore
- Swap

The following root class operation takes about 20 seconds.

- Removing volumes
- Joining proxy groups

A.12 Operating When There Is Disk in DISABLE Status or There Is Class not Displayed with sdxinfo Command

If there is a disk in DISABLE status or there is a class not displayed with the sdxinfo command, do not perform the creation, change, or deletion of any classes, groups, or volumes.

Perform creation, changes, or deletion only after performing recovery of the disk in DISABLE status or of the class not displayed with the sdxinfo command.

- For recovering the disk in DISABLE status
See the section "(1) Class becomes closed status during operation." in "D.1.4 Class Status Abnormality."
- For recovering classes not displayed with the sdxinfo command
See the section "(2) Class cannot be started when booting the system." in "D.1.4 Class Status Abnormality."

A.13 Notes on Mirroring among Servers

To use Mirroring among servers, check the following notes when designing the system.

Items	Reference
Design	"3.9 Mirroring among Servers"
Disk swap	"7.3.3 Swapping Disks of Netmirror Group"

Items	Reference
Operation and Maintenance	"7.16 Mirroring Among Servers"

Appendix B Command Reference

This appendix discusses the commands provided by GDS, GDS Snapshot, GDS I/O Monitor Option, and PRIMECLUSTER.

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 [EFI]
sdxparam	Configuration parameter operations
sdxconfig	Object configuration operations
sdxdevinfo	Display device information of internal disk

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 C GDS Messages](#)" check the meaning of the error message and take necessary action.

Note

GDS commands that are not described in this command reference must be used according to the procedures described in this manual or related manuals.

B.1 GDS Commands

B.1.1 sdxclass - Class Operations

SYNOPSIS

```
sdxclass -R -c class
```

DESCRIPTION

Use sdxclass 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.

Check that all nodes in the cluster domain are started before executing the command.

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, see "[B.1.2 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.

B.1.2 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.

Spare disks cannot be connected to groups. Keep disks and single disks cannot be connected to existing groups. Additionally, a keep disk and a single disk, multiple keep disks, or multiple single disks cannot be connected to one group together.

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, volumes are created for every physical slice within the keep disk. If the physical slices are not open, the created volumes are started and synchronization copying is performed automatically 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 groups.

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 the 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 scope 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 `sdxgroup -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.

Check that all nodes in the cluster domain are started before executing the command.

If no class with the name specified by *class* already exists, then one is automatically created.

A root type *class* can include *device* of the keep type. However, when registering multiple keep type devices together, as many or more undef type *devices* must be registered.

Do not execute the command if there is a closed class or a disk in SWAP status in the cluster domain.



Since the `sdxdisk` command initializes the registered physical disks (excluding *devices* with "keep" assigned as the type attributes), 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.

Check that all nodes in the cluster domain are started before executing the command.

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=netmirror, type=stripe, type=concat or type=switch (default is **mirror**)

Sets the type attribute of *group*. If *class* is the root class, specifying "netmirror," "stripe," or "concat" will result in an error. If *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."

netmirror

Sets type attribute to "netmirror."

stripe

Sets type attribute to "stripe."

concat

Sets type attribute to "concatenation."

switch

Sets type attribute to "switch."



"stripe" and "concat" cannot be specified in RHEL8.

width=*blks* (default is 32)

Sets the stripe width of *group*. The *blks* indicates the stripe width in block number (base 10). One block is 512 bytes. For *blks*, you can indicate an integer (from 1 to 1,073,741,824) 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 *group* 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 **shared**)

Sets the attributes of *class* type.

root [EFI]

Sets the type attribute to "root."

Objects within *class* of the root type can be used only on the current node.

In the root type *class*, the following disks can be registered: system disks which include root file systems, disks to mirror with the system disk, spare disks, and disks on which proxy volumes of the system volumes are created.

Only one root type class can be created within the system. If a class of the root type already exists, specifying another *class* of the root type will result in an error.

For the scope attribute, the node identifier of the current node is set automatically.

local

Sets the type attribute to "local."

Objects within *class* of the local type can be used only on the current node.

For the scope attribute, the node identifier of the current node is set automatically.

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 that uses class objects.

This option can be omitted for the root class (when the `-a type=root` option is specified) or the local class (when the `-a type=local` option is specified). If a node is specified other than `self-node`, an error occurs.

For a shared class (when the `-a type=shared` option is specified or the `-a type` option is omitted), specify a node group with 2 nodes or more and 4 nodes or less including `self-node`.

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.

iotimeout=on[:time[:panic]] or iotimeout=off (default is off) [IOMonitor]

It sets the operation of the I/O Response Time Assurance.

on

Enables the I/O Response Time Assurance.

off

Disables the I/O Response Time Assurance.

time

Indicates the I/O response time (sec).

The value is set to 77 by default.

Values from 6 to 400 may be specified for time.

If a value smaller than the default value is set, the I/O may terminate by I/O Response Time Assurance Function before the path switches when an error of a single path of the multi-path configuration occurs.

panic

It sets the system panic mode of I/O Response Time Assurance Function.

System panic mode is a mode that sets whether the system has panicked when an I/O error occurred in all of the slices that configures the volume and when the class is closed.

The value is set to on by default. Set the value to on for cluster system.

on

Enables the system panic mode.

off

Disables the system panic mode.

Note

Systems with both mirror volume and single volume

In the system with mirror volume, it is recommended to set the system panic mode to disable for all the class with single volume. When the system panic mode is set to enable for a class with single volume, the redundancy by mirroring has no effect, because the system will panic with just an occurrence of I/O error or I/O delay in only one storage unit.

-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. *device* which can be specified is up to 400.

The physical disk name can be specified in either the following formats:

sdX	(for normal hard disks)
mpathX	(for mpath devices of DM-MP)
emcpowerX	(for emcpower disks)
vdX	(for virtual disks on a KVM guest)

X indicates the device identifier.

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 [EFI]

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 with "keep" assigned as its *type*, the specified *device* must have a reserved physical slice number and sufficient free disk space so that the private slice can be created on the *device*. When multiple *devices* are specified, *devices* with "keep" assigned to *type* must have reserved physical slice numbers and sufficient free disk space or swap devices with sufficient disk space.

With a system disk (disk with running `/`, `/usr`, `/var`, `/boot`, or `/boot/efi` or a swap area) with "keep" assigned as its *type*, even if all or some of the physical slices are currently open, the `sdxdisk` command ends normally. However, if "keep" is specified for a disk other than a system disk, this command will result in an error where open physical slices exist. For example, if any of the physical slices are being used as a file system, unmount the file system to free up the physical slice, and then execute 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**) [EFI]

Specifies the created volume's attribute value when connecting *disk* of the keep type. This option setting is simply ignored if *disk* of the keep type is not specified.

Always use an equal sign (=) after *num*, and separate *volume* and *jrm* with a colon (:). When specifying the attributes values of multiple volumes, specifiers should be combined using commas (,) as the delimiter.

Specify the physical disk slice number (an integer 1 to 15) of a keep type *disk* storing volume data for *num*, the volume name for

volume, and the created volume's just resynchronization mode ("on" or "off") for *jrm*.

When the keep type *disk* contains multiple physical slices of which size is nonzero, it is necessary to specify the corresponding volume attribute values for all the physical slices.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.
Otherwise, a non-zero value is returned.

B.1.3 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 groups, 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 (from 1 to 1,073,741,824) that is two raised to the power, 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.

-I *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.

B.1.4 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

offline

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 14 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 4 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/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 disk label, 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 disk label, 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.

**Note**

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. *volume* which can be specified is up to 400.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.

Otherwise, a non-zero value is returned.

B.1.5 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.

Slice operations are impossible for the root class.

PRIMARY OPTIONS

You can use one of the following options.

-F

oFfline

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.

Note

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/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.

Note

Be aware that as long as a slice is not attached using the `-R` option, the degree of mirror multiplexing stays reduced.

-N

oNline

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 effective only for a 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 local or shared 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 specified 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* which can be specified is up to 400.

Slice name should be specified 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.

B.1.6 **sdxinfo** - Display Object Configuration and Status Information

SYNOPSIS

```
sdxinfo [-ACDGSV] [-c class] [-o object] [-e label,long]  
sdxinfo -I [-c class] [IOmonitor]
```

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. Time required for display varies depending on the configuration of the objects.

The objects are displayed at random order. Even when the objects are created in the same order, the display order may vary.

Note

- Do not execute the `sdxinfo` command multiple times and continuously. Executing it multiple times and continuously may make execution take more time.
- When the `sdxinfo` command is executed for the first time after starting the system, among the by-id files (symbolic link files created under the `/dev/disk/by-id` directory) managed by the `udev` function of Linux, the by-id files of disks which are not registered in classes may not exist temporarily because OS creates the by-id files again triggered by the GDS processes for obtaining the device information. When access to the by-id files of disks which are not registered in classes fails, retry access. Alternatively, when the `sdxinfo` command is executed for the first time after starting the system, do not access to the by-id files of disks which are not registered in classes.

PRIMARY OPTIONS

Primary options indicate the type of objects or the information of I/O Response Time Assurance Function to be displayed. If nothing is specified, only information on the pertinent object will be displayed.

Object types can be any combination of the following (except -I).

-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.

-I [IOmonitor]

IOmonitor

Displays the information of I/O Response Time Assurance Function.

When the shadow class is set with `-c` option, information will not be displayed.

-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 label

Add the disk label type and output it to Class information.

-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.



Note

For viewing the COPY status of slices in proxy volumes, do not specify this option.

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 [EFI]

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.

LABEL

Displays one of the following values to indicate the disk label type of a disk registered with a class when the `-e label` option is specified. For a root class, an asterisk (*) is displayed.

gpt

GPT type

msdos

MSDOS type (MBR type)

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.

netmirror

Netmirror. It is connected to a netmirror 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 [EFI]

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.

```
sdX           (for normal hard disks)
mpathX       (for mpath devices of DM-MP)
emcpowerX   (for emcpower disks)
vdX         (for virtual disks on a KVM guest)
```

X indicates the device identifier

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.

Note

When an I/O error has occurred in the root class, the E field of the disk information will not display "1" which indicates the I/O error status. Check the I/O error information using system log etc.

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, if it is 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.

netmirror

Netmirror 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.

netmirror

Netmirror. It belongs to a netmirror 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

.....
This mode is normally off. If Symfoware Server(Native) is used, Symfoware Server(Native) changes the mode. Note that the open 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 the disk label 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 disk label.

For a shadow volume it always displays off no matter whether the shadow slice is registered with the disk label.

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

Proxy JRM is "on."

off

Proxy JRM is "off."

Slice information:

OBJ

Displays slice as an object classification.

NAME

When `-e` long option is used, this 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, 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, 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 an asterisk (*) is displayed.

Note

When using GDS Snapshot

- For differential copy processing with EC
Displays the total number of blocks that have been equivalent status at the start time of copy processing and the blocks that have already been copied.
 - For copy processing with QuickOPC
When the type of ETERNUS Disk storage system and the firmware version are older than the following; ETERNUS DX90 S2, DX410 S2, DX440 S2, DX8100 S2, DX8700 S2 (Firmware version V10L30), displays the number of blocks that are equivalent to a copy source volume. If the copy source volume or copy destination volume are updated during the copy processing, CURBLKS at the completion time of copy processing is smaller than COPYBLKS for the updated blocks.
 - For copy processing with TimeFinder or SRDF
Displays an asterisk (*).
-

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 sfdsk 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.

QOPC

Copying is in process using the QuickOPC function.

REC

Copying is in process using the Remote Equivalent Copy function.

TF

Copying is in process with Dell EMC TimeFinder.

SRDF

Copying is in process with Dell 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.

QOPC

Source or target of the QuickOPC session.

REC

Source or target of the Remote Equivalent Copy session.

TF

Source or target of Dell EMC TimeFinder's BCV pair.

SRDF

Source or target of Dell EMC SRDF's SRDF pair.

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.

execute

QuickOPC session is on the copy processing and the "tracking" state.

suspend

EC or REC session suspended.

split

BCV pair or SRDF pair split.

track

QuickOPC session is the "tracking" state.

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.

Information of I/O Response Time Assurance Function [IOmonitor]

CLASS

Displays the class name.

MODE

Displays one of the following as an operation of I/O Response Time Assurance.

on

Enables I/O Response Time Assurance.

off

Disables I/O Response Time Assurance.

TIME

Displays the I/O response time (seconds).

Displays the asterisk (*), when the I/O Response Time Assurance is disabled.

PANIC

Displays one of the following as a system panic mode.

Displays the asterisk (*), when I/O Response Time Assurance is disabled.

on

Enables system panic mode.

off

Disables system panic mode.

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
```

Displays the information of I/O Response Time Assurance for all of the class within the local node.

```
# sdxinfo -I
CLASS  MODE TIME PANIC
-----
class0001 off  *   *
class0002 on   77  on
```

Displays the information of I/O Response Time Assurance of the class called "class0002."

```
# sdxinfo -I -c class0002
CLASS  MODE TIME PANIC
-----
class0002 on   77  on
```

CAUTION

Additional information may be displayed in accordance with new functions provided.

B.1.7 sdxattr - Change 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**

Changes the class type attribute.

When *class* includes active volumes, it can be changed from "local" to "shared" but cannot be changed from "shared" to "local." To change the class type from "shared" to "local", stop all the volumes and then execute this command.

Changes from "root" to "local" or "shared" and similarly from "local" or "shared" to "root" are impossible. Additionally, the class types cannot be changed with shared classes that include disks of which physical scope is one node or that include switch groups.

local

Change the type attribute to "local."

Objects within local type classes can be used only on the current node.

For the scope attribute, the node identifier of the current node is set automatically.

shared

Changes the type attribute to "shared."

By specifying this along with the scope attribute, objects within the class can be shared with multiple nodes including the current node within the class.

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.

The scope can specify a node group with 2 nodes or more and 4 nodes or less including self-node.

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.

iotimeout=on[:time[:panic]] or iotimeout=off [IOmonitor]

Change the operation of the I/O Response Time Assurance.

The changes become active right away. It is not necessary to restart the system to activate the changes.

on

Enables the I/O Response Time Assurance.

An error occurs if a concat, stripe or netmirror type group exists in *class*.

off

Disables the I/O Response Time Assurance.

time

Set the I/O response time (seconds).

The default value is 77.

The values possible to set for time is from 6 to 400.

If setting a value lower than the default value, then when an irregularity occurs on one path of the multi-path disk, there is a chance that the I/O Response Time Assurance Function aborts the I/O before the path is switched over.

panic

It sets the system panic mode of I/O Response Time Assurance Function.

System panic mode is a mode that sets whether the system has panicked when an I/O error occurred in all of the slices that configures the volume and when the class is closed.

The value is set to on by default. Set the value to on for cluster system.

on

Enables the system panic mode.

off

Disables the system panic mode.



Systems with both mirror volume and single volume

In the system with mirror volume, it is recommended to set the system panic mode to disable for all the class with single volume. When the system panic mode is set to enable for a class with single volume, the redundancy by mirroring has no effect, because the system will panic with just an occurrence of I/O error or I/O delay in only one storage unit.

-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 [EFI]

Sets the type attribute to "keep."

This disk will then be handled as a keep disk, and its format and data will be retained when it is connected to a group. Single disks cannot be changed to keep disks.

single

Sets the type attribute to "single."

Single volume may be created within *disk* thereafter. Keep disks cannot be changed to single disks.

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=volume_name

Sets the volume name to *volume_name*.

When there is an activated *volume*, it results in an error. First, you must stop the *volume*.

When changing a volume 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/fstab*.

/dev/sfdsk/classname/dsk/volume_name

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 disk label.

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 (four) 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.

B.1.8 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 the sdxswap command returns the control 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.



When setting the I/O Response Time Assurance Function, an error may happen when there is not enough memory.

In this case, the following message is displayed to the standard error output.

```
ERROR: sfdsk driver returned an error, errno=12
```

Also, a warning message 22018 (WARNING: sfdsk: failed to open disk, no enough address space: ...) is output to the console.

Perform any of the following action before executing this command.

- (a) Restart the system.
- (b) Wait until the memory is freed up.
- (c) Increase the memory.

Refer to the installation guide for the necessary memory.

Plan the memory increase if selecting (a) or (b).

-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.

B.1.9 sdxfix - Restore Failed Object

SYNOPSIS

```
sdxfix -C -c class
```

```
sdxfix -D -c class -d disk [-e online] [-x Devlabel, NoRdchk]
```

```
sdxfix -V -c class {-g group|-d disk} -v volume [-e force] [-x NoRdchk]
```

DESCRIPTION

Use `sdxfix` to restore failed objects (excluding shadow objects). Data on restored disks or volume objects may no longer have consistency, and after the restoration, it is necessary to restore consistency using backup data or checking with 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. Response from the command requires some time depending on the *disk* size because the entire disk area is read.

For restoring the object status without reading the *disk* area, specify the sub option -x NoRdchk.

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 a 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 force (when using **-V**)

When *volume* is a netmirror volume, even if the stopped node exists, restoration is performed. You cannot use this option by specifying volumes other than a netmirror volume.

Do not use this option other than the procedure described in this manual.

-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.

-x Devlabel (when using **-D**)

Restores the device information of the disk used for the mirroring among servers that is specified by *disk*.

Do not use this option other than the procedure described in this manual.

-x NoRdchk (when using **-D**)

Does not perform a read check in *disk* area.

When a read check is not necessary, for example, if a disk enters the I/O error state because the path has been removed, you can shorten the recovery process time.

Use this option only when a read check is clearly unnecessary.

-x NoRdchk (when using **-V**)

Does not perform a read check of the slice specified by the combination of *disk* and *volume*, or the combination of *group* and *volume*.

Do not use this option other than the procedure described in this manual.

RETURNED VALUE

When it is normally terminated, "0" is returned.
Otherwise, a non-zero value is returned.

B.1.10 `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. *volume* which can be specified is up to 400.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.

Otherwise, a non-zero value is returned.

B.1.11 sdxroot - Root File System Mirroring Definition and Cancellation [EFI]

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 Linux operating system 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, /boot, /boot/efi, or swap`



Do not use this command on this version. Setting and cancellation of system disk mirroring is performed using the GDS Management View. For details, see "[6.5 System Disk Mirroring Settings \[EFI\]](#)" and "[9.4 System Disk Mirroring Cancellation \[EFI\]](#)."

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*. After returning from this command, reboot the system immediately. After the system is rebooted, system disk mirroring will commence.

*) Update system files (`fstab`, `grub.cfg`, and `dracut.conf`).

Specify a disk that includes a slice currently running as `/ (root), /usr, /var, /boot, /boot/efi`, or a swap area for *disk*. Among `/ (root), /usr, /var, /boot, /boot/efi`, and swap areas, although it is not required to specify disks with swap areas only, be sure to specify disks with `/ (root), /usr, /var, /boot, and /boot/efi`. 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*. After returning from this command, reboot the system immediately. After the system is rebooted, system disk mirroring will be canceled.

*) Update system files (`fstab`, `grub.cfg`, and `dracut.conf`).

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, /var, /boot, /boot/efi`, 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, /var, /boot, and /boot/efi 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



Information Collection and Environment Configuration Before and After Setting the System Disk

Information collection and environment configuration are required before and after setting the system disk.

For details, see "[6.5 System Disk Mirroring Settings \[EFI\]](#)."

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, /var, /boot, and /boot/efi is "sda", and the disk allocated as a swap area is "sdb."

```
# sdxdisk -M -c System -a type=root -d sda=Root1:keep,  
sdc=Root2:undef,sdb=Swap1:keep,sdd=Swap2:undef
```

3. Connect the system disks to a group.

```
# sdxdisk -C -c System -g Group1 -d Root1,Root2 -v 1=root:on,  
2=usr:on,3=var:on,4=home:on,5=boot:on,6=efi:on  
# sdxdisk -C -c System -g Group2 -d Swap1,Swap2 -v 1=swap:on
```



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 not displayed with the `mount(8)` command may not be open.

4. Check that mirroring definition is completed.

```
# sdxroot -M -c System -d Root1,Swap1
```

5. Reboot the system.

```
# shutdown -r now
```

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 -r now
```

5. Check that the mirroring has been canceled.

Using the mount command and the sdxinfo command, make sure that the system disk mirroring has been canceled 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 boot
# sdxvolume -F -c System -v efi
# 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 boot
# sdxvolume -R -c System -v efi
# 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
```

B.1.12 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 one 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. Execute on all the nodes for cluster system.

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.

B.1.13 sdxconfig - Object Configuration Operations

SYNOPSIS

```
sdxconfig Backup -c class[-o outfile] [-e update]
```

```
sdxconfig Convert -e remove[,update] -c class -d disk,...
[-i infile] [-o outfile]
```

```
sdxconfig Convert -e remove[,update] -c class -g group,...
[-i infile] [-o outfile]
```

```
sdxconfig Convert -e rename[,update] -c class=classname
[-i infile] [-o outfile]
```

```
sdxconfig Convert -e replace[,update] -c class -d disk=device[,disk=device,...]
[-i infile] [-o outfile]
```

```
sdxconfig Convert -e replace[,update] -c class
-p device=newdevice[,device=newdevice,...]
[-i infile] [-o outfile]
```

```
sdxconfig Remove -c class [-e keepid]
```

```
sdxconfig Restore -c class -i infile [-e chkps|mklablel,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.

Do not execute the command if there is a closed class or a disk in SWAP status in the cluster domain.

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. Also, to restore the shared class registered in the cluster application, after this command execution it is necessary to reboot the node where this command

was executed.

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.

For a cluster system, it is necessary to register the physical disk contained in the configuration table with the resource database of the cluster system before executing this option.

Note

- The device number (the minor number) of the restored volume, and the owner and the access permission of a device special file of the volume cannot be restored to the same value as the time of executing the Backup option. It will be the same value as the case where a new volume is created with the `sdxvolume -M` command. The device number, the owner and the access permission can be checked with the following command:

```
# ls -l /dev/sfdisk/class_name/disk/volume_name
```

If the device number at the time of executing the Backup option is set to the application which uses the restored volume, you need to modify the application configuration.

For restoration of the ownership and access permission of the device special file, change the settings by using the `chown(1)` command or the `chmod(1)` command.

- The settings of the I/O Response Time Assurance Function are not restored. If setting the I/O Response Time Assurance Function, set it by executing the `sdxattr` command after executing the Restore option.
- Configurations that include concatenation groups or stripe groups cannot be restored in RHEL8. Remove concatenation groups and stripe groups from the configuration table with the `sdxconfig Convert` command.

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.

```
sdX      (for normal hard disks)
mpathX   (for mpath devices of DM-MP)
emcpowerX (for emcpower disks)
vdX      (for virtual disks on a KVM guest)
```

X indicates the device identifier.

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.

If a class name *class* is stored in the private slices of physical disks which are not contained in the configuration table, do not specify this option.

This option cannot be used along with the `-e mklable` option.



Information

The following are examples of when the `-e chkps` option is required or not required when executing the `sdxcfg Restore` command.

- The `-e chkps` option is required
 - When specifying the `-e keepid` option in the `sdxcfg Remove` command
 - When the data is copied per LUN from the migration source storage to the migration destination storage
- The `-e chkps` option is not required
 - When not specifying the `-e keepid` option in the `sdxcfg Remove` command
 - When executing the `sdxcfg Restore` command for a disk where only volume areas on the physical disks are copied

-e keepid (when using **Remove**)

Deletes object configuration information of *class* while retaining 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 `sdxcfg Restore` command.

Note

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, restore the object configuration using the Restore option once, and then execute the Remove option again without this option.

-e mklabel (when using **Restore**)

If the following two disk label types; a disk label type of a physical disk which will be registered with *class*, and a disk label type of a class which described in a configuration table of specified configuration file by *infile*, are different, change the physical disk label type to the disk label type of the class described in the configuration table.

When this option is used, it is not ensured for data in volume area on the physical disk registered with the *class*.

This option cannot be used along with the `-e chkps` 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.

sd <i>X</i>	(for normal hard disks)
mpath <i>X</i>	(for mpath devices of DM-MP)
emcpower <i>X</i>	(for emcpower disks)
vd <i>X</i>	(for virtual disks on a KVM guest)

X indicates the device identifier.

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.

B.1.14 sdxdevinfo - Display Device Information of Internal Disk

SYNOPSIS

```
sdxdevinfo -c class -d disk,...
```

DESCRIPTION

When **sdxdevinfo** is executed, it displays a physical disk name and by-id name at the time the internal disk *disk* was registered with the root class or local class *class*.

The **sdxdevinfo** command can be executed with superuser access privileges only.

PRIMARY OPTIONS

None

SUB OPTIONS

Sub options are as follows:

-c class

For *class*, specify a class name of root class or local class whose information will be displayed.
Do not specify a shared class.

-d disk,...

For *disk*, specify the disk name of the internal disk whose information will be displayed.
Use comma-delimited format to specify more than one disk.
Do not specify a shared disk. When a shared disk is specified, the correct information will not be displayed.

DISPLAYED INFORMATION

sdxddevinfo command displays the following information.

class

Displays a class name.

disk

Displays a disk name.

device

Displays a physical disk name when *disk* was registered with *class*.

by-id

Displays a by-id name when *disk* was registered with *class*.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.
Otherwise, a non-zero value is returned.

USAGE EXAMPLES

Displays the physical disk name and by-id name of the internal disk rootDisk0001 registered in the root class RootClass when it was registered in RootClass.

```
# /etc/opt/FJSVsdx/bin/sdxdevinfo -c RootClass -d rootDisk0001
class    disk        device  by-id
-----  -
RootClass rootDisk0001 sda      3500000e111c01810
```

B.2 GDS Snapshot Commands

B.2.1 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,OPC,QOPC]

sdxproxy Root -c class -m master,... [-e boot] [EFI]

sdxproxy Root -c class -p proxy,... [-e boot] [EFI]

sdxproxy Root -c class -m master,... -p proxy,... [-e boot] [EFI]

sdxproxy Swap -c class -p proxy

sdxproxy Update -c class -p proxy,... [-e instant,nowaitsync,OPC,QOPC]

```

DESCRIPTION

Use sdxproxy to perform operations on proxy objects.
The sdxproxy command can be executed with superuser access privileges only.



See

.....

- [2.3.12 Proxy Configuration Preconditions](#)
- [2.3.13 Number of Proxy Volumes](#)
- [2.3.14 Proxy Volume Size](#)
- [2.3.15 Proxy Group Size](#)
- [3.18 Instant Snapshot by OPC](#)
- [3.19 Using Dell EMC TimeFinder or Dell EMC SRDF in Proxy Configuration](#)
- [A.8 Using Advanced Copy Function in 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 canceled. A proxy volume within a proxy group cannot be specified for cancellation.

Even after the relationship is canceled, 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 canceled. For example, if a master volume is being used as a file system, you must use the `umount(8)` command to unmount 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 canceled 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 attributes of proxy volumes created in the proxy group are 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 specified, 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/sfdsk/class/dsk/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 `umount(8)` command to unmount 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. 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

[EFI]

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 volumes.

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 specified, 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 or QuickOPC function, the proxy data at the moment of starting the copy process 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 or QuickOPC function is unavailable, the command fails.

When neither `-e OPC` nor `-e QOPC` options are specified, the copy function to use depends on whether a QuickOPC session is present or not. Use OPC function if QuickOPC session is not present, and use QuickOPC function if QuickOPC session is present.

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, unmount the file system with the `umount(8)` 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.
However, when all the following conditions are met, this command will not result in an error:
 - Copying is in process from a specified proxy volume with QuickOPC function.
 - `-e OPC` is not specified.
- A slice being copied exists in 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 [EFI]

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/fstab` 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 copy destinations.
- 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 system files* 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 `fstab` on the alternative root volume without change. After parting the current root volume and the alternative root volume using the `Part` option, to edit `fstab` on the current root volume, to edit `fstab` 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 `fstab` 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 system files (`fstab`, `grub.cfg`, and `dracut.conf`) 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 selecting the boot device name for the alternative boot environment in the EFI boot manager's boot option selection screen, 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, an EC session, a BCV pair, or an SRDF pair exists.

Update

Copies data from a master to a parted proxy and updates contents of the proxy. With the OPC or QuickOPC function, the master data at the moment of starting the copy process 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). If the OPC or QuickOPC function is unavailable, the command fails.

When neither `-e OPC` nor `-e QOPC` options are specified, the copy function to use depends on whether a QuickOPC session is present or not. Use OPC function if QuickOPC session is not present, and use QuickOPC function if QuickOPC session is present.

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, unmount the file system with the `umount(8)` 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.
However, when all the following conditions are met, this command will not result in an error:
 - Copying is in process from a specified proxy volume with QuickOPC function.
 - `-e OPC` is not specified.
- A slice being copied exists in the proxy volume.
However, when all the following conditions are met, this command will not result in an error:
 - Copying is in process from a master volume with QuickOPC function.
 - `-e OPC` is not specified.
- 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**) [EFI]

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.

When the `sdxproxy` command ends normally, the previous and new boot device names are output to standard output. Be sure to take a note of the output boot device names. By selecting a previous boot device name in the EFI boot manager's boot option selection screen, it is possible to boot in the original boot environment. Additionally, by moving a previous boot device name to the top with the Change Boot Order menu on the EFI boot manager's maintenance menu, 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 the copy function of a disk unit, 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 canceled and the status of the proxy volume will become INVALID. When Dell EMC SRDF is used for the copying process from master to proxy, the relationship between master and proxy cannot be canceled.

-e instant (when using **Part**)

Parts proxy volumes and creates virtual snapshots of master volumes with the OPC or QuickOPC 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. If copying is in progress from proxy to master, or if the OPC function is unavailable, 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 or QuickOPC 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 OPC (when using Restore, Update)

OPC function is used for the copying process between master volume and proxy volume.
This option cannot be specified at the same time with -e QOPC option.

-e QOPC (when using Restore, Update)

QuickOPC function is used for the copying process between master volume and proxy volume.
When this option is specified, a QuickOPC session will remain between master volume and proxy volume after completing the copying process.
This option cannot be specified at the same time with -e OPC option.

-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 to be handled. When specifying multiple volumes or groups, they must belong to the same class.

Specify the master volume name or the master group name for *master*.

When specifying multiple volumes or groups, these specifiers should be combined using commas (,) as the delimiter.

-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. *proxy* which can be specified is up to 400.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.

Otherwise, a non-zero value is returned.

B.2.2 sdxshadowdisk - Shadow Disk Operations



.....
This command is not supported in this version.
.....

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 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 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*. The *blks* indicates the stripe width in block number (base 10). One block is 512 bytes. For *blks*, you can indicate an integer (from 1 to 1,073,741,824) that is two raised to the power, which is equal to or smaller than the available size of the smallest shadow disk specified by *disk*,... If *group* is not a stripe type, this option will result 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 (,). *device* which can be specified is up to 400.

The physical disk name can be specified in either the following formats:

```
sdX          (for normal hard disks)
mpathX      (for mpath devices of DM-MP)
emcpowerX   (for emcpower disks)
vdX         (for virtual disks on a KVM guest)
```

X indicates the device identifier.

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.

B.2.3 sdxshadowgroup - Shadow Group Operations



.....
This command is not supported in this version.
.....

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 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 specified by *hgroup*. Shadow groups specified 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 specified by *hgroup* already exists, stripe columns will be added after the existing stripe columns in *hgroup*, in the order they are specified 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*. The *blks* indicates the stripe width in block number (base 10). One block is 512 bytes. For *blks*, you can indicate an integer (from 1 to 1,073,741,824) that is two raised to the power, which is equal to or smaller than the available size of the smallest shadow group specified by *lgroup*,... If *hgroup* is not a stripe type, this option will result in 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.

B.2.4 sdxshadowvolume - Shadow Volume Operations

**Note**

.....
This command is not supported in this version.
.....

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 to (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 the disk label. 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/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 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 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 (.). *volume* which can be specified is up to 400.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.
Otherwise, a non-zero value is returned.

B.3 Command Reference

This section explains the following commands.

- **clautoconfig(8)**
- **cldelrsc(8)**

B.3.1 clautoconfig(8)-Execute resource registration

SYNOPSIS

```
/etc/opt/FJSVcluster/bin/clautoconfig -f diskfile  
/etc/opt/FJSVcluster/bin/clautoconfig -r -n
```

DESCRIPTION

This command enters information of the following hardware connected to a cluster system in the resource database.

- Shared disk devices

Information of shared disk devices specified in the shared disk definition file is entered in the resource database.

- Network devices

Network devices are automatically detected and entered in the resource database.

You can execute this command on any single node in the cluster system with the Administrator's authority.

While this command is running, you may not execute this command on the given node or another node.

OPTIONS

-f *diskfile*

Registers the shared disk unit in the resource database. Specify the full path to a predefined shared disk definition file for *diskfile*.

-n

Registers the network device in the resource database. The detected network device is active. The inactive devices are not detected. Execute the `ifconfig(8)` command to confirm if the network device is running.

-r

Registers the resource automatically.

EXIT STATUS

0:

Normal termination.

Other than 0:

Abnormal termination.

NOTES

You need to check the following before executing the command:

- CF, CIP, and resource database settings are complete on all nodes
- All nodes and disk units are connected.
- All nodes are up in multi-user mode.
- There is no closed class in the cluster domain.
- There is no disk in SWAP status in the cluster domain.

B.3.2 `cldelrsc(8)`-Delete resource

SYNOPSIS

`/etc/opt/FJSVcluster/bin/cldelrsc -r rid`

DESCRIPTION

The `"/etc/opt/FJSVcluster/bin/cldelrsc -r rid"` command deletes resources from the resource database. Resources of the following resource classes can be deleted.

Resource class name	Meaning
DISK	Resource of the shared disk unit
SHD_DISK	Shared resource of the shared disk unit
Ethernet	Resource of the network device

Use the `clgettree(1)` command to check the resources to be deleted.

Among shared disk unit resources, if resources in the DISK class are deleted and shared resources in the SHD_DISK class remain, it is also necessary to delete the shared resources.

Execute this command on any single node in the cluster system using system administrator access privileges.

OPTIONS

-r *rid*

Specify *rid* for the resource ID of the resource to be deleted. Check the resource ID executing the `clgettree(1)` command.

EXIT STATUS

0:

Normal termination.

Other than 0:

Abnormal termination.

NOTES

If a shared disk of which resources are to be deleted is registered with a GDS class, delete the shared disk from the GDS class, and then delete the resources.

EXAMPLES

The following example shows the procedure for deleting resources of a shared disk unit `"/dev/sdc"` shared between node1 and node2.

Perform the following procedure on any node.

1. Before deleting the resource, change the device name set in the resource to the current device name. Execute the following command.

For *filepath*, specify an empty file with an absolute path.

```
# /etc/opt/FJSVcluster/bin/clautoconfig -f filepath
```

Note

If `SDX_UDEV_USE=off` is described in the GDS configuration file `/etc/opt/FJSVsd/sdx.cf`, you must not execute the above command.

2. Check the resource IDs of `/dev/sdc` resources using the `clgettree(1)` command. In the following example, the resource ID of shared disks is 17, the `/dev/sdc` resource ID on node1 is 18, and the `/dev/sdc` resource ID on node2 is 19.

```
# /etc/opt/FJSVcluster/bin/clgettree
Cluster 1 cluster
  Domain 2 PRIME
    Shared 7 SHD_PRIME
      SHD_DISK 14 SHD_Disk14 UNKNOWN
        DISK 15 sdb UNKNOWN node1
        DISK 16 sdb UNKNOWN node2
      SHD_DISK 17 SHD_Disk17 UNKNOWN
        DISK 18 sdc UNKNOWN node1
        DISK 19 sdc UNKNOWN node2
    Node 3 node1 ON
      DISK 15 sdb UNKNOWN
```

```
DISK 18 sdc UNKNOWN
Node 5 node2 ON
DISK 16 sdb UNKNOWN
DISK 19 sdc UNKNOWN
```

3. Delete the resources of the shared disk unit "/dev/sdc" using the `cldelrsc(8)` command.

```
# /etc/opt/FJSVcluster/bin/cldelrsc -r 17
# /etc/opt/FJSVcluster/bin/cldelrsc -r 18
# /etc/opt/FJSVcluster/bin/cldelrsc -r 19
```

Appendix C GDS Messages

This section explains the messages for GDS drivers, daemons and commands.



`/var/log/messages` file

When `$EscapeControlCharactersOnReceive` off has not been specified in `/etc/rsyslog.conf`, control characters included in the message of `/var/log/messages` may be displayed in the form of "#XXX(X indicates digits)." However, no operational failure will occur and no action is required.

Example:

```
SDX: sdxservd: INFO: command executed:#012#011sdxvolume -M -c Class1 -g Group1 -s 32768 -v Volume1
SDX: sdxservd: INFO: command exited, exit-status=0:#012#011sdxvolume -M -c Class1 -g Group1 -s 32768
-v Volume1
```

C.1 Web-Based Admin View Messages (0001-0099)

For message numbers 0001 to 0099, see "PRIMECLUSTER Web-Based Admin View Operation Guide" for details.

C.2 Driver Messages

Driver messages are output to a log file or console via the `rsyslogd` interface.

Message Numbers

The message numbers used to identify messages described in E.2.1 to E.2.3 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>driver</i>	Driver name
<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>pdev</i>	Physical disk device name. e.g.) sda
<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)

Variable names	Descriptions
<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 four 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
ERROR	This message is output when an event that will stop some services is detected.	kern.err
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 facility and priority of a message passed from the GDS driver to rsyslogd(8).

The output destinations of messages are defined in the `/etc/rsyslog.conf` configuration file of rsyslogd(8) and can be changed by modifying the definitions in `/etc/rsyslog.conf`. For details, see the manual page of rsyslog.conf(5).

C.2.1 Panic Messages (20000-20099)

20002

```
sfdsk: I/O error on last ACTIVE slice
```

Explanation

Caused a panic after an I/O error occurred on the only active slice in the volume.

Resolution

Disk failure may have occurred. Check if the physical disk is operating normally.

20003

```
sfdsk: class closed down
```

Explanation

Caused a panic after the class was closed.

Resolution

The following are the possibilities:

- Disk failure.
- An abnormality occurred in the communication between nodes in the cluster system.

Check if the physical disk and the cluster system are operating normally.

20004

`sfdsksys: I/O error on last ACTIVE slice`

Explanation

Caused a panic after an I/O error occurred on the only active slice in the volume.

Resolution

Disk failure may have occurred. Check if the physical disk is operating normally.

20005

`sfdsksys: class closed down`

Explanation

Caused a panic after the class was closed.

Resolution

Disk failure may have occurred. Check if the physical disk is operating normally.

C.2.2 Error Messages (21000-21099)

21000

`ERROR: sfdsksys: configuration database corrupted, details`

Explanation

Invalid information is detected in the configuration database of root class. If this message is output, the system may not be started.

Resolution

Restore it according to "[\(4\) System cannot be booted. \(Failure of all boot disks\)](#)" or "[D.1.5 System Disk Abnormality \[EFI\]](#)." Even if the system can be started, this resolution needs to be performed.

21001

`ERROR: driver: this kernel version is not supported`

Explanation

The GDS driver *driver* could not be loaded because it does not support this kernel version.

Resolution

If the system cannot be started, collect the console logs and contact field engineers.

If the system can be started, take the following actions:

Check that GDS patch supporting the kernel version is applied. Apply it if not.

If the patch is applied or doesn't exist, collect investigation material and contact field engineers.

C.2.3 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=pdev

    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=pdev

    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=pdev

    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=pdev
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=pdev
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=pdev
```

```
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=pdev
           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=pdev
           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=pdev
           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=pdev
```

```
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=pdev
```

```
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=pdev
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 field engineers.

22012

```
WARNING: sfdsk: hook for device is not GDS entry.
```

Explanation

Inconsistency was found in management information of disks registered with GDS, but it does not affect the operation.

Resolution

No resolution is required.

22013

```
WARNING: access protection of physical special files disabled, major=major
```

Explanation

Accessing physical special files of the device for the major number *major* cannot be restrained due to unsuccessful memory allocation.

Resolution

Be careful not to access physical special files directly.

To restrain access physical special files, check whether you have sufficient memory or swap area, and then reboot the system after performing the procedure if necessary.

22014

```
WARNING: sdxfunc: no enough address space, ofunc=ofunc, errno=errno
```

Explanation

Memory allocation failed.

Resolution

Check whether you have sufficient memory or swap area. If not, expand it and reboot the system.

22015

```
WARNING: sfdsksys: cannot get device identifier, details
```

Explanation

Could not get the disk device identifier. Thus, operations such as setting or canceling the system disk mirroring, or swapping the disk cannot be performed.

Resolution

Collect investigation material and contact field engineers.

22016

```
WARNING: sfdsksys: illegal disk specific information
```

Explanation

The disk specific information of the root class that GDS registered to NVRAM (non-volatile memory) of the server itself is invalid.

If this message is output, the system may not be started.

This message is output when the system is started, if either operation below was performed in the environment where the system disk mirroring is performed on a VMware guest by using virtual disk (VMDK) or RDM in virtual compatible mode.

- The 4.3A40 or earlier environment that the patch for PH06344 was not applied was upgraded to the 4.4A00 or later environment.
- The 4.3A40 or earlier environment that the patch for PH06344 was applied without canceling the system disk mirroring was upgraded to the 4.4A00 or later environment.

Resolution

Cancel the system disk mirroring once, and set it again.

By this operation, the appropriate disk specific information is registered to NVRAM.

22017

```
WARNING: sfdsk: failed to create volume, no enough address space:
```

```
volume info: devno(maj,min)=v_devno(v_maj,v_min)
```

Explanation

Because the memory allocation used in the settings for the I/O Response Time Assurance Function failed, the creation of the volume shown in volume info failed.

Resolution

Check if the memory is insufficient, see "MemFree" of /proc/meminfo. If the memory is insufficient it is necessary to change the system configuration and allocate more memory, or to cancel the settings for the I/O Response Time Assurance Function.

Refer to the installation guide for the memory necessary for the I/O Response Time Assurance Function.

22018

```
WARNING: sfdsk: failed to open disk, no enough address space:
```

```
device info: devno(maj,min)=p_devno(p_maj,p_min)
```

```
devname=pdev
```

Explanation

Because the memory allocation used in the settings for the I/O Response Time Assurance Function failed, the open request of disk shown in device info failed.

Resolution

Check if the memory is insufficient, see "MemFree" of /proc/meminfo. If the memory is insufficient it is necessary to change the system configuration and allocate more memory, or to cancel the settings for the I/O Response Time Assurance Function.
Refer to the installation guide for the memory necessary for the I/O Response Time Assurance Function.

22019

WARNING: sfdsk: failed to change attribute value of iotimeout in class, no enough address space

Explanation

Because the memory allocation used in the settings for the I/O Response Time Assurance Function failed, the command to change class attribute failed.

Resolution

Check if the memory is insufficient, see "MemFree" of /proc/meminfo. If the memory is insufficient it is necessary to change the system configuration and allocate more memory, or to cancel the settings for the I/O Response Time Assurance Function.
Refer to the installation guide for the memory necessary for the I/O Response Time Assurance Function.

22020

WARNING: sfdsk: startup failed, failed to create a new thread

Explanation

sfdsk driver ended because the creation of thread has failed and continuing of the process is impossible.

Resolution

Check if the memory is insufficient, see "MemFree" of /proc/meminfo. If the memory is insufficient it is necessary to change the system configuration and allocate more memory, or to cancel the settings for the I/O Response Time Assurance Function.
Refer to the installation guide for the memory necessary for the I/O Response Time Assurance Function.

22021

WARNING: sfdsksys: failed to create a new thread, the I/O monitor was inactivated

Explanation

I/O Response Time Assurance Function is disabled because the creation of thread has failed.

Resolution

Check if the memory is insufficient, see "MemFree" of /proc/meminfo. If the memory is insufficient it is necessary to change the system configuration and allocate more memory, or to cancel the settings for the I/O Response Time Assurance Function.
Refer to the installation guide for the memory necessary for the I/O Response Time Assurance Function.

22022

WARNING: sfdsksys: no enough address space, the I/O monitor was inactivated

Explanation

I/O Response Time Assurance Function is disabled because the memory allocation has failed.

Resolution

Check if the memory is insufficient, see "MemFree" of /proc/meminfo. If the memory is insufficient it is necessary to change the system configuration and allocate more memory, or to cancel the settings for the I/O Response Time Assurance Function. Refer to the installation guide for the memory necessary for the I/O Response Time Assurance Function.

22023

WARNING: sfdsk: slice is pre-detached by an I/O error:

volume info: devno(maj,min)=v_devno(v_maj,v_min)

device info: devno(maj,min)=p_devno(p_maj,p_min)

Explanation

The slice in which I/O error has occurred was detached from the netmirror volume prior to changing the slice status.

Resolution

If the message 44005 is output on one of the nodes after this message, no action is required for this message.

Otherwise, perform the following resolution:

- If the cluster node is in LEFTCLUSTER

Restore the LEFTCLUSTER. For details on how to check the LEFTCLUSTER and restoration procedure, see "[7.16.8 Restoration when the Latest Disk Cannot Be Selected Automatically](#)" and "PRIMECLUSTER Cluster Foundation Configuration and Administration Guide."

- If the cluster node is not in LEFTCLUSTER

Closed class may occur due to disk failure. Identify the cause by referring to the object status, GDS log message, and rsyslogd message. For procedures on how to check for closed class and restoration procedure, see "[\(1\) Class becomes closed status during operation.](#)" in "[D.1.4 Class Status Abnormality.](#)"

C.2.4 Information Messages (24000-24099)

24000

sfdsk: driver started up

Explanation

The driver has been installed into the system.

24001

sfdsk: received shutdown request

Explanation

A shutdown request from the sdxservd daemon has been received.

24002

`sforsk: 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: sforsk: 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: sforsk: 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=pdev

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: sforsk: 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=pdev
```

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=pdev

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=pdev

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=pdev
```

Explanation

The request to cancel I/O requests on the physical device indicated in device info (an ioctl request for the mphpd or mplb driver) ended normally.

24011

```
NOTICE: sfdsk: HBA='adapter' varyio=enabled
```

Explanation

Variable length IO is available for *adapter*, a SCSI host bus adapter, and SCSI disks connected to *adapter*.

24012

```
NOTICE: sfdsk: HBA='adapter' varyio=disabled
```

Explanation

Variable length IO is unavailable for *adapter*, a SCSI host bus adapter, and SCSI disks connected to *adapter*.

24013

```
NOTICE: sfdsk: variable length IO is enabled
```

Explanation

The sfdsk driver can make use of variable length IO.

24014

```
NOTICE: sfdsk: variable length IO is disabled
```

Explanation

The sfdsk driver cannot make use of variable length IO.

24015

```
NOTICE: sfdsk: I/O break on disk:
```

```
device info: devno(maj,min)=p_devno(p_maj,p_min)
           devname=pdev
```

Explanation

The I/O request for the physical disk displayed in device info was aborted.

24016

```
NOTICE: sfdsk: disable error disk:
           device info: devno(maj,min)=p_devno(p_maj,p_min)
           devname=pdev
```

Explanation

An abnormality occurred in the open request to the physical disk, or the I/O request to the physical disk was delayed and became abnormal. Due to this, further access to the physical disk is prevented.

24017

```
NOTICE: sfdsk: access is disabled on disk:
           device info: devno(maj,min)=p_devno(p_maj,p_min)
           devname=pdev
           error info: func=sdxfunc
```

Explanation

Because access to the physical disk is prevented the I/O request terminated abnormally.

C.2.5 Internal Error Messages (26000-26001)

26000

```
severity: sfdsk: internal error, details
```

Explanation

Internal error occurred. *details* indicates the cause of the error.
severity indicates the severity of the message.

a) If part of or all of the disks registered with the root class are not recognized by the OS at OS startup, the following error message may be output.

```
ERROR: sfdsk: internal error, func=gds_dev_open(*,*,*) FAIL
```

b) If the major number 487 is already in use, the following message may be output at system startup.

```
ERROR: sfdsk: internal error, func=register_blkdev(sfdsk), errno=16
```

c) In the environment where the mirroring among servers is used, the following message may be output at node startup. (Values will be output to *X*)

```
ERROR: sfdsk: internal error,func=ioctl(dev=XXX,XXX, cmd=0xXXXXXXXX) FAIL(-19) proc=sdxservd(XXX)
```

d) When the `sdxfix -C` command is executed, the following message may be output.

`dev` indicates the device number of the disk.

```
WARNING: sfdsk: internal error, cannot get disk offset. dev=0xNNNNNNNN, cmd=HDIO_GETGEO, errno=5
```

Resolution

In the event of a), the error has no effect on the system. Restore the state where the disks are recognized by the OS and reboot the system, and the message will no longer be output.

In the event of b), change the `sfdsk` driver major number to a number other than 487. For the major number changing method, see "[8.8.1 Changing sfdsk Driver Major Number](#)."

In the event of c), the system may have been in any one of the following status:

(c-1)

Status: Another node has been stopped.

Resolution: The error has no effect on the system. No action is required.

(c-2)

Status: An error has occurred in the network used in the mirroring among servers.

Resolution: Restore the network error.

In the event of d), a disk failure may have occurred. Check the disk of the displayed device number output to `dev`, and restore the disk if it has an error.

In other events, collect investigation material and contact field engineers.

26001

```
sfdsklog: internal error, details
```

Explanation

Internal error occurred. *details* indicates the cause of the error.

a) If the major number 234 is already in use at system startup, the following message may be output on the system log.

```
sfdsklog: internal error, func=register_chrdev(sfdsklog) fail(-16)
```

At that time, loading the `sfdsklog` driver in GDS fails and GDS is not started. Therefore, the shared disk resource registration (`clautoconfig` command), GDS operation and accessing to the GDS volume fail.

Resolution

In the event of a), change the `sfdsklog` driver major number to 300 or higher.

For the `sfdsklog` driver major number changing method, see "[8.8.2 Changing sfdsklog Driver Major Number](#)."

In events other than a), collect investigation material and contact field engineers.

C.3 Daemon Messages

Daemon messages are output to the GDS log file.

```
/var/opt/FJSSVsdx/msglog/sdxservd.log
```

Depending on the settings of `rsyslogd`, daemon message will also be output to the `rsyslogd` log file and console.

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.

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 (sd <i>X</i> , mpath <i>X</i> , emcpower <i>X</i> , and vd <i>X</i>) <i>X</i> indicates the device identifier.
<i>pslice</i>	Physical slice name (sd <i>Xn</i> , mpath <i>Xpn</i> , emcpower <i>Xn</i> , and vd <i>Xn</i>) <i>X</i> indicates the device identifier and <i>n</i> indicates the slice number. The device tree file name may also be used. e.g.) "sd@2,0:a,raw"
<i>psdevtree</i>	Device tree path name to a physical slice e.g.) /dev/sd@2,0: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)
<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)

Variable names	Descriptions
<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
<i>val</i>	Parameter value

GDS: Global Disk Services

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

GDS: Global Disk Services

facility.level is the facility and priority of a message passed from the GDS daemon to rsyslogd(8).

The output destinations of messages are defined in the /etc/rsyslog.conf configuration file of rsyslogd(8) and can be changed by modifying the definitions in /etc/rsyslog.conf. For details, see the manual page of rsyslog.conf(5).

C.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 field engineers.

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 field engineers.

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 rsyslogd.

Resolution

When error number information is insufficient for identifying the cause, collect investigation material, and contact field engineers.

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.

C.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=*blknos1*,

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=*blknos1*,

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.

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 "[D.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 ["D.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 ["D.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 ["D.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 ["D.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 field engineers.

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 field engineers.

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 field engineers.

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 "[D.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 "[D.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 "[D.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 "[D.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 rsyslogd message.
For information on recovery, see "[D.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 rsyslogd message.
For information on recovery, see "[D.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 rsyslogd message.

For information on recovery, see "[D.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 rsyslogd message.

For information on recovery, see "[D.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 field engineers.

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 field engineers.

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 field engineers.

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 resource registration was performed on a node where resource registration was not complete.
- b. In a cluster environment where 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 "D.1.10 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 ETERNUS Disk storage system 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 ETERNUS Disk storage system 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, the log messages for disk drivers, and the log messages for Dell EMC storage units 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, the log messages for disk drivers, and the log messages for Dell EMC storage units 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, ETERNUS Disk storage system 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, ETERNUS Disk storage system 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, ETERNUS Disk storage system 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 ETERNUS Disk storage system 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 ETERNUS Disk storage system 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 ETERNUS Disk storage system 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 the log messages for GDS, the log messages for disk drivers, and the log messages for Dell EMC storage units that were output right before the occurrence of the error, and restore.

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 the log messages for GDS, the log messages for disk drivers, and the log messages for Dell EMC storage units that were output right before the occurrence of the error, and restore.

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 the log messages for GDS, the log messages for disk drivers, and the log messages for Dell EMC storage units that were output right before the occurrence of the error, and restore.

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 the log messages for GDS, the log messages for disk drivers, and the log messages for Dell EMC storage units that were output right before the occurrence of the error, and restore.

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 the log messages for GDS, the log messages for disk drivers, and the log messages for Dell EMC storage units that were output right before the occurrence of the error, and restore.

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 the log messages for GDS, the log messages for disk drivers, and the log messages for Dell EMC storage units that were output right before the occurrence of the error, and restore.

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 the log messages for GDS, the log messages for disk drivers, and the log messages for Dell EMC storage units that were output right before the occurrence of the error, and restore.

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 the log messages for GDS, the log messages for disk drivers, and the log messages for Dell EMC storage units that were output right before the occurrence of the error, and restore.

42044

```
ERROR: disk is bound to RAW device. disabled disk, class=class
```

Explanation

The disk *disk* registered with GDS was found that it was bound to a RAW device, and the disk was disabled.

Resolution

Cancel the bind to a RAW device using the raw(8) command for the disk registered with GDS. You may not bind disks registered with GDS to RAW devices.

42049

```
ERROR: proxy: failed to copy with QuickOPC, source=disk.volume, target=disk.volume, class=class
```

Explanation

Copying failed due to an I/O error caused while conducting copying between the proxy volume *proxy* and the master volume with QuickOPC function.

"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 by referring to the log messages for GDS, disk driver and ETERNUS Disk storage system that were output right before the occurrence of the error, and restore it.

42050

```
ERROR: proxy: failed to start QuickOPC session, source=disk.volume, target=disk.volume, class=class
```

Explanation

An error occurred between proxy volume *proxy* and the master volume when QuickOPC session 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 by referring to the log messages for GDS, disk driver and ETERNUS Disk storage system that were output right before the occurrence of the error, and restore it.

42051

ERROR: proxy: failed to stop QuickOPC session, source=disk.volume, target=disk.volume, class=class

Explanation

An error occurred between proxy volume *proxy* and the master volume when QuickOPC session was stopped.

"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 by referring to the log messages for GDS, disk driver and ETERNUS Disk storage system that were output right before the occurrence of the error, and restore it.

42052

ERROR: failed to detach slice:

volume info: devno(maj,min)=v_devno(v_maj,v_min)

device info: devno(maj,min)=p_devno(p_maj,p_min)

Explanation

The slice detachment by the I/O error failed..

Resolution

Recover the LEFTCLUSTER state. For how to recover the LEFTCLUSTER state, see "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide."

42054

ERROR: volume: failed to lock netmirror volume, class=class

Explanation

Failed to lock the netmirror volume. After outputting this message, the class *class* is closed down.

Resolution

Stop both the nodes first, and then start from a node other than the nodes of which the message was output.

42055

ERROR: check the iSCSI device has timeout, iSCSI=num

Explanation

This message is output on the system log when the iSCSI disk registered in a class cannot be recognized after the number of seconds specified by the tuning parameter SDX_STARTUP_WAIT_TIME. *num* is the number of the iSCSI devices that cannot be recognized.



See

.....
For SDX_STARTUP_WAIT_TIME, see "6.2 Using iSCSI Disk."
.....

Resolution

Check the iSCSI disk. If the disk has a failure, resotre the failed part. If the disk does not have a failure, take the following procedure.

1. On the system log, check what time the following message has been output.

```
Mon Day HH:MM:SS node devlabel: devlabel service started/restarted
```

2. The following message is output on the system log after the above message in procedure 1 has been output. Check what time the last message has been output.

[sd device]

```
Mon Day HH:MM:SS node kernel: sd host:channel:id:lun:[device] Attached SCSI disk
```

[mpath device]

```
Mon Day HH:MM:SS node multipathd: device [maj:min]: path added to devmap dmpath
```

Each variable indicates the following meaning.

Variable	Meaning
<i>Mon</i>	month
<i>Day</i>	date
<i>HH</i>	hour
<i>MM</i>	minute
<i>SS</i>	second
<i>node</i>	node name
<i>host</i>	host number
<i>channel</i>	channel number
<i>id</i>	id number
<i>lun</i>	lun number
<i>device</i>	sd <i>X</i> (<i>X</i> indicates the device identifier)
<i>maj</i>	major number
<i>min</i>	minor number
<i>dmpath</i>	mpath <i>X</i> (<i>X</i> indicates the device identifier)

3. Calculate the time by the following formula and set the time for the tuning parameter SDX_STARTUP_WAIT_TIME.

[Formula]

```
Time difference between the messages 1 and 2 above + 10 + 1800 (seconds)
```



See

For how to set up SDX_STARTUP_WAIT_TIME, see "6.2 Using iSCSI Disk."

42056

```
ERROR: netmirror DB file is invalid, class=class, volume=volume
```

Explanation

The disk containing the latest data cannot be determined because inconsistency exists in the information to specify the disk containing the latest data.

Resolution

See ["7.16.8 Restoration when the Latest Disk Cannot Be Selected Automatically."](#)

42057

ERROR: failed to add boot entry *entry*, exit-status=*exitstat*

Explanation

The boot entry *entry* cannot be added to the boot option selection screen of the EFI boot manager.

The cause for this error message is as follows:

If the exit-status is other than 137, the `efibootmgr(8)` command has ended abnormally with the exit status *exitstat*.

If the exit-status is 137, the `efibootmgr(8)` command has not ended within 300 seconds.

Resolution

Collect investigation material, and contact field engineers.

42058

ERROR: failed to delete boot entry *entry*, exit-status=*exitstat*

Explanation

The boot entry *entry* cannot be deleted from the boot option selection screen of the EFI boot manager.

The cause for this error message is as follows:

If the exit-status is other than 137, the `efibootmgr(8)` command has ended abnormally with the exit status *exitstat*.

If the exit-status is 137, the `efibootmgr(8)` command has not ended within 300 seconds.

Resolution

Collect investigation material, and contact field engineers.

42059

ERROR: failed to get boot entry information, exit-status=*exitstat*

Explanation

The boot entry information of the boot option selection screen of the EFI boot manager cannot be obtained.

The cause for this error message is as follows:

If the exit-status is other than 137, the `efibootmgr(8)` command has ended abnormally with the exit status *exitstat*.

If the exit-status is 137, the `efibootmgr(8)` command has not ended within 300 seconds.

Resolution

Collect investigation material, and contact field engineers.

42060

ERROR: failed to change display order of boot entry, exit-status=*exitstat*

Explanation

The display order of boot entries of the boot option selection screen of the EFI boot manager cannot be changed.

The cause for this error message is as follows:

If the exit-status is other than 137, the efibootmgr(8) command has ended abnormally with the exit status *exitstat*.

If the exit-status is 137, the efibootmgr(8) command has not ended within 300 seconds.

Resolution

Collect investigation material, and contact field engineers.

42061

ERROR: failed to open temporary file, sdxfunc=sdxfunc, errno=errno

Explanation

A temporary file cannot be opened.

Resolution

Collect investigation material, and contact field engineers.

42062

ERROR: class: timeout on waiting configuration database relocation

Explanation

Timeout occurred while waiting for the completion of the configuration database relocation of the class *class*.

Resolution

If this message is output when starting the system, start the system again.

If this message is output when restoring the class, execute the command again.

If this message is still output, collect investigation material and contact field engineers.

42063

ERROR: class: closing down class, failed to update netmirror NVURM information

Explanation

Class *class* was closed because NVURM information of the netmirror group could not be updated.

Resolution

Take the following actions:

- When no capacity available under /var/opt
Delete unnecessary files to secure 100MB or more free capacities under /var/opt. After that, restart the system.
 - When enough capacity available under /var/opt
Output the /var/opt/FJSVsdx/log/.sdxnetmirror_nvurm_ class file by the cat command, and check that the output is not garbled.
If it is garbled, delete the file and then restart the system.
If it is not garbled, collect investigation material and contact field engineers.
-

42064

ERROR: failed to update netmirror NVURM information, class=class, volume=volume

Explanation

NVURM information of the netmirror group could not be updated.

Resolution

Take the following actions:

- When no capacity available under /var/opt
Delete unnecessary files to secure 100MB or more free capacities under /var/opt. After that, restart the system.
- When enough capacity available under /var/opt
Output the /var/opt/FJSVsdX/log/.sdXnetmirror_nvurm_*class* file by the cat command, and check that the output is not garbled.
If it is garbled, delete the file and then restart the system.
If it is not garbled, collect investigation material and contact field engineers.

42065

```
ERROR: class: closing down class, failed to update netmirror device information
```

Explanation

Class *class* was closed because the device information of the netmirror group could not be updated.

Resolution

Take the following actions:

- When no capacity available under /var/opt
Delete unnecessary files to secure 100MB or more free capacities under /var/opt. After that, restart the system.
- When enough capacity available under /var/opt
Output the /var/opt/FJSVsdX/log/.sdXnetmirror_disable.db file by the cat command, and check that the output is not garbled.
If it is garbled, delete the file and then restart the system.
If it is not garbled, collect investigation material and contact field engineers.

42066

```
ERROR: hangup detected
```

Explanation

Either a hang-up in the GDS daemon or the driver was detected.
The system panics after this message is output.

42067

```
ERROR: class: failed to start volume volume, status slice in local node must be recovered
```

Explanation

The netmirror volume *volume* of class *class* cannot start because the slice of the operating node is detached (is on *status*).

Resolution

Restore the slice status of the operating node before forcibly starting the cluster application.
Refer to "[7.16.6 Operation on Only One Node](#)" for restoration procedure.

42068

ERROR: failed to update netmirror device information, class=class

Explanation

The device information of netmirror group in class *class* has failed to be updated.

Resolution

Take the following actions:

- When no capacity available under /var/opt

Delete unnecessary files to secure 100MB or more free capacities under /var/opt. After that, restart the system.

- When enough capacity available under /var/opt

Output the /var/opt/FJSVsdx/log/.sdxnetmirror_disable.db file by the cat command, and check that the output is not garbled. If it is garbled, delete the file and then restart the system.

If it is not garbled, collect investigation material and contact field engineers.

42069

ERROR: lock is set on volume volume, node=node

Explanation

Failed in starting volume due to the starting lock of volume *volume* is set in the node *node*.

Resolution

To start the volume, turn off the "Lock volume" mode, or use the -e unlock option.

However, if *volume* is the netmirror volume, lock the netmirror volume (including specifying -e unlock option) only as the restoration procedure in "[7.16 Mirroring Among Servers](#)."

If this message is output during the startup of cluster application that includes netmirror volume, see "[7.16.8 Restoration when the Latest Disk Cannot Be Selected Automatically](#)" and restore.

For details, see "[7.16.5 Lock Volume of the Netmirror Volume](#)."

42070

ERROR: class: failed to start volume volume, slice in local node must be recovered

Explanation

Netmirror volume *volume* of the class *class* cannot be started due to failure in the slice state of the started node.

Resolution

Check the slice status on the disk of the started node before forcibly starting the cluster application and then restore.

- When the slice is in ACTIVE or STOP status

Disk failure may have occurred.

For resolutions, see "(1) [Mirror slice configuring the mirror volume is in INVALID status.](#)" in "[D.1.1 Slice Status Abnormality](#)."

- For others

See "[D.1.1 Slice Status Abnormality](#)."

42071

ERROR: val: interval must be integer

Explanation

The hang-up monitoring function cannot be started because the value of monitoring interval (SDX_HCHK_INTERVAL) of hang-up monitoring function configured in the GDS settings file /etc/opt/FJVSdx/sdx.cf is not a numerical value. This message is not output to the log file (sdxservd.log) of daemon indicated in "[C.3 Daemon Messages](#)."

Resolution

1. Refer to "How to change parameters" in "[6.12.2 Hang-up Monitoring Function](#)" in order to edit the value configured in SDX_HCHK_INTERVAL of GDS settings file /etc/opt/FJVSdx/sdx.cf.
2. Restart the system or restart the service using the following command.

```
# /usr/bin/systemctl restart fjsvdxhealth.service
```

3. Check that the service is in running state with the following command.

```
# /usr/bin/systemctl status fjsvdxhealth.service
* fjsvdxhealth.service - PRIMECLUSTER Global Disk Services hang up monitor service
...
Active: active (running) ...
...
```

42072

ERROR: val: too small interval

Explanation

The hang-up monitoring function cannot be started because the value of monitoring interval (SDX_HCHK_INTERVAL) of hang-up monitoring function configured in the GDS settings file /etc/opt/FJVSdx/sdx.cf is smaller than the allowed value. This message is not output to the log file (sdxservd.log) of daemon indicated in "[C.3 Daemon Messages](#)."

Resolution

1. Refer to "How to change parameters" in "[6.12.2 Hang-up Monitoring Function](#)" in order to edit the value configured in SDX_HCHK_INTERVAL of GDS settings file /etc/opt/FJVSdx/sdx.cf.
2. Restart the system or restart the service using the following command.

```
# /usr/bin/systemctl restart fjsvdxhealth.service
```

3. Check that the service is in running state with the following command.

```
# /usr/bin/systemctl status fjsvdxhealth.service
* fjsvdxhealth.service - PRIMECLUSTER Global Disk Services hang up monitor service
...
Active: active (running) ...
...
```

42073

ERROR: val: too large interval

Explanation

The hang-up monitoring function cannot be started because the value of monitoring interval (SDX_HCHK_INTERVAL) of hang-up monitoring function configured in the GDS settings file /etc/opt/FJVSdx/sdx.cf is larger than the allowed value. This message is not output to the log file (sdxservd.log) of daemon indicated in "[C.3 Daemon Messages.](#)"

Resolution

1. Refer to "How to change parameters" in "[6.12.2 Hang-up Monitoring Function](#)" in order to edit the value configured in SDX_HCHK_INTERVAL of GDS settings file /etc/opt/FJVSdx/sdx.cf.
2. Restart the system or restart the service using the following command.

```
# /usr/bin/systemctl restart fjsvdxhealth.service
```

3. Check that the service is in running state with the following command.

```
# /usr/bin/systemctl status fjsvdxhealth.service
* fjsvdxhealth.service - PRIMECLUSTER Global Disk Services hang up monitor service
...
Active: active (running) ...
...
```

42074

ERROR: val: criterion must be integer

Explanation

The hang-up monitoring function cannot be started because the value of conclusion time (SDX_HCHK_CRITERION) of hang-up monitoring function configured in the GDS settings file /etc/opt/FJVSdx/sdx.cf is not a numerical value. This message is not output to the log file (sdxservd.log) of daemon indicated in "[C.3 Daemon Messages.](#)"

Resolution

1. Refer to "How to change parameters" in "[6.12.2 Hang-up Monitoring Function](#)" in order to edit the value configured in SDX_HCHK_CRITERION of GDS settings file /etc/opt/FJVSdx/sdx.cf.
2. Restart the system or restart the service using the following command.

```
# /usr/bin/systemctl restart fjsvdxhealth.service
```

3. Check that the service is in running state with the following command.

```
# /usr/bin/systemctl status fjsvdxhealth.service
* fjsvdxhealth.service - PRIMECLUSTER Global Disk Services hang up monitor service
...
Active: active (running) ...
...
```

42075

ERROR: val: too small criterion

Explanation

The hang-up monitoring function cannot be started because the value of conclusion time (SDX_HCHK_CRITERION) of hang-up monitoring function configured in the GDS settings file /etc/opt/FJSVsdx/sdx.cf is smaller than the allowed value. This message is not output to the log file (sdxservd.log) of daemon indicated in "[C.3 Daemon Messages](#)."

Resolution

1. Refer to "How to change parameters" in "[6.12.2 Hang-up Monitoring Function](#)" in order to edit the value configured in SDX_HCHK_CRITERION of GDS settings file /etc/opt/FJSVsdx/sdx.cf.
2. Restart the system or restart the service using the following command.

```
# /usr/bin/systemctl restart fjsvdxhealth.service
```

3. Check that the service is in running state with the following command.

```
# /usr/bin/systemctl status fjsvdxhealth.service
* fjsvdxhealth.service - PRIMECLUSTER Global Disk Services hang up monitor service
...
Active: active (running) ...
...
```

42076

ERROR: val: too large criterion

Explanation

The hang-up monitoring function cannot be started because the value of conclusion time (SDX_HCHK_CRITERION) of hang-up monitoring function configured in the GDS settings file /etc/opt/FJSVsdx/sdx.cf is larger than the allowed value. This message is not output to the log file (sdxservd.log) of daemon indicated in "[C.3 Daemon Messages](#)."

Resolution

1. Refer to "How to change parameters" in "[6.12.2 Hang-up Monitoring Function](#)" in order to edit the value configured in SDX_HCHK_CRITERION of GDS settings file /etc/opt/FJSVsdx/sdx.cf.
2. Restart the system or restart the service using the following command.

```
# /usr/bin/systemctl restart fjsvdxhealth.service
```

3. Check that the service is in running state with the following command.

```
# /usr/bin/systemctl status fjsvdxhealth.service
* fjsvdxhealth.service - PRIMECLUSTER Global Disk Services hang up monitor service
...
Active: active (running) ...
...
```

42077

ERROR: interval val isn't larger than criterion val

Explanation

The hang-up monitoring function cannot be started because the value of conclusion time (SDX_HCHK_CRITERION) of hang-up monitoring function configured in the GDS settings file /etc/opt/FJVSdx/sdx.cf is greater than or equal to the value of monitoring interval (SDX_HCHK_INTERVAL).

This message is not output to the log file (sdxservd.log) of daemon indicated in "[C.3 Daemon Messages](#)."

Resolution

1. Refer to "How to change parameters" in "[6.12.2 Hang-up Monitoring Function](#)" in order to edit the value configured in either of SDX_HCHK_CRITERION or SDX_HCHK_INTERVAL of GDS settings file /etc/opt/FJVSdx/sdx.cf.
2. Restart the system or restart the service using the following command.

```
# /usr/bin/systemctl restart fjsvdxhealth.service
```

3. Check that the service is in running state with the following command.

```
# /usr/bin/systemctl status fjsvdxhealth.service
* fjsvdxhealth.service - PRIMECLUSTER Global Disk Services hang up monitor service
...
Active: active (running) ...
...
```

42078

ERROR: failed to update netmirror device information:

volume info: devno(maj,min)=v_devno(v_maj,v_min)

device info: devno(maj,min)=p_devno(p_maj,p_min)

Explanation

The device information of the netmirror group failed to be updated.

Resolution

Take the following actions:

- When no capacity available under /var/opt

Delete unnecessary files to secure 100MB or more free capacities under /var/opt. After that, restart the system.

- When enough capacity available under /var/opt

Output the /var/opt/FJVSdx/log/.sdxnetmirror_disable.db file by the cat command, and check that the output is not garbled.

If it is garbled, delete the file and then restart the system.

If it is not garbled, collect investigation material and contact field engineers.

42079

ERROR: failed to resume suspended volume by I/O error:

volume info: devno(maj,min)=v_devno(v_maj,v_min)

device info: devno(maj,min)=p_devno(p_maj,p_min)

Explanation

The suspended I/O failed to be resumed due to an I/O error.

Resolution

Collect investigation material and contact field engineers.

C.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=*blkno*, 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=*blkno*, 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 for the physical slice *pslice* terminated abnormally.

The following messages may be output when a node is booted, investigation material is collected (*pclsnap* 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 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 *volume* was unsuccessful.

The following potential causes can be considered:

- a) A disk malfunction occurred and there are an insufficient number of normally operating disks which can store the logs (NVURM) for the Just Resynchronization Mechanism (JRM).
- b) The size of the private slice of the class *class* is insufficient and it is impossible to store the logs (NVURM) for the JRM to the private slices.

If one has added to the class a disk with a capacity larger than the disk first registered to the class, it is possible that b) is the cause. For details on insufficient private slice size, refer to "[2.3.9 Disk Size](#)."

Resolution

Check the disk status within *class*. If a disk failure has occurred, identify the cause by referring to disk driver log message, and recover the disk.

If *volume* is not a mirror volume (such as a single volume), the Just Resynchronization Mechanism is not necessary. This being the case, if *volume* is not a mirror volume and the cause is b), turning the *volume's* JRM attribute to off or ignoring this message poses no problems. For turning the JRM attribute off, refer to "[B.1.7 sdxattr - Change Objects Attributes](#)."

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.

This message is output if starting a node at the time that all disks registered in the class cannot be found due to disk array case down or cable disconnection.

Resolution

Add *class* which is output at the message to `/etc/opt/FJSVsdx/sysdb.d/class.db` so that this message will not be output.

Example) If *class* is `class0001`

Add `class0001` at the position before "# Disk Class List."

```
# cat class.db
class0001 <--- Add
# Disk Class List
```

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 "[D.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 is output via rsyslogd. Although the message will not be output on GDS logfile, it does not affect other processes.

Resolution

Collect investigation material, and contact field engineers 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 is output via rsyslogd. Although the message will not be output on GDS logfile, it does not affect other processes.

Resolution

Collect investigation material and contact field engineers.

44025

WARNING: failed to reply message to node node, details

Explanation

Replying to remote node *node* was unsuccessful.

Resolution

Collect investigation material and contact field engineers.

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 field engineers.

44027

WARNING: sdxfunc: no enough address space, ofunc=ofunc, errno=errno

Explanation

The OS *ofunc* 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 is output via rsyslogd.

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 field engineers.

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 parted(8) command to recover format information.

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 disk label of *device* failed. A disk failure may have occurred.

Resolution

To recover the disk label of *device*, use the parted(8) 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 disk label of *device* failed. A disk failure may have occurred.

Resolution

To recover the disk label of *device*, use the parted(8) command.

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 ETERNUS Disk storage system 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 ETERNUS Disk storage system that were output right before the occurrence of the error, and restore.

44039

WARNING: backup_efi_files_grub: details

Explanation

In the process of setting or canceling the system disk mirroring, the `grub.conf` file or the `grub.cfg` file cannot be backed up. *details* indicates the cause of the error.

Resolution

If no other message was output, the error has no effect on the system. There is no need to work around.
If some other message was output, work around the error message.

44040

WARNING: backup_efi_files_initrd: details

Explanation

In the process of setting or canceling the system disk mirroring, an initial RAM disk cannot be backed up.
details indicates the cause of the error.

Resolution

If no other message was output, the error has no effect on the system. There is no need to work around.
If some other message was output, work around the error message.

44041

WARNING: failed to update netmirror DB file, class=class, volume=volume

Explanation

The device information of the netmirror group has failed to be updated.

Resolution

- If the INVALID status slice exists within the netmirror volume *volume*, and other nodes are LEFTCLUSTER or DOWN.
Synchronization copying of netmirror volume fails.
Restore the node first. After that, execute the `sdxcopy -B` command, for example, to perform the synchronization copying.
For how to recover from the LEFTCLUSTER state, see "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide."
- Otherwise
Take the following procedure to restore.
 1. Restore the file system or file with all of the nodes.
 - When no available capacity under `/var/opt`
Delete unnecessary files to secure 100MB or more free capacities under `/var/opt`. After that, restart the system.
 - When enough capacity available under `/var/opt`
Output the `/var/opt/FJSVsdx/log/.sdxnetmirror_disable.db` file by the `cat` command, and check that the output is not garbled.
If it is garbled, delete the file and then restart the system.
If `/var/opt` on all of the nodes has capacity and the file is not garbled, collect investigation material and contact field engineers.
 2. Depending on the slice status within the netmirror volume *volume*, take the following procedure.
 - When the NOUSE status slice does not exist in the netmirror volume *volume*.
Synchronization copying of netmirror volume fails.
Perform the synchronization copying by using the `sdxcopy -B` command, for example.

- When the NOUSE status slice exists in the netmirror volume *volume*.

Restore NOUSE status of the slice. For details, see "(6) Slice configuring the volume is in NOUSE status." in "D.1.1 Slice Status Abnormality."

C.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*.

However, the just resynchronization which is executed after the following events copies the entire slice.

- When an I/O error occurs in the memory area of difference information inside the private slice
- When the just resynchronization in process is canceled due to shutting the node, switching the cluster application, the `sdxcopy` command, etc.

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 *psdevtree* 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 canceled.

"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 canceled.

"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.

46113

INFO: physical disk not found, devno(maj,min)=p_devno(p_maj,p_min)

Explanation

The physical disk registered with GDS was not found.

p_devno(p_maj,p_min) is the device number (major number, minor number) of the physical disk used when it was registered with a class.

46114

INFO: restarting sdxservd to cancel outstanding event operations, because the shutdown shared class request has been waited for details seconds

Explanation

While operations on the shared class to be separated from the cluster configuration were waited for the completion, the operations did not end within the set time *details*. Therefore, the `sdxservd` daemon is restarted and the operations are canceled. *details* is preset time to wait for the completion of operations on a shared class.

46115

INFO: identify SDX_EFI_DISK as "on", GPT labeled disks exist in class

Explanation

A node where `SDX_EFI_DISK=on` is not described in the GDS configuration parameter file `/etc/opt/FJSVsdx/sdx.cf` exists, however the node is identified as `SDX_EFI_DISK=on` because GPT labeled disks are registered with a shared or local class. Even if this message is output, the operations of the system and GDS are not affected.

46116

INFO: proxy: QuickOPC session started, source=disk.volume, target=disk.volume, class=class

Explanation

A QuickOPC session between proxy volume *proxy* and master volume started. "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.

46117

INFO: proxy: completed copying with QuickOPC, source=disk.volume, target=disk.volume, class=class

Explanation

Copying between proxy volume *proxy* and master volume with the QuickOPC function is completed. "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.

46118

INFO: proxy: canceled copying with QuickOPC, source=disk.volume, target=disk.volume, class=class

Explanation

Copying between proxy volume *proxy* and master volume with the QuickOPC function was canceled. "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.

46119

INFO: proxy: QuickOPC session stopped, source=disk.volume, target=disk.volume, class=class

Explanation

QuickOPC session between proxy volume *proxy* and master volume with the QuickOPC function was stopped. "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.

46120

```
INFO: class: recovered configuration database replicas:
      psdevtree
```

Explanation

The configuration database of the class *class* stored in the private slice *psdevtree* was recovered.

C.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.

a) The following messages may be output when an I/O error occurs.

```
sdxservd: ERROR: module: internal error, sdxfunc=dbrw_read_dc(), rval=-1, sdxerrno=4120
sdxservd: ERROR: module: internal error, sdxfunc=dbrw_read_dc(), rval=-1, sdxerrno=4121
```

b) The following message 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).

```
sdxservd: WARNING: module: internal error, ofunc=read, rval=0, errno=2
```

c) If part of or all of the disks registered with the root class are not recognized by the OS at OS startup, the following message may be output.

```
sdxservd: ERROR: sendtd_root: internal error, sdxfunc=sv_pd_find_bytd, rval=-1, sdxerrno=2
```

d) If an attempt to boot a proxy volume is made when synchronization copying between the joined master and proxy is in process, the following message may be output.

```
sdxservd: WARNING: mv_start: internal error, sdxfunc=mv_copy_dispatch, sdxerrno=11
```

e) The following message may be output if restoring physical disks of the swapped internal disk when there is a device name change.

```
sdxservd:ERROR: internal error (pd_find_bytd,sdxerrno=2)
```

f) If the mirroring among servers is used when the shutdown facility (SF) of PRIMECLUSTER is not set properly, the following message may be output when one node is restarted. In this case, the synchronization copying fails, and the slice of the restarted node becomes INVALID.

```
sdxservd: INFO: tdmcopy_begin: internal error, can't find COPY slice
```

g) If the cluster application is started in the configuration of mirroring among servers, the following messages may be output by the GDS process running in the Online process.


```
sdxservd: ERROR: dev_rename_minor_node: internal error, /dev/sfdsk/class/dsk/volume -> /dev/sfdsk/class/dsk/volume: rename error, errno=2
sdxservd: ERROR: rcv_dev_rename_minor_node: internal error, sdxfunc=dev_rename_minor_node, sdxerrno=1026, class=class
```

Resolution

In the event of a), 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 b), if no error message was output around the same period of time, GDS was behaving normally and the message b) 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 the event of c), the error has no effect on the system. Restore the state where the disks are recognized by the OS and reboot the system, and the message will no longer be output.

In the event of d), the error has no effect on the system. There is no need to work around.

In the event of e), reboot the system to solve the device name change.

In the event of f), use the `sdxcopy -B` command to perform the synchronization copying. Set the shutdown facility (SF) properly as well.

In the event of g), check the status of class `class`. If the class is not started, restore it.

Execute the following command on the node where the message is output.

```
# /etc/opt/FJSVsdx/bin/sdxcdcdwn
```

If the `class` information is not displayed, the class is not started.

In this case, see the note "Shutting down the node" of the section "(2) Class cannot be started when booting the system." in "D.1.4 Class Status Abnormality", stop the node where the message is output, and then restart it.

If the `class` information is still not output, see the section "(2) Class cannot be started when booting the system." in "D.1.4 Class Status Abnormality" and restore.

After outputting the message, if the node is already restarted and the class is started, restoration work is not necessary.

In other events, collect investigation material and contact field engineers.

C.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/FJSVsdx/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
<i>device</i>	Physical disk name (sd <i>X</i> , mpath <i>X</i> , emcpower <i>X</i> , and vd <i>X</i>) <i>X</i> indicates the device identifier.
<i>pslice</i>	Physical slice name (sd <i>Xn</i> , mpath <i>Xp</i> <i>n</i> , emcpower <i>Xn</i> , and vd <i>Xn</i>) <i>X</i> indicates the device identifier. <i>n</i> indicates the slice number.
<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

GDS: Global Disk Services

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.

Level of importance	Descriptions
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.

C.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 B Command Reference](#)."

60003

`ERROR: syntax error`

Explanation

Executed command has a syntax error.

Resolution

See the Fix Messages described later, or the "[Appendix B 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 B 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.

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 field engineers.

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 B 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 B 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 B Command Reference](#)."

If "pslice=off: invalid attribute value" is output, the possible cause is that the FJSVsdX-bas package of GDS may not be installed normally. If this is the cause, reinstall FJSVsdX-bas.

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 4 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 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.

Possible cause:

Cause a: Attempted to register the native device (sd device) *object* that configures the mpath device of DM-MP.

Resolution

Change to unused status and execute the command again.

When the possible cause is (Cause a), perform the following resolutions:

- When not using DM-MP, cancel the settings of DM-MP.
- When using DM-MP, register mpath device for the class. Do not register the native device at this time. Also register the native device to the Excluded Device List. For the Exclude Device List, see "[6.1 Creating Excluded Device List.](#)"

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. In the event of c) or d), collect investigation material and contact field engineers.

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 B 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 B 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 sdxcltd
```

Explanation

Connecting to sdxcltd, a GDS cluster linked daemon was unsuccessful.

Resolution

When unable to identify the cause, collect investigation material and contact field engineers.

60046

```
ERROR: physical device driver returned an error, errno=errno
```

Explanation

Physical disk driver returned an error.

An error occurs in software lower than GDS (multipath driver , sd driver, for example) or in a hardware such as a disk device.

Resolution

Identify the cause by referring to error number and message log.

If this message is output during the operation for the class where the netmirror volume exists, the possible case is that the disk of another node cannot be accessed due to a node stop or network error. Check the status of the node or the network.

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 rsyslogd 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 rsyslogd message. For information on recovery, see "[D.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 rsyslogd message. For information on recovery, see "[D.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.

60056

ERROR: device contains overlapping slices

Explanation

The physical disk *device* contains physical slices with overlapping cylinders.

Resolution

Correct the physical slice configuration using the parted(8) command.

60058

ERROR: device : one or more free slice numbers are required, delete a slice by parted(8) command

Explanation

The number of slices created on the physical disk *device* has reached the upper limit (15). The slice numbers are all in use and *device* cannot be registered as a keep disk.

Resolution

Use the parted(8) command to remove one or more slices from *device*.

60060

ERROR: device : no enough unassigned disk space, reserve enough space by parted(8) command

Explanation

The physical disk *device* does not have sufficient free disk space.

Resolution

Use the parted(8) command and create sufficient free disk space.
For the size, see "[2.3.9 Disk Size](#)."

60061

```
ERROR: device : no unassigned disk space nor swap space, reserve enough space by parted(8) command
```

Explanation

The physical disk *device* does not have free disk space or swap area.

Resolution

Use the parted(8) command and create sufficient free disk space or swap area.
For the size, see "[2.3.9 Disk Size](#)."

60062

```
ERROR: device: no enough unassigned disk space nor swap space, reserve enough space by parted(8) command
```

Explanation

Physical disk *device* does not have sufficient free disk space or swap area.

Resolution

Use the parted(8) command and create sufficient free disk space or swap area.
For the size see "[2.3.9 Disk Size](#)."

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.

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 B 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 "[2.3.3 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 B 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

Check the following causes:

- a) The size of *object* is too small.
- b) An attempt to register the disk *object* where the disk label is not the GPT type to a root class.

Resolution

- a) Check the necessary object size and specify a larger object.
- b) Change the disk label of the disk *object* to the GPT type with the parted(8) command.

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

The specified physical slice number *pslice_num* is invalid.

Resolution

For the physical slice number *pslice_num*, specify an integer from 1 to 15.

60093

ERROR: *pslice* is private slice

Explanation

Physical slice *pslice* is a private slice.

Resolution

If this message is output at the time of `sdxdisk -M` command execution, check whether the disk to which the physical slice *pslice* belongs is registered with a class in another domain. If it is registered, the disk cannot be registered with a class. If not, to register the disk to a class, it is necessary to delete the physical slice *pslice*, for example, with the `parted(8)` command.

If this message is output at the time of `sdxdisk -C` command execution, re-execute the `sdxdisk -C` command. When re-executing, do not specify the slice number of the physical slice *pslice* for the `-v` option.

60094

ERROR: *pslice*: corresponding volume attributes must be specified

Explanation

Volume attributes corresponding to physical slice *pslice* is not specified.

Resolution

See "[Appendix B Command Reference](#)."

60095

ERROR: *object*: invalid size

Explanation

Object size of *object* is invalid.

Resolution

See "[Appendix B 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 B 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 B 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 B 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 B 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 B 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 FJSVsdx-bas package of GDS is not installed, multiple disks cannot be connected to a group.

Resolution

If FJSVsdx-bas has not been installed successfully, 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 field engineers.

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 "[D.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 "[D.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 B 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 B 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.

After expanding the physical disk size of the group or single disk specified by *object*, if the volume size in the group or single disk is expanded by the `sdxvolume -S` command, this message is output. In this case, take one of the following actions.

a) When *object* is a mirror group and the number of disks in the group is two or more

Perform the following operations for the physical disk of which the size was expanded.

- a1) Make the physical disk exchangeable with the `sdxswap -O` command or [Swap Physical Disk].
- a2) Restore the status with the `sdxswap -I` command or [Restore Physical Disk].

b) When multiple disks are registered in the class to which *object* belongs

b1) Perform the following operations for the group or single disk specified by *object*.

b1-1) Back up the data of all of the volumes in *object*.

b1-2) Delete all of the volumes from *object*.

b2) Perform the following operations for the disk of which the size was expanded.

b2-1) When the disk is connected to the group, disconnect the disk from the group.

b2-2) Make the physical disk exchangeable with the `sdxswap -O` command or [Swap Physical Disk].

b2-3) Restore the status with the `sdxswap -I` command or [Restore Physical Disk].

b2-4) If the disk is disconnected from the group in step b2-1), connect the disk to the group again.

b3) Perform the following operations for the group or single disk specified by *object*.

b3-1) Recreate the volume deleted in step b1-2).

b3-2) Restore the data of volume backed up in step b1-1).

c) When only one disk is registered in the class to which *object* belongs

c1) Perform the following operations for the group or single disk specified by *object*.

c1-1) Back up the data of all of the volumes in *object*.

- c1-2) Delete all of the volumes from *object*.
- c2) Perform the following operations for the disk of which the size was expanded.
 - c2-1) When the disk is connected to the group, disconnect the disk from the group.
 - c2-2) Delete the disk from the class.
 - c2-3) Register the disk to the class again.
 - c2-4) If the disk is disconnected from the group in step c2-1), connect the disk to the group again.
- c3) Perform the following operations for the group or single disk specified by *object*.
 - c3-1) Recreate the volume deleted in step c1-2).
 - c3-2) Restore the data of volume backed up in step c1-1).

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 B 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 rsyslogd message.

For information on recovery, see "[D.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.

However, if volume is the netmirror volume, lock the netmirror volume (including specifying `-e unlock` option) only as the restoration procedure in ["7.16 Mirroring Among Servers."](#) For details, see ["7.16.5 Lock Volume of the Netmirror Volume."](#)

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 rsyslogd message.

For information on recovery, see "[D.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.

If *volume* indicates the netmirror volume, this message may be output after an error of the network used for the mirroring among server is restored. In this case, if *status* indicates COPY, no resolution is required.

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 "[3.12 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 "[3.12 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 B 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 B 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

Perform any of the following procedures:

- a) Confirm that physical disk *device* is operating normally.
- b) If this message is output when restoring a physical disk of the swapped internal device, there may be a device name change. For checking a device name change, see "[7.3.5.4 Swapping Internal Disks Registered with Root Classes or Local](#)." If there is a device name change, reboot the system to solve it.

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/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

Perform any of the following procedures:

a) Specify a physical disk with sufficient size.

b) If this message is output when restoring a physical disk of the swapped internal device, there may be a device name change. For checking a device name change, see "Swapping Internal Disks Registered with Root Classes or Local Classes [RHEL6]" in "[7.3.5.4 Swapping Internal Disks Registered with Root Classes or Local](#)." If there is a device name change, reboot the system to solve it.

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/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 B 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 "[D.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 "[8.5 Migration from Single Nodes to 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 "[8.5 Migration from Single Nodes to 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 resource registration.

For the resource registration, see "[6.3 Shared Disk Resource Registration](#)."

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 FJSVsdx-bas package of 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 FJSVsdx-bss package of 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 B Command Reference](#)" and specify a correct class name. If it is c) or d), install FJSVsdx-bas or FJSVsdx-bss normally.

60216

ERROR: *disk*: not a keep disk

Explanation

Disk *disk* is not a keep disk.

Resolution

See "[Appendix B 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 B 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, /var, /boot, or /boot/efi, and to perform system disk setting, it is necessary to register *device* with *class*.

Resolution

Complete preparations properly referring to "6.5.2.1.1 Operation Flow" or "Appendix B 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, /var, /boot, or /boot/efi, and to cancel disk mirroring, it is necessary to specify *disk*.

Resolution

Specify all system disks with the slices currently operating as / (root), /usr, /var, /boot, or /boot/efi, 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 B 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 B 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 B Command Reference](#)."

60232

ERROR: param =val: invalid parameter value

Explanation

Parameter value *val* is invalid.

Resolution

See "[Appendix B 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 B 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 a combination of volumes or groups.

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 "[2.3.13 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 "[B.2.1 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 ["D.1.8 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 B 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 field engineers.

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 "D.1.8 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 "D.1.8 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.

Resolution

As the proxy group, specify a group whose volume layout is consistent with the master group's. For the top block (sector) number and size 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 ETERNUS Disk storage system 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 ETERNUS Disk storage system 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 field engineers.

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 field engineers.

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/fstab

Explanation

The volume *volume* is not the alternative volume of a volume described in /etc/fstab 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/fstab file
- Master volume of a proxy volume described in the /etc/fstab file
- Another proxy volume related to the master volume of a proxy volume described in the /etc/fstab file

If this message was output when a master or proxy group was specified and the sdxproxy Root command was executed, see "(2) [Alternative boot environment setup fails resulting in an error numbered 60305. \[EFI\]](#)" in "[D.1.8 Proxy Object Abnormality.](#)"

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, command command failed, exit-status=exitstat
```

Explanation

The command *command* (`fsck(8)` or `xfs_repair(8)`) for the volume *volume* failed, and a file system on *volume* could not be inspected or repaired.

The exit status of the command *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 command *command* that was output right before the occurrence of the error and to the manual for the command *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(8) command failed, exit-status=exitstat
      details
```

Explanation

The `mount(8)` command for the volume *volume* failed, and the file system on *volume* could not be mounted.

The exit status of the `mount(8)` command is *exitstat*. *details* is the error message for the `mount(8)` command.

Resolution

Remove the error after identifying the cause based on *details*, and try the `sdxproxy Root` command again.

60311

```
ERROR: volume: cannot change boot disks, efibootmgr(8) command failed, exit-status=exitstat
```

details

Explanation

The `efibootmgr(8)` command failed and the boot manager setting to boot from the volume *volume* was not set. The exit status of the `efibootmgr(8)` command is *exitstat*. *details* is the error message for the `efibootmgr(8)` 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(8) command failed, exit-status=exitstat

details

Explanation

The `umount(8)` command for the volume *volume* failed, and the file system on *volume* temporarily mounted could not be unmounted. The exit status of the `umount(8)` command is *exitstat*. *details* is the error message for the `umount(8)` 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 B 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 B 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.

60318

ERROR: device is bound to RAW device.

Explanation

The device to be registered has been bound to a RAW device. With GDS, you cannot register disks bound to a RAW device.

Resolution

Register a disk not bound to a RAW device. Alternatively, cancel the bind to a RAW device using the `raw(8)` command and then register.

60319

ERROR: no license

Explanation

The command cannot be executed. The possible causes are as follows.

- a) In a system where the FJSVsdxcmd package of GDS is not installed normally, the sdxconfig command was used.
- b) In a system where the FJSVsdcss package of GDS Snapshot is not installed normally, the sdxshadowdisk command was used.

Resolution

To use the sdxconfig command, install FJSVsdxcmd normally. To use the sdxshadowdisk command, install FJSVsdcss 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

If a message is output right before this error has occurred, follow the message to take corrective actions. If a message is not output, collect investigation material and contact field engineers.

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 field engineers.

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 `sdxconfig 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* 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 field engineers.

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, partition table, 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 `sdxcfg 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 of *device* must have *size* blocks.

Resolution

Exchange physical disk *device* for a physical disk whose private slice 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 `sdxcfg 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 `sdxcfg 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 `sdxcfg Convert` command.

60346

ERROR: class : restoration based on configuration file failed

Explanation

Restoration of the *class* configuration based on configuration file *file* failed.

Resolution

When using multipath, register the native device to each node's Excluded Device List. For the Exclude Device List, see "[6.1 Creating Excluded Device List](#)"

Otherwise, collect investigation material and contact field engineers.

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.

60357

```
ERROR: found invalid device name at the shared disk definition file (rid=va1)
```

Explanation

Shared disk unit resource information contains an error.
va1 is the misdefined shared disk unit's resource ID.

Resolution

The shared disk unit managed as a *va1* resource is not one disk unit if seen from separate nodes.
Remove the incorrect resource information, and create a correct resource.
For shared disk unit removal and registration, see "[6.3 Shared Disk Resource Registration](#)."

60358

```
ERROR: node : physical disk not found
```

Explanation

Some physical disks are no longer recognized by the OS after the node *node* restarted.
Confirm unrecognizable physical disks with the following method.

1. Execute the `sdxinfo -D` command on the relevant *node*.
2. Unrecognizable disks are those indicated by the asterisk (*) in the DEVNAM field in the command output.

Resolution

Unrecognizable physical disks may be failed disks. Check whether they are failed and perform certain maintenance such as disk swap, and then retry the relevant operation.

60359

ERROR: found node(s) which udev environment setting has not been set correctly details

Explanation

Although the use of udev was attempted during the GDS processing, udev was not available due to an invalid environment.
This is a *details* environment issue.

Resolution

Collect investigation material and contact field engineers.

60360

ERROR: node : no enough address space

Explanation

Memory allocation failed on *node*.

Resolution

Identify the cause of the memory or swap space insufficiency on *node* and take corrective action, and then retry the relevant operation.

60361

ERROR: node : SWAP status disk is found

Explanation

A SWAPPED disk exists on *node*.

Resolution

Recover the SWAPPED disk on *node*.

60362

ERROR: node : class closed down is found

Explanation

A closed class exists on *node*.

Resolution

Recover the closed class on *node*.

For the recovery method, see "[D.1.4 Class Status Abnormality](#)."

60363

ERROR: you may not run this command while another command is running

Explanation

This command cannot be executed while other commands are running.

On a node in a cluster domain, the following commands cannot run together.

- clautoconfig -f
- sdxdisk -M
- sdxdisk -R
- sdxswap -O
- sdxswap -I
- sdxclass -R
- sdxconfig Restore
- sdxconfig Remove

Resolution

Execute the command again after the command is executed on another node.

When unmirroring the system disk of the GDS Management View on multiple nodes in the cluster system, and restarting the nodes at the same time, this message may be output on the system log of some nodes by the sdxclass command after the nodes are restarted. In this case, the object in the root class of the nodes are left without removed. Follow "[9.3 Removing Volumes, Groups, and Classes](#)" to remove the object in the root class.

60364

ERROR: node : failed to access an internal file *file*, errno=*errno*

Explanation

Accessing the file *file* on *node* failed.

Resolution

Collect investigation material and contact field engineers.

60365

ERROR: node : failed to obtain udev information from one of the devices

Explanation

A device whose udev information cannot be obtained exists on *node*.

Resolution

Collect investigation material and contact field engineers.

60366

ERROR: node : failed to update configuration database

Explanation

Updating the configuration database failed.

Resolution

The possible cause of this failure is that all of or most of the disks registered with the class are no longer accessible.

Identify the cause and recover the disks to gain normal access.

If this message is output, the class is closed immediately. After recovering the disks, also restart the class and then retry the relevant operation.

60367

ERROR: device : mklavel command failed

Explanation

Setting the *device* label failed.

The parted command is not installed, or failed.

Resolution

Install the parted command and re-execute it.

When an error occurs, collect investigation material and contact field engineers.

60373

ERROR: volume: related to proxy volume proxy with QuickOPC

Explanation

QuickOPC 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.

60374

ERROR: volume: related to master volume master with QuickOPC

Explanation

QuickOPC session exists between volume *volume* and master volume *master*.

Resolution

Use the `sdxproxy Cancel` command and cancel the QuickOPC 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.

60375

ERROR: device: too large for class of MSDOS labeled disks

Explanation

The capacity of a physical disk *device* is 2 TB or larger, so it cannot be registered with the class of MSDOS labeled disks.

Resolution

Register disks with the size smaller than 2 TB. Or, register disks to a new class or a class of GPT labeled disks. The disk label type of a class can be viewed in the LABEL field that is displayed with the `sdxinfo -C -e label` command.

60376

ERROR: volume: copying with QuickOPC

Explanation

Volume *volume* is in the process of QuickOPC copy.

Resolution

After the copying process is complete, execute again.

60377

ERROR: QuickOPC not available

Explanation

The QuickOPC function is unavailable. Check the following causes:

- a) Multiple proxy volumes are related to the master volume and the number of proxy volume to which QuickOPC sessions are set reaches the upper limit(15).
- b) Other causes.

Resolution

For a), take one of the following actions and perform the operation again.

- Cancel the QuickOPC session set to the operational master volume with the `sdxproxy Cancel volume`.
- Cancel the relation of the operational master volume and the proxy volume to which QuickOPC sessions are set with the `sdxproxy Break` command.

For b), see "[\(1\) The Advanced Copy function cannot be used in master-proxy copying.](#)" in "[D.1.8 Proxy Object Abnormality.](#)"

60378

ERROR: proxy: failed to start QuickOPC session, source=disk.volume, target=disk.volume, class=class

Explanation

An error occurred between proxy volume *proxy* and the master volume when QuickOP 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, ETERNUS Disk storage system and such, and promptly recover the disk.

60379

ERROR: patch levels are different between nodes in the class scope

Explanation

This message is output on one of the following:

a) The node where the setting of the parameter (SDX_EFI_DISK=on) is described in the GDS configuration parameter file /etc/opt/FJSVsdx/sdx.cf and the node not described exist together in the cluster system.

Conducting the following operation in this state will output this message.

- Registration of disk to the shared class.
- Changing the type from local class to the shared class.

b) Cannot install GDS I/O Monitor Option to the cluster system.

When setting the I/O Response Time Assurance Function for the shared class in this state, this message is output.

Resolution

Take the following action according to the cause, and then execute the operation again.

- Use the same parameter value of SDX_EFI_DISK described in the GDS configuration parameter file for all nodes in the cluster system.
- Install GDS I/O Monitor Option to all the nodes within the cluster system.

60380

ERROR: device: disk label is not matched with class class

disk label of class = class-label

disk label of device = device-label

Explanation

Restoration of *class* cannot be performed because the disk label type of a physical disk *device* (*device-label*) is different from that of a class *class* (*class-label*) described in the configuration table.

Resolution

If a wrong configuration file is specified, specify the correct configuration file name.

If a correct configuration file is specified, perform one of the following operations. When you perform one of the following operations, data in volume area on the physical disk registered with *class* is not ensured. For restoring volume data, restore the backup data after the following operation:

- Change the *device* disk label to *class-label* type with the parted(8) command and restore *class* again.
- Restore *class* with -e mklable option again.

60383

ERROR: device: too large for root class

Explanation

The capacity of a physical disk *device* is 1 TB or larger, so it cannot be registered with a root class.

Resolution

Register disks with the size smaller than 1 TB with a root class.

60384

ERROR: option1 and option2 cannot be specified at the same time

Explanation

The command options, *option1* and *option2*, cannot be specified simultaneously.

Resolution

See "[Appendix B Command Reference](#)" and specify the applicable option.

60385

ERROR: command is already running

Explanation

The command cannot be executed because the command *command* has already been run.

Resolution

Check that the process of *command* has stopped with the ps command, and then execute the command again.

60386

ERROR: file already exists

Explanation

The processing cannot be executed because the file *file* has already been existed.

Resolution

Either delete the file *file* or move it to the other directory.

60387

ERROR: device: failed to get multipath information on some node in physical scope

Explanation

The multipath information could not be obtained on any one of nodes where the physical disk *device* is connected.

Resolution

Check if the multipath software setting is correct on all nodes where the physical disk *device* is connected.

60388

ERROR: device: failed to get device information on some node in physical scope

Explanation

The device information could not be obtained on any one of nodes where the physical disk *device* is connected.

Resolution

Check if the physical disk *device* operates normally on all nodes where it is connected.

60389

ERROR: device: by-id file not found on some node in physical scope

Explanation

The by-id file cannot be found on any one of nodes where the physical disk *device* is connected.

Resolution

Check if the physical disk *device* operates normally on all nodes where it is connected.

60390

ERROR: device: failed to get multipath information, details

Explanation

The multipath information of the physical disk *device* failed to be obtained.

Resolution

Check if the multipath software is set properly.

60391

ERROR: device: failed to get device information, iscsiadm ret=ret

Explanation

The device information of the physical disk *device* failed to be obtained.

ret is the return value of the iscsiadm command.

Resolution

See the iscsiadm command manual to take corrective actions.

60392

ERROR: device: by-id file not found

Explanation

The by-id file of the physical disk *device* cannot be found.

Resolution

Check if the physical disk *device* works properly.

60393

ERROR: sfdsk driver is not starting

Explanation

The sfdsk driver is not started.

Resolution

Execute the command again after the sfdsk driver is started.

60394

ERROR: environment error, details

Explanation

The command cannot be executed due to an error in the environment.

Resolution

Collect investigation material and contact field engineers.

60396

ERROR: netmirror group cannot be created in this environment

Explanation

The mirroring among servers is available when one of the following products is used.

- PRIMECLUSTER Enterprise Edition
- PRIMECLUSTER HA Server
- PRIMECLUSTER Clustering Base and PRIMECLUSTER GDS

When attempting to create the netmirror group in the system where the above products are not properly installed, this message is output and the netmirror group cannot be created.

Resolution

Check that any product described in "Explanation" above is properly installed.

If this message is still output after any one of the above products is properly installed, collect investigation material and contact field engineers.

60397

ERROR: one disk in each of nodes needs to be connected to a netmirror group

Explanation

This message is output under the following conditions when the second disk is connected to the netmirror group.

- a. A disk is to be connected to the netmirror group where the other disk on the same node is already connected.
- b. A disk already connected to the netmirror group or a disk to be connected to the netmirror group is not the device that has been created in the procedure "[4.8 Disk Setting for Performing Mirroring among Servers](#)."

Resolution

For the condition a, change the netmirror group configuration so that the iSCSI device for the disk of any node in the class scope can be included.

For the condition b, connect the iSCSI device that is set up according to "[4.8 Disk Setting for Performing Mirroring among Servers](#)" to the netmirror group.

60398

```
ERROR: failed to initialize NVURM
```

Explanation

The log (NVURM) for the Just Resynchronization Mechanism (JRM) failed to be initialized.

Resolution

Check the I/O error message of NVURM. If the disk has a failure, replace the disk.

60399

```
ERROR: volume: file system uuid cannot be changed, xfs_admin(8) command failed, exit-status=exitstat
```

Explanation

The `xfs_admin(8)` command for the volume `volume` failed, and UUID of `volume` cannot be changed.

The exit status of the `xfs_admin(8)` 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 `xfs_admin(8)` command that was output right before the occurrence of the error and to the manual for the `xfs_admin(8)` command. If required, restore the data on `volume`, for example, using the backed up data, and try the `sdxproxy Root` command again.

60400

```
ERROR: cannot create netmirror group in root class
```

Explanation

The netmirror group cannot be created in the root class.

Resolution

Create the netmirror group in the shared class.

60401

```
ERROR: cannot connect three or more disks to netmirror group
```

Explanation

Three or more disks cannot be connected to the netmirror group.

Resolution

Connect maximum two disks to the netmirror group.

60402

ERROR: cannot create netmirror group in class including other type group, single disk, or spare disk

Explanation

The netmirror group cannot be created in a class where a group other than the netmirror group, the single disk, or the spare disk exists.

Resolution

Check if the specified class name, the group type attribute, and the object configuration in the class are correct.

If the specified class name or the group type attribute are incorrect, specify the correct class name and the correct group type attribute and then perform the operation again.

If the object configuration in the class is incorrect, correct the configuration.

60403

ERROR: cannot create netmirror group in class other than shared class

Explanation

The netmirror group cannot be created in a class other than the shared class.

Resolution

Create the netmirror group in the shared class.

60404

ERROR: cannot create netmirror group in class when three or more nodes exist in class scope

Explanation

The netmirror group cannot be created in the shared class where the class scope includes 3 or more nodes.

Resolution

Create the netmirror group in the shared class where the class scope includes 2 nodes.

60405

ERROR: cannot create group other than netmirror group in class including netmirror group

Explanation

A group other than the netmirror group cannot be created in a class where the netmirror group exists.

Resolution

Check if the specified class name, the group type attribute, and the object configuration in the class are correct.

If the specified class name or the group type attribute are incorrect, specify the correct class name and the correct group type attribute and then perform the operation again.

If the object configuration in the class is incorrect, correct the configuration.

60406

ERROR: cannot disconnect disk from netmirror group including volume

Explanation

A disk cannot be disconnected from the netmirror group where the volume exists.

Resolution

When disconnecting the disk from the netmirror group, delete the volume first.

60407

ERROR: cannot register disk other than undef disk with class including netmirror group

Explanation

In a class where the netmirror group exists, only the undef type disk can be registered.

Resolution

Check if the specified class name, the disk type attribute, and the object configuration in the class are correct.

If the specified class name or the disk type attribute are incorrect, specify the correct class name and the correct disk type attribute and then perform the operation again.

If the object configuration in the class is incorrect, correct the configuration.

60408

ERROR: cannot connect group to netmirror group

Explanation

A group cannot be connected to the netmirror group.

Resolution

Connect a disk to the netmirror group.

If the specified group name is incorrect, specify the correct group name, and then perform the operation again.

60409

ERROR: cannot connect netmirror group to other group

Explanation

The netmirror group cannot be connected to other groups.

Resolution

Specify the group other than the netmirror group, and then perform the operation again.

60410

ERROR: cannot set jrm to off for netmirror volume

Explanation

The JRM attribute of the netmirror volume cannot be set to off.

Resolution

Specify on for the JRM attribute when creating the netmirror volume.

60411

ERROR: cannot create volume in netmirror group including only one disk

Explanation

The volume cannot be created in the netmirror group where only one disk is connected.

Resolution

Before creating the volume, connect 2 disks to the netmirror group.

60412

ERROR: cannot detach slice temporarily from netmirror volume

Explanation

A slice of the netmirror volume cannot be temporarily detached.

Resolution

If the specified volume name is incorrect, specify the correct volume name, and then perform the operation again.

60413

ERROR: cannot set three or more nodes to scope of class including netmirror group

Explanation

The class scope where the netmirror group exists cannot be changed to 3 or more nodes.

Resolution

If the specified class name is incorrect, specify the correct class name, and then perform the operation again.

60414

ERROR: cannot change attribute value of disk in class including netmirror group

Explanation

If the netmirror group exists in a class, the disk type attribute in the class cannot be changed.

Resolution

If the specified class name is incorrect, specify the correct class name, and then perform the operation again.

60416

ERROR: cannot perform proxy operation for netmirror object

Explanation

Join and Relate operations cannot be performed to the netmirror volume or the netmirror group.

Resolution

If the specified object is incorrect, specify the correct object, and then perform the operation again.

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)

60419

ERROR: cannot change jrm to off for netmirror volume

Explanation

The JRM attribute of the netmirror volume cannot be changed to off.

Resolution

If the specified volume name is incorrect, specify the correct volume name, and then perform the operation again.

60420

ERROR: cannot find local device

Explanation

The disk without the iSCSI device setting is connected to the netmirror group.

Resolution

Follow "[4.8 Disk Setting for Performing Mirroring among Servers](#)" to set the iSCSI device, and then, execute the operation again.

60422

ERROR: file: failed to delete file

Explanation

The `/var/opt/FJSVsdX/log/file` file could not be deleted.

Resolution

Delete the `/var/opt/FJSVsdX/log/file` file if it exists. For the cluster configuration, check if the `/var/opt/FJSVsdX/log/file` file exists on all the nodes that configure the cluster.

If the file does not exist or this message is still output after deleting the file, collect investigation material, and then contact field engineers.

60423

ERROR: cannot register *number* or more disks to class including netmirror group

Explanation

Cannot register *number* or more disks to the class where the netmirror group exists.

Resolution

Register up to *number* or less disks to the class where the netmirror group exists.

60424

ERROR: cannot create netmirror group in class including *number* disks

Explanation

If *number* or more disks are registered in the class, the netmirror group cannot be created.

Resolution

Register up to *number* or less disks to the class where the netmirror group is to be created.

60425

ERROR: cannot operate -e force option in class not including netmirror group

Explanation

The -e force option can only be used when specifying the disk connected to the netmirror group.

Resolution

Use the -e force option only when specifying the disk connected to the netmirror group.

60426

ERROR: cannot open /etc/fstab

Explanation

The /etc/fstab file cannot be opened, and the setting of the system disk mirroring cannot be confirmed. In this case, the parameter cannot be set.

Resolution

The /etc/fstab file may have an error. Check that the /etc/fstab file is normal, and then execute the command again.

60427

ERROR: system disk is already mirrored

Explanation

The parameter cannot be set because the system disk mirroring is already configured.

Resolution

Cancel the system disk mirroring once, and then execute the command again.

60428

ERROR: cannot find root filesystem information in /etc/fstab

Explanation

The information of the root file system is not described in the `/etc/fstab` file, and the configuration of the system disk mirroring cannot be confirmed. In this case, the parameter cannot be set.

Resolution

The `/etc/fstab` file may have an error. Check that the `/etc/fstab` file is normal, and then execute the command again.

60429

ERROR: device: failed to update disk partition (errno=errno)

Explanation

The process to update a disk partition failed.

Resolution

When `errno=16`, make sure that the disk *device* has no error. After that, perform the operation again.

When `errno` is other than 16, contact field engineers.

60430

ERROR: device: failed to delete a partition from the disk (errno=errno)

Explanation

The process to delete a partition from the disk failed.

Resolution

When `errno=16`, make sure that the disk *device* has no error. After that, perform the operation again.

When `errno` is other than 16, contact field engineers.

60431

ERROR: device: failed to setup netmirror device

Explanation

The settings related to the mirroring among servers of the physical disk *device* failed.

Resolution

Collect investigation material and contact field engineers.

60432

ERROR: cannot swap disk connected to netmirror group in local class

Explanation

The disk connected to the netmirror group in the local class can not be swapped.

Resolution

Change the class to the shared class, and then perform the operation again.

60433

ERROR: *command* command failed

Explanation

The command *command* cannot be executed.

Resolution

Collect investigation material, and contact field engineers.

60434

ERROR: /boot file system not mounted

Explanation

/boot file system is not mounted.

Resolution

Mount /boot file system first, and then execute the command again.

60435

ERROR: cannot open /etc/opt/FJSVsdX/modules/sfdsk.conf

Explanation

/etc/opt/FJSVsdX/modules/sfdsk.conf cannot be opened.

Resolution

Collect investigation material, and contact field engineers.

60436

ERROR: failed to create temporary file

Explanation

The temporary file cannot be created.

The possible causes are as follows:

- a) /var file system is not mounted.
- b) Insufficient capacity in /var/opt.
- c) The file system is corrupted or the disk fails.

Resolution

Take the following steps, and then execute the command again:

- 1) Mount /var file system if it is not mounted.
- 2) If storage capacity in /var/opt is less than 100MB, delete unnecessary files to secure 100MB or more storage capacity in /var/opt. After that, restart the system.

If the cause of the error cannot be specified or the error occurs again after the command is executed, collect investigation material, and contact field engineers.

60437

ERROR: cannot get partition number of file system including /boot

Explanation

The partition number of the device of the file system including /boot cannot be obtained.

Resolution

Take the following steps:

1. Execute the following command to update the partition number information managed by GDS.

```
# /etc/opt/FJSVsdx/bin/sdxbootpart -S
```

2. Execute the command that causes this error.

Execute sdxbootpart command only in the steps described in this manual or other GDS documents.

60438

ERROR: failed to update /boot/efi/EFI/redhat/grub.cfg

Explanation

/boot/efi/EFI/redhat/grub.cfg file cannot be updated.

Resolution

Collect investigation material, and contact field engineers.

60439

ERROR: class: closing down class on node node, failed to update netmirror device information

Explanation

The class *class* has closed in the node *node* due to the device information of the netmirror group has failed to be updated.

Resolution

Take the following steps, and then execute the command again:

1. Restore the file or file system.

- When no available capacity under /var/opt

Delete unnecessary files to secure 100MB or more free capacities under /var/opt. After that, restart the system.

- When enough capacity available under /var/opt

Output the /var/opt/FJSVsdx/log/.sdxnetmirror_disable.db file by the cat command, and check that the output is not garbled.

If it is garbled, delete the file and then restart the system.

If it is not garbled, collect investigation material and contact field engineers.

2. Check that the class is not closed on the node which has not been restarted.

Execute the following command on the node which has not been restarted. If "yes" is displayed in the DOWN field, the class is closed.

```
# /etc/opt/FJSVsdx/bin/sdxcdown
```

3. Restore the class if it is closed in step 2.

Execute the following command on the node which has not been restarted.

```
# sdxfix -C -c class_name
```

60440

ERROR: class: failed to update netmirror device information

Explanation

The command has failed due to the device information of the netmirror group of class *class* has failed to be updated.

Resolution

Take the following steps, and then execute the command again:

- When no available capacity under /var/opt

Delete unnecessary files to secure 100MB or more free capacities under /var/opt. After that, restart the system.

- When enough capacity available under /var/opt

Output the /var/opt/FJSVsdx/log/.sdxnetmirror_disable.db file by the cat command, and check that the output is not garbled.

If it is garbled, delete the file and then restart the system.

If it is not garbled, collect investigation material and contact field engineers.

60441

ERROR: I/O monitor option is not installed

Explanation

GDS I/O Monitor Option was not installed correctly.

Resolution

Install the GDS I/O Monitor Option correctly when using the sdsdisk -M command or the sdxattr -C command to set up the I/O Response Time Assurance, or using the sdxinfo -I command to check the information of the I/O Times Response Assurance.

60442

ERROR: class: cannot change iotimeout attribute value to on, concat or stripe group exist

Explanation

The iotimeout attribute of the class cannot be changed to on because concat or stripe type groups exist in the class *class*.

Resolution

Check if there is no mistake in the specified class name or the object configuration in the class *class*.

If there is a mistake in the specified class name, specify the correct class name and execute again.

If there is a mistake in the object configuration in the class, change it to the correct configuration.

60443

ERROR: class: cannot create concat or stripe group, iotimeout attribute value is on

Explanation

Concat and stripe type groups cannot be created because the iotimeout attribute in class *class* is on.

Resolution

Check if there is no mistake in the specified class name or the iotimeout attribute value.

If there is a mistake in the specified class name, specify the correct class name and execute again.

If there is a mistake in the iotimeout attribute value of the class, use the sdxattr command to change the iotimeout attribute to off.

60444

ERROR: only iotimeout attribute can be specified

Explanation

Only the iotimeout attribute can be specified in the root class.

Resolution

Check if there is no mistake in the specified class name or the object configuration in the class *class*.
If there is a mistake in the specified class name, specify the correct class name and execute again.
If there is a mistake in the specified attribute, specify the correct attribute and execute again.

60445

ERROR: class: cannot change iotimeout attribute value to on, netmirror group exist

Explanation

The iotimeout attribute of the class cannot be changed to on because netmirror type of group exists in the class *class*.

Resolution

Check if there is no mistake in the specified class name or the object configuration in the class *class*.
If there is a mistake in the specified class name, specify the correct class name and execute again.
If there is a mistake in the object configuration in the class, change it to the correct configuration.

60446

ERROR: class: cannot create netmirror group, iotimeout attribute value is on

Explanation

Netmirror type group cannot be created because the iotimeout attribute in the class *class* is on.

Resolution

Check if there is no mistake in the specified class name or the iotimeout attribute value in the class *class*.
If there is a mistake in the specified class name, specify the correct class name and execute again.
If there is a mistake in the iotimeout attribute value of the class, use the `sdxattr` command to change the iotimeout attribute to off.

60447

ERROR: netmirror DB file is invalid, volume=volume

Explanation

For the disks that make up the volume *volume*, the disk containing the latest data cannot be determined because there is an inconsistency in information to determine the disk containing the latest data.

Resolution

See "[7.16.8 Restoration when the Latest Disk Cannot Be Selected Automatically.](#)"

60448

ERROR: object.volume is only not pre-detached slice

Explanation

The slice *object.volume* is the only slice that is not pre-detached in the netmirror volume *volume*.
Operation cannot be performed because continuing operation may result in data loss of *volume*.

Resolution

Wait until the other slice in volume is detached and it becomes INVALID.
After that, if the INVALID slice is restored, operation of *object.volume* can be performed.
When restoring the INVALID slice, see "[\(1\) The netmirror slice that configures the netmirror volume is in INVALID state.](#)" in "[D.1.6 Errors in the Mirroring Among Servers.](#)"

60449

ERROR: volume: netmirror volume must be specified

Explanation

The netmirror volume must be specified.

Resolution

Specify the netmirror volume.

60450

ERROR: failed to update netmirror DB file, volume=volume

Explanation

Update of netmirror device information has failed.

Resolution

- If the INVALID status slice exists within the netmirror volume *volume*, and other nodes are LEFTCLUSTER or DOWN.

Synchronization copying of netmirror volume fails.

Restore the node first. After that, execute the `sdxcopy -B` command, for example, to perform the synchronization copying.

For how to recover from the LEFTCLUSTER state, see "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide."

- Otherwise

Take the following procedure to restore.

1. Restore the file system or file with all of the nodes.

- When no available capacity under `/var/opt`

Delete unnecessary files to secure 100MB or more free capacities under `/var/opt`. After that, restart the system.

- When enough capacity available under `/var/opt`

Output the `/var/opt/FJSVsdx/log/.sdxnetmirror_disable.db` file by the `cat` command, and check that the output is not garbled.

If it is garbled, delete the file and then restart the system.

If `/var/opt` on all of the nodes has capacity and the file is not garbled, collect investigation material and contact field engineers.

2. Depending on the slice status within the netmirror volume *volume*, take the following procedure.

- When the NOUSE status slice does not exist in the netmirror volume *volume*.

Synchronization copying of netmirror volume fails.

Perform the synchronization copying by using the `sdxcopy -B` command, for example.

- When the NOUSE status slice exists in the netmirror volume *volume*.

Restore NOUSE status of the slice. For details, see "[\(6\) Slice configuring the volume is in NOUSE status.](#)" in "[D.1.1 Slice Status Abnormality.](#)"

60451

ERROR: node: physical disk read error

Explanation

A read error occurred on a physical disk 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.

60452

`ERROR: no valid slice left in sfdsk driver on node node`

Explanation

Operation cannot be continued because no valid slices are left in the sfdsk driver of the node *node*.

Resolution

After verifying that the status of the network for mirroring among servers is normal, restart the node *node* and perform the operation again.

60453

`ERROR: group: type group not supported`

Explanation

The class object configuration cannot be restored in RHEL8 because a *type* type group *group* exists in the configuration table.

Resolution

Remove concatenation groups and stripe groups from the configuration table with the `sdxconfig Convert` command.

For details on the removal method, see "[Appendix B Command Reference](#)."

C.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 rsyslogd 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 B Command Reference](#)."

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/fstab not updated, unsupported file system type fstype

Explanation

Entry in the `/etc/fstab` 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/fstab` file or change the file system configuration according to need.

62018

WARNING: correct `/etc/fstab` and file system configuration before rebooting, or system may not be booted

Explanation

The `/etc/fstab` file contains an invalid entry. When an entry for a file system, such as `/` (root), `/usr`, and `/var`, essential to system start-up is invalid, it will be impossible to start the system if rebooted.

Resolution

Refer to a GDS's WARNING message output just before this message and pinpoint the invalid entry in the `/etc/fstab` file. Be sure to repair the `/etc/fstab` file and the file system configuration as needed prior to rebooting the system.

62019

WARNING: volume: entry in `/etc/fstab` not updated, unsupported file system type *fstype*

Explanation

Since file system type *fstype* is not supported, entry in `/etc/fstab` file for *volume* could not be updated.

Resolution

Directly edit `/etc/fstab` 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.

62025

WARNING: physical disk not found

Explanation

After the system is restarted, some physical disks cannot be recognized from OS.

When the physical disk cannot be recognized, the asterisk (*) is displayed in DEVNAM field of the `sdxinfo -D` command output result.

Resolution

The unrecognized physical disk may have a failure. Check if the disk has a failure. Restore the failed disk by replacing it with the new disk, for example, and then execute the command again.

62026

WARNING: *device*: failed to delete partition table

Explanation

The partition table failed to be deleted.

Resolution

The physical disk *device* may have a failure.

Check if the disk has a failure or not.

If *device* has a failure, restore the disk by replacing it with another disk, for example.

If *device* is normal, delete *device* from the class before restarting the system. After that, delete all the partitions of *device* by using a command such as the `parted(8)` command.

C.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/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/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/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/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, shut 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 canceled.

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/fstab* file

Explanation

The */etc/fstab* file was updated.

64045

INFO: updated /etc/opt/FJSVsdx/modules/sfdsk.conf file

Explanation

The /etc/opt/FJSVsdx/modules/sfdsk.conf file was updated.

64047

INFO: updated *filepath* file

Explanation

The file was updated. *filepath* is the path name of the updated file.

64048

INFO: completed definitions of system disk(s) mirroring

Explanation

Definition for system disk mirroring is complete.

64049

INFO: system must be rebooted:

```
shutdown -r now
```

Explanation

Reboot the system by executing the following command.

```
shutdown -r now
```

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/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/dsk/disk.volume

Explanation

The current *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 canceled.

64060

INFO: *proxy*: forced to break relation with master

Explanation

Relationship between master and proxy *proxy* was forcibly canceled.

64061

INFO: *proxy*: must be recovered by `sdxfix(8)` command

Explanation

In order to restore data of proxy *proxy*, volume must be restored with `sdxfix(8)` command.

64062

INFO: *volume*: started volume on node *node*,...

`/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(8) command:

details

Explanation

The file system on the alternative root volume *volume* will be inspected and repaired with the fsck(8) command. *details* is a message of the fsck(8) command.

64076

INFO: updated /etc/fstab file on alternative root volume *volume*

Explanation

The `/etc/fstab` 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.

64079

`INFO: updated filepath file on alternative root volume volume`

Explanation

The file on the alternative root volume *volume* was updated.
filepath is the path name of the updated file.

64081

`INFO: updated boot devices:`

`old=old-boot-device ...`

`new=new-boot-device ...`

Explanation

The boot devices changed from *old-boot-device ...* to *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, reverse the boot order using the Change Boot Order option in the EFI boot manager's Boot option maintenance menu.

64082

`INFO: completed definitions of switching boot environment`

Explanation

The boot environment switching is now defined.

64084

`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* ... Reverse the boot order using the Change Boot Order menu in the EFI boot manager's Boot option maintenance menu to switch to the alternative boot environment.

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*.

64096

INFO: device: creating by-id file

Explanation

The by-id file of the physical disk *device* is created in the `/dev/disk/by-id` directory.

C.4.4 Fix Messages (66000)

66000

TO FIX: Usage: usage

Resolution

Follow the *usage* syntax and execute the command.

C.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.

a) The following message 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.

```
sdxcopy: ERROR: SDX_COPY_STATUS_WAIT2: internal error, sdxerrno=145
```

b) If the `sdxvolume -M` command is executed when I/O to the physical disk is disabled due to a failure or other causes, the following message may be output.

```
sdxvolume: ERROR: internal error (mv_create,sdxerrno=0)
```

c) If the `sdxdisk -M` command is executed when I/O to the physical disk is disabled due to a failure or other causes, the following message may be output.

```
sdxdisk: ERROR: internal error (td_regist,sdxerrno=4207)
```

d) If the `sdxdisk -R` command, the `sdxconfig Remove` command, or the `sdxclass -R` command is executed when I/O to the physical disk is disabled due to a failure or other causes, the following message may be output.

```
command: ERROR: internal error (td_delete,sdxerrno=4120)
```

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/FJSVsdx/msglog/sdxservd.log`.

```
SDX:sdxservd: INFO: command exited, exit-status=exitstat:  
sdxcopy -C -c class -v volume,...
```

In the event of b) or c), check the physical disk status and the I/O route status. If an error exists, restore the error and then execute the command again.

In the event of d), the class is closed. Follow "(1) Class becomes closed status during operation." in "D.1.4 Class Status Abnormality."

In other events, collect investigation material and contact field engineers.

C.5 Operation Management View Messages

This section describes the GDS Management View messages.

C.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 field engineers.

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 field engineers.

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 field engineers.

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 "[2.3.1 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 "[2.3.1 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 "[2.3.1 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 "[2.3.1 Object Name.](#)"

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

The system disk's slice numbers are all in use. Remove at least one slice before setting the system disk.

Explanation

The number of slices on the system disk has reached the upper limit (15) and system disk setting cannot be performed.

Resolution

Remove one or more system disk's slices, for example, with the parted (8) 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 field engineers.

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 field engineers.

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 field engineers.

C.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 field engineers.

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.

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.

Explanation

The volume cannot be removed because of a file system configured on the volume.

Resolution

Remove the description of the corresponding volume from the /etc/fstab file, 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

An attempt was made to change the volume name for the volume mounted or described in the /etc/fstab file.

Resolution

Change the volume name after performing the following resolutions:

- 1) If the file system created on the volume is mounted, unmount it.
- 2) If the state of the volume is active, stop the volume in all of the nodes within the class scope.
- 3) If the volume name is described in the entry of /etc/fstab file, change the description of volume name or delete the entry.

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.

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.

5148

The physical disks will be registered with a class, but some of them are not in the cluster resource database. These disks can be registered with a local class but not with a shared class. Continue?

Explanation

The disk resources are not included in the cluster resource database or physical disk information was not updated after the disk resource registration. These disks can be registered with a local class but not with a shared class.

Resolution

- a. To register the disks with a local class, click <Yes>.
- b. To register the disks with a shared class, click <No> to close the message, and click <Exit> to close the Class Configuration screen once.
 - b1) If the disk resource registration is done, select [Operation]:[Update Physical Disk Information] in the main screen and then create the class.
 - b2) If the disk resource registration is not done, perform resource registration first, and select [Operation]:[Update Physical Disk Information] in the main screen and then create the class.

For details on resource registration, see the "[6.3 Shared Disk Resource Registration](#)."

C.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/dsk/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 field engineers 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.

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/fstab
filepath

Explanation

The system disk setting process completed mirroring definition. *filepath* is the path name of the updated file.

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 "[2.3.12 Proxy Configuration Preconditions](#)" and the description of the Join option in "[B.2.1 sdxcproxy - 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 "[2.3.12 Proxy Configuration Preconditions](#)" and the description of the Join option in "[B.2.1 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 "[2.3.12 Proxy Configuration Preconditions](#)" and the description of the Relate option in "[B.2.1 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 "[2.3.12 Proxy Configuration Preconditions](#)" and the description of the Relate option in "[B.2.1 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.

Appendix D 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.

D.1 Resolving Problems

D.1.1 Slice Status Abnormality

If the slice status is one of the following statuses, take the actions 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 canceled as the result of a cause such as [Cancel Copying] selection from the GDS Management View, sdxcopy command execution, or a power outage.
- An I/O error occurred due to a SCSI timeout. Whether a SCSI timeout has occurred can be determined from the following message:

```
kernel: mptscsih: iocn: attempting task abort! (sc=e00001401a0dce80)
```

Resolution

- 1) Identify the physical disk name of a faulty disk using the sdxinfo command.

(Example A1)

```
# sdxinfo -G -o Volume1
OBJ  NAME  CLASS  DISKS  BLKS  FREEBLKS  SPARE
-----
group Group1 Class1 Object1:Object2  17596416  17498112  0

# sdxinfo -D -o Volume1 -e long
OBJ  NAME  TYPE  CLASS  GROUP  DEVNAM  DEVBLKS  FREEBLKS  DEVCONNECT  STATUS  E
-----
disk  Object1 mirror Class1  Group1  sda      17596416          * node1:node2
ENABLE *
disk  Object2 mirror Class1  Group1  sdb      17596416          * node1:node2
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 sdb 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)

```

# sdxinfo -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 sda     17596416      * node1:node2
ENABLE  0
disk     Disk2     stripe Class1  Object1 sdb     17596416      * node1:node2
ENABLE  0
disk     Disk3     stripe Class1  Object2 sdc     17682084      * node1:node2
ENABLE  1
disk     Disk4     stripe Class1  Object2 sdd     17682084      * node1:node2
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 sdc 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 field engineers 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 "7.3.1.2 Operation Procedure."

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 field engineers 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 [EFI]

B. For the swap area [EFI]

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  sda     17596416      * node1:node2
ENABLE 1
disk   Disk2   mirror Class1  Group1  sdb     17596416      * node1:node2
ENABLE 0
```

A. For /(root), /usr, or /var [EFI]

Follow "(4) System cannot be booted. (Failure of all boot disks)" in "D.1.5 System Disk Abnormality [EFI]" for restoration. In the procedure of "Resolution", replace only the faulty copy source disk. Do not replace all the disks that are registered in the root class.

B. For the swap area [EFI]

B.1) Check the swap volume.

```
# swapon -s
Filename                Type      Size    Used    Priority
/dev/sfdsk/gdssys32    partition 4194296 0        -1
```

B.2) Remove the volume from the swap area.

```
# swapoff /dev/sfdsk/gdssys32
```

Depending on the part or severity of failure in disks that constitute the volume, the swapoff(8) 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.1.2).

B.2.1) Comment out the swap line to prevent use of the volume as a swap area after the system is rebooted.

```
# vim /etc/fstab

Before edit:
/dev/sfdsk/gdssys32  swap                swap  defaults  0 0

After edit:
#/dev/sfdsk/gdssys32  swap                swap  defaults  0 0
```

B.2.2) Reboot the system.

```
# shutdown -r now
```

B.3) Stop the volume.

```
# sdxvolume -F -c Class1 -v Volume1
```

B.4) Restore the status of the copy destination slice in INVALID status.

```
# sdxfix -v -c Class1 -d Disk2 -v Volume1
```

B.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
```

B.6) Start the volume.

```
# sdxvolume -N -c Class1 -v Volume1
```

B.7) Add the volume to the swap area again.

```
# swapon /dev/sfdsk/gdssys32
```

When step B.2.1) was performed, undo the edit that was made in the /etc/fstab file.

```
# vim /etc/fstab

Before edit:
#/dev/sfdsk/gdssys32  swap                swap  defaults  0 0

After edit:
/dev/sfdsk/gdssys32  swap                swap  defaults  0 0
```

B.8) Remove the faulty copy source disk from GDS management to give it a replaceable status.

```
# sdxswap -O -c Class1 -d Disk1
```

B.9) Swap the faulty copy source disk.

B.10) 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.

In this example, the volume has been used as an ext4 file system.

```
# umount /dev/sfdsk/Class1/dsk/Volume1
```

Depending on the part or the severity of failure in disks that compose the volume, the umount command may fail due to an I/O error. In this event, unmount the file system performing steps C.2.1) through C.2.3).

C.2.1) If class Class1 is registered with a cluster application, remove the cluster application.

C.2.2) Comment out the /dev/sfdsk/Class1/dsk/Volume1 line in the /etc/fstab file to prevent mounting of the volume after the system is rebooted.

C.2.3) Reboot the system.

C.3) Stop the volume.

```
# sdxvolume -F -c Class1 -v Volume1
```

C.4) Restore the status of the copy destination slice in the INVALID status.

```
# sdxfix -V -c Class1 -d Disk2 -v Volume1
```

C.5) Verify that the restored copy destination slice is in the STOP status and the copy source slice is the 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 lost. Restore the backup data or perform repair using the fsck(8) command if necessary.



If an I/O error occurred in the copy source slice during just resynchronization after the system went down, restoration may possibly be performed with the fsck(8) command.

When step C.2.2) was performed, undo the edit that was made in the /etc/fstab file.

When step C.2.1) was performed, re-create the cluster application removed in step C.2.1).

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 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 "[D.1.2 Disk Status Abnormality](#)."

D.1.2 Disk Status Abnormality

If the disk status is one of the following statuses, take the actions 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 "[B.1.8 sdxswap - Swap Disk](#)" and "[7.3.1.2 Operation Procedure](#)."
- c) If the possible cause is (Cause c), follow the procedures and swap the disks.

The following action is necessary.

- When the disk is a single disk
Delete the single volume in INVALID state on a single disk, and then swap the disk.
- When the disk is connected to the highest level group
Delete the volume in INVALID state in the highest level group, and then swap the disk.
- When the disk is connected to a lower level group
When swapping the disk, delete the INVALID volume first, and then disconnect the lower level group from the higher level group.

(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.

Note

When an I/O error has occurred in the root class, the E field of the disk information shown by the `sdxinfo -e` long command will not display "1" which indicates the I/O error status. Check the I/O error information using system log etc.

Resolution

Check the statuses of volumes and slices relative to the disk with an I/O error, and perform restoration, referring to "[D.1.3 Volume Status Abnormality](#)" or "[D.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
```

D.1.3 Volume Status Abnormality

If the volume status is one of the following statuses, take the actions 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 canceled 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  sda     17596416  node1:node2  ENABLE
disk   Disk2   mirror  Class1  Group1  sdb     17596416  node1:node2  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  sda     17596416  node1:node2  ENABLE
disk   Disk2   stripe  Class1  Group2  sdb     17596416  node1:node2  DISABLE
disk   Disk3   stripe  Class1  Group3  sdc     17682084  node1:node2  ENABLE
disk   Disk4   stripe  Class1  Group3  sdd     17682084  node1:node2  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 "[D.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 canceled 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  *      sda     8355840  node1:node2  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 "D.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  ACTDISK
-----
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   sda       17596416  node1:node2   ENABLE
disk     Disk2     concat  Class1   Group2   sdb       17596416  node1:node2   DISABLE
disk     Disk3     concat  Class1   Group3   sdc       17682084  node1:node2   ENABLE
disk     Disk4     concat  Class1   Group3   sdd       17682084  node1:node2   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 "[D.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    sda        8421376   *           ENABLE
disk     Disk2     mirror    Class1    Group1    sdb        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 "[D.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 "[D.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 `umount`. When using the proxy volume as a file system, execute `umount`.

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 "[B.1.8 sdxswap - Swap Disk](#)" and "[7.3.1.2 Operation Procedure](#)."

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 "[D.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 unmount.

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 "D.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 sda     8421376  *          ENABLE
disk  Disk2  mirror Class1  Group1 sdb     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 "D.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 "D.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 "[B.1.8 sdxswap - Swap Disk](#)," or section "[7.3.1.2 Operation Procedure](#)."

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 "[B.1.8 sdxswap - Swap Disk](#)," or section "[7.3.1.2 Operation Procedure](#)."

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 message "ERROR: *disk: volume* in INVALID status" will be output. 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 unmount.

d3) Stop all volumes in the group.

```
# sdxvolume -F -c Class1 -v Volume2
```

Note

When the volume is in INVALID status, the message "ERROR: *disk: volume* in INVALID status" will be output. 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 "[B.1.8 sdxswap - Swap Disk](#)," or section "[7.3.1.2 Operation Procedure](#)."

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 "D.1.10 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 `umount` command. When I/O error occurs on the `umount` command, execute `-f` option of the `umount` 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 "[D.1.1 Slice Status Abnormality](#)."
- a5) If there is a NOUSE status slice within the volume, attempt recovery following the procedures given in "[D.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 "[7.3.1.2 Operation Procedure](#)" and "[B.1.8 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 "[D.1.1 Slice Status Abnormality](#)."
- b5) If there is a NOUSE status slice within the volume, attempt recovery following the procedures given in "[D.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 "7.3.1.2 Operation Procedure" and "B.1.8 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 (sda) has a failure.

```
# sdxinfo -D -o Disk1
OBJ      NAME      TYPE      CLASS     GROUP     DEVNAM     DEVBLKS     DEVCONNECT     STATUS
-----
disk     Disk1     single   Class1    *         sda        8493876     node1:node2     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.

Note

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 sda=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, recover the failed component, and then reboot the system. Slice data is valid and there is no need to restore the data.

However, as the E field of the disk information shown by the sdxinfo -e long command will display "1" due to an I/O error, run the sdxfix -D command to cancel the I/O error status.

```
# sdxfix -D -c Class_Name -d Disk_Name -e online -x NoRdchk
```

(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  DEVMAM  DEVBLKS  FREEBLKS  DEVCONNECT  STATUS  E
-----
disk  Disk1  concat  Class1  Group2  sda     17596416          *  node1:node2  ENABLE  0
disk  Disk2  concat  Class1  Group2  sdb     17596416          *  node1:node2  ENABLE  1
disk  Disk3  concat  Class1  Group3  sdc     17682084          *  node1:node2  ENABLE  0
disk  Disk4  concat  Class1  Group3  sdd     17682084          *  node1:node2  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 sdb, as shown in the DEVMAM field.

Note

When an I/O error has occurred in the root class, the E field of the disk information shown by the sdxinfo -e long command will not display "1" which indicates the I/O error status. Check the I/O error information using system log etc.

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  ACTDISK
-----
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 unmount command. When I/O error occurs on unmount command, execute -f option of unmount 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 sdb=Disk2
```


a1) 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
```

a2) 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
```

a3) 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, recover the failed component, and then reboot the system. Slice data is valid and there is no need to restore the data.

However, as the E field of the disk information shown by the sdxinfo -e long command will display "1" due to an I/O error, run the sdxfix -D command to cancel the I/O error status.

```
# sdxfix -D -c Class_Name -d Disk_Name -e online -x NoRdchk
```

D.1.4 Class Status Abnormality

If the class status is one of the following statuses, take the actions as indicated for the relevant situation.

- (1) Class becomes closed status during operation.
- (2) Class cannot be started when booting the system.

(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 Dell EMC storage unit 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 sda:sdb:sdc:sdd:sde:sdf:sdg:sdh
Class2 yes Comm 10 10 8 0 sdi:sdj:sdk:sdl:sdm:sdn:sdo:sdp
Class3 yes FewDB 10 10 1 7 sdq
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 "[D.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 field engineers.

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    sda        8847360     *              ENABLE
disk     Disk32    mirror    Class3    Group1    sdb        8847360     *              ENABLE
disk     Disk33    mirror    Class3    Group2    sde        8847360     *              ENABLE
disk     Disk34    mirror    Class3    Group2    sdf        8847360     *              ENABLE
disk     Disk35    mirror    Class3    Group3    sdc        17793024    *              ENABLE
disk     Disk36    mirror    Class3    Group3    sdg        17793024    *              ENABLE
disk     Disk37    mirror    Class3    Group4    sdd        17793024    *              ENABLE
disk     Disk38    mirror    Class3    Group4    sdh        17793024    *              ENABLE
disk     Disk39    spare     Class3    Group1    sdr        17793024    *              ENABLE
disk     Disk40    spare     Class3    *         sds        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.

```
# vim /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 field engineers. 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 "[7.3.1.2 Operation Procedure](#)," or "[B.1.8 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/sdxcdown
CLASS   DOWN REASON NDK  NEN  NDB  NLDB  DEVNAM
-----
Class1  no   -       10  10   8    0  sda:sdb:sdc:sdd:sde:sdf:sdg:sdh
Class2  no   -       10  10   8    0  sdi:sdj:sdk:sdl:sdm:sdn:sdo:sdp
Class3  no   -       10  10   8    0  sdq:sdr:sds:sdv:sdu:sdv:sdw:sdx
Class4  no   -       10  10   8    0  sdaa:sdab:sdac:sdad:sdae:sdaf:sdag:sdah
```

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.

```
# vim /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 field engineers.

(2) Class cannot be started when booting the system.

Explanation

When booting the system, if the configuration database that contains the configurations and status of objects in the class is not accessible due to an I/O error of the disk or similar causes, the class remains unstarted.

All objects in the class that remains unstarted are not accessible. Also, the class which remains unstarted will not display related objects or any other class information even running the sdxinfo command.

Resolution

- 1) Identify the cause by referring to the sfdisk and disk driver message output on the console, and recover the class.
- 2) See (Cause a) and (Cause b) in section "(1) Disk is in DISABLE status." in "D.1.2 Disk Status Abnormality", and take the corresponding resolution measures.



Shutting down the node

If the class is registered in the cluster application, perform the following procedure before shutting down the node.

If the following procedure is not performed, the Offline process of userApplication may not be completed and the shutdown process may be waited.

1. Check the status of userApplication on each node.

```
# hvdisp -a
```

If there is no userApplication of which the status is "Unknown", step 2. and subsequent steps do not need to be performed.

Shut down the node as usual.

2. Stop RMS and then shut down the node.

The procedure differs depending on the class status.

a) If unstarted classes exist in all the nodes

a-1) Stop userApplication of which the status is not "Unknown" on all the node.

```
# hvutil -f userApplication_name
```

a-2) Stop RMS on all the nodes.

Execute the following command on all the nodes.

```
# hvshut -L
```

Enter "yes" for the warning message displayed during execution.

a-3) Shut down all the nodes.

b) If there is a node in which all the classes are started

b-1) If there is a userApplication of which the status is not "Unknown" on the node where no class is started, switch userApplication to the node where the class is started.

```
# hvswitch userApplication_name SysNode_name_where_the_class_is_started
```

b-2) Stop RMS.

Execute the following command on the node where no class is started.

```
# hvshut -L
```

Enter "yes" for the warning message displayed during execution.

b-3) Shut down the node where no class is started.

For details of each command, see the manual of each command.

In any of the following situations, contact field engineers, as the corresponding class must be forcibly removed and then recreated.

- If you want to use a disk other than the normal hard disk (sdX):
Add the following text to the /etc/opt/FJSVsdx/sdx.cf file on all nodes.

```
SDX_DEVLABEL_USE=off
```

- If you have mistakenly formatted a disk that is registered with GDS when operating an ETERNUS Disk storage system or other server case.

D.1.5 System Disk Abnormality [EFI]

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 the actions 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. (Failure of all boot disks)
- (5) System cannot be booted. (Boot disk data damage)
- (6) System cannot be booted. (Removing errata kernel)

(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 "7.4.2 Backing Up [EFI]."

(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 "[7.4.5 Restoring \(When System Can Be Booted\) \[EFI\]](#)."

(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) From boot disk units displayed in the EFI boot manager's boot option selection screen, select a mirror disk and boot.

Among the SDX disk names, the top one is the original boot disk unit name and the others are the mirror disk names. When mirroring is configured with three-way or more multiplexing, multiple mirror disks exist. In the following example, Root1 is the original boot disk name and Root2 and Root3 are the mirror disk names.

```
Root1
Root2
Root3
```

2) If booting is successful, find out what caused the original disk to fail, and recover the disk by swapping disks.



See

For details on disk swap, see "[B.1.8 sdxswap - Swap Disk](#)" and resolutions described in "[7.3.1.2 Operation Procedure](#)."

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 "[\(5\) System cannot be booted. \(Boot disk data damage\)](#)."

(4) System cannot be booted. (Failure of all boot disks)

Explanation

Even when you have the system disk mirrored, you may experience unsuccessful booting.

The following describes the restoration procedure when all the boot disks physically fail simultaneously.

For the following procedure, swap all the physical disks registered in a root class and restore the system disk data based on collected backups. Then, re-construct the mirroring configuration.

Also, this section explains how to perform the mirroring cancellation to boot the system without swapping a physical disk or restoring data of the system disk by assuming that booting cannot be performed by some abnormality other than a physical disk failure.



Note

Cancel the mirroring of the normal system disk by the procedure in "[9.4 System Disk Mirroring Cancellation \[EFI\]](#)."

Note

To perform restoration according to this procedure, you must collect the following information beforehand when setting the system disk mirroring:

- Partition configuration of the system disk
- Backup of the system disk

For details, see "[6.5 System Disk Mirroring Settings \[EFI\]](#)."

Resolution

Perform restoration according to the following procedure.

1) Swap all the physical disks registered in the root class. You must swap them according to the procedure defined for each device.

Skip this step if the disks do not have a failure.

2) Boot the system from a CD-ROM device.

3) Set the new system disk to the original partition configuration.

See

For details on the parted command, see the manual for the parted(8) command.

3-1) Start the parted command.

```
# parted /dev/sda
```

3-2) Create the disk label (partition table) of the GPT type with the mklable subcommand.

```
(parted) mklable gpt
```

3-3) Before setting the system disk mirroring, create a partition according to the partition configuration checked in step 1) in "[6.5 System Disk Mirroring Settings \[EFI\]](#)."

3-3-1) Create the partition with the mkpart subcommand.

- For RHEL7

```
(parted) mkpart "EFI System Partition" fat16 1M 630M
(parted) mkpart " " xfs 630M 15G
(parted) mkpart " " xfs 15G 17G
(parted) mkpart " " xfs 17G 19G
(parted) mkpart " " xfs 19G 21G
(parted) mkpart " " xfs 21G 23G
(parted) mkpart " " linux-swap 23G 25G
```

- For RHEL8 or later

```
(parted) mkpart "EFI System Partition" fat32 1M 630M
(parted) mkpart " " xfs 630M 15G
(parted) mkpart " " xfs 15G 17G
(parted) mkpart " " xfs 17G 19G
(parted) mkpart " " xfs 19G 21G
(parted) mkpart " " xfs 21G 23G
(parted) mkpart " " linux-swap 23G 25G
```

3-3-2) Set the boot partition.

Example: When the file system type is FAT and the partition number is 1

```
(parted) set 1 boot on
```

- For RHEL7

```
# mkfs.vfat -v -c -F 16 /dev/sda1
```

- For RHEL8 or later

```
# mkfs.vfat -v -c -F 32 /dev/sda1
```

3-4) Confirm that the partition configuration is the same as the one checked in step 1) in "[6.5 System Disk Mirroring Settings \[EFI\]](#)" with the print subcommand.

```
(parted) print
Disk /dev/sdf: 34.4GB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start   End     Size    File system  Name                  Flags
  1      1049kB  630MB  629MB   fat32        EFI System Partition  boot, esp
  2      630MB   15.0GB 14.4GB   xfs
  3      15.0GB  17.0GB 2000MB   xfs
  4      17.0GB  19.0GB 2001MB   xfs
  5      19.0GB  21.0GB 2000MB   xfs
  6      21.0GB  23.0GB 2001MB   xfs
  7      23.0GB  25.0GB 2000MB   linux-swap(v1)      swap
```

Note

If a disk was not swapped in step 1), the private slice of GDS remains. At this time, if you execute the print subcommand, the private slice of GDS is displayed. Among the partitions displayed by the print subcommand, the partitions excluding the partitions that were checked in step 1) of "[6.5 System Disk Mirroring Settings \[EFI\]](#)" is the private slice of GDS. In this case, delete the private slice of GDS with the rm subcommand of the parted command.

Example: When the partition number (displayed in Minor filed) is 8

```
(parted) rm 8
```

Note

When a mirror destination disk is not replaced in step 1), remove all the partitions of the mirror destination disk.

Example: When the mirror destination disk is /dev/sdc

```
parted /dev/sdc
(parted) rm 1
(parted) rm 2
(parted) rm 3
(parted) rm 4
(parted) rm 5
(parted) rm 6
```



```
(parted) rm 7
(parted) rm 8
```

3-5) Exit the parted command.

```
(parted) quit
```

4) Restore the backup data collected in a tape medium and so on to each physical slice.

Take this step if the physical disk is swapped in step 1) or if the system disk data becomes invalid.



See

For the restoration, see the manuals of the OS, file systems to be restored and used commands.

5) Cancel the setting of the system disk mirroring.

5-1) Mount the EFI system partition and the root partition.

In the following example, the /work/efi, /work/root, and /work/boot directories are used as temporary mount points.

```
# mkdir /work
# mkdir /work/efi
# mkdir /work/root
# mkdir /work/boot
# mount /dev/sda7 /work/efi
# mount /dev/sda5 /work/root
# mount /dev/sda6 /work/boot
```



Information

If you have never performed the operation to update the initrd file (such as updating kernels) after mirroring the system disk, you can cancel the setting easily and quickly by performing the procedure "[Editing procedure of the grub.cfg file when the initrd file is not updated.](#)" in Information described later instead of steps 5a-2) and 5a-3).

Execute this procedure only when you can ensure that the operation to update the initrd file has not been performed. In other cases, perform the following steps 5a-2) through 5a-12).

5-2) Delete the information of GDS from the initrd file.

5-2-1) Check the entry name.

Check the entry name described in the set default entry in the /work/efi/EFI/redhat/grub.cfg file.

```
# more /work/efi/EFI/redhat/grub.cfg
### BEGIN /etc/grub.d/00_header ###
set pager=1

if [ -s $prefix/grubenv ]; then
  load_env
fi
if [ "${next_entry}" ] ; then
  set default="${next_entry}"
  set next_entry=
  save_env next_entry
  set boot_once=true
else
  set default="GDS_sysvol_entry"
fi
```

5-2-2) Check the initrd file name.

Check the initrd file name of the entry checked in step 5a-2-1) among the entries in the grub.cfg file.

In the following example, the initrd file name is `initramfs-3.10.0-123.el7.x86_64.img`.

```
# more /work/efi/EFI/redhat/grub.cfg
menuentry 'GDS_sysvol_entry' --class red --class gnu-linux --class gnu --class os --unrestricted {
    (*1)
    load_video
    set gfxpayload=keep
    insmod gzio
    insmod part_gpt
    insmod xfs
    insmod regexp
    regexp -s device '(hd[^\,]+)'
    set root="$device,gpt2"
    linuxefi /vmlinuz-3.10.0-123.el7.x86_64 root=/dev/sfdsk/gdssys2
    ro vconsole.keymap=jp106 crashkernel=auto
    vconsole.font=latarcyrheb-sun16 rhgb quiet LANG=ja_JP.UTF-8
    rd.driver.post=sfdsksys
    initrdefi /initramfs-3.10.0-123.el7.x86_64.img
    (*2)
}
```

(*1) Entry name checked in step 5a-2-1)

(*2) initrd file name

5-2-3) Create two work directories in the /work/root directory.

Create directories with different names from existing ones.

```
# mkdir /work/root/work_directory1
# mkdir /work/root/work_directory2
```

5-2-4) Extract the initrd file to each work directory.

In the following procedure, extract the initrd file to the work directory created in step 5a-2-3.

If in the following procedure, the message "no space left on device" is output because the file system mounted on the /work/root does not have enough space, mount another file system with enough space. Also create work directories and then extract the initrd file.

1) Check that the initrd file contains Early CPIO image.

If "Early CPIO image" is output when the following command is executed, the initrd file contains Early CPIO image.

```
# /work/root/usr/bin/lsinitrd /work/boot/initramfs-3.10.0-123.el7.x86_64.img \
    (*1)
2>/dev/null | grep -e "^Early CPIO image"
```

(*1) initrd file name checked in step 5a-2-2)

2) If the initrd file contains Early CPIO image, take the procedure below to extract the initrd file.

```
# cd /work/root/work_directory1
# cpio -idm < /work/boot/initramfs-3.10.0-123.el7.x86_64.img
    (*1)
# cd /work/root/work_directory2
# /work/root/usr/lib/dracut/skipcpio /work/boot/initramfs-3.10.0-123.el7.x86_64.img \
| gunzip -c - | cpio -idm
    (*1)
```

(*1) initrd file name checked in step 5a-2-2)

3) If the initrd file does not contain Early CPIO image, take the procedure below to extract the initrd file.

```
# cd /work/root/work_directory2
# gunzip -c /work/boot/initramfs-3.10.0-123.el7.x86_64.img | cpio -idm
(*1)
```

(*1) initrd file name checked in step 5a-2-2)

5-2-5) Remove GDS drivers from extracted initrd.

```
# rm /work/root/work_directory2/lib/modules/3.10.0-123.el7.x86_64/extra/FJVSsdm-drvcore/
sfdsksys.ko
# rm /work/root/work_directory2/lib/modules/3.10.0-123.el7.x86_64/extra/FJVSsdm-drvcore/
sfdsk_lib.ko
```

5-2-6) Remove shell scripts to read GDS drivers from extracted initrd.

Delete the following files if they exist:

- /work/root/work_directory2/sbin/ins_sfdsksys.sh
- /work/root/work_directory2/usr/lib/dracut/hooks/cmdline/01-parse-ins_sfdsksys.sh

```
# rm /work/root/work_directory2/sbin/ins_sfdsksys.sh
# rm /work/root/work_directory2/usr/lib/dracut/hooks/cmdline/01-parse-ins_sfdsksys.sh
```

5-2-7) Compress edited initrd.

1) If the initrd file contains Early CPIO image in 1) of step 5a-2-4), take the procedure below to compress initrd.

```
# cd /work/root/work_directory1
# find . -print0 | cpio --null -H newc -o --quiet > /work/boot/initrd-new.img
# cd /work/root/work_directory2
# find . -print0 | cpio --null -H newc -o --quiet | gzip -9 >> /work/boot/initrd-
new.img
```

2) If the initrd file does not contain Early CPIO image in 1) of step 5a-2-4), take the procedure below to compress initrd.

```
# cd /work/root/work_directory2
# find . -print0 | cpio --null -H newc -o --quiet | gzip -9 > /work/boot/initrd-new.img
```

5-2-8) Delete the work directories created in the step 5a-2-3).

```
# cd /
# rm -rf /work/root/work_directory1
# rm -rf /work/root/work_directory2
```

5-2-9) Place the initrd file compressed in step 5a-2-7) in the EFI system partition.

Back up the initrd file checked in step 5a-2-2), and then place the initrd file compressed in step 5a-2-7) as the same name of the initrd file checked in step 5a-2-2).

```
# cd /work/boot
# mv initramfs-3.10.0-123.el7.x86_64.img initramfs-3.10.0-123.el7.x86_64.img.gdsbak
# mv /work/boot/initrd-new.img /work/boot/initramfs-3.10.0-123.el7.x86_64.img
```

5-3) Edit the grub.cfg file.

```
# vim /work/efi/EFI/redhat/grub.cfg
```

Change the entry checked in step 5a-2-1) among the grub.cfg file entries as follows:

- Example before edit

```
menuentry 'GDS_sysvol_entry' --class red --class gnu-linux --class gnu --class os --unrestricted {
    (*1)
    load_video
    set gfxpayload=keep
    insmod gzio
    insmod part_gpt
    insmod xfs
    insmod regexp
    regexp -s device '(hd[^,]+)'
    set root="$device,gpt2"
    linuxefi /vmlinuz-3.10.0-123.el7.x86_64 root=/dev/sfdsk/qdssys2
                                                    (*2)
    ro vconsole.keymap=jp106 crashkernel=auto
    vconsole.font=lataarcyrrheb-sunl6 rhgb quiet LANG=ja_JP.UTF-8
    rd.driver.post=sfdsksys
    (*3)
    initrdefi /initramfs-3.10.0-123.el7.x86_64.img
}
```

- Example after edit

```
menuentry 'GDS_sysvol_entry' --class red --class gnu-linux --class gnu --class os --unrestricted {
    (*1)
    load_video
    set gfxpayload=keep
    insmod gzio
    insmod part_gpt
    insmod xfs
    insmod regexp
    regexp -s device '(hd[^,]+)'
    set root="$device,gpt2"
    linuxefi /vmlinuz-3.10.0-123.el7.x86_64 root=UUID=ef9cab98-5b1d-4fe2-89e6-4e1548778338
                                                    (*2)
    ro vconsole.keymap=jp106 crashkernel=auto
    vconsole.font=lataarcyrrheb-sunl6 rhgb quiet LANG=ja_JP.UTF-8
    initrdefi /initramfs-3.10.0-123.el7.x86_64.img
}
```

(*1) Entry name checked in step 5a-2-1)

(*2) Change the root device name from the path of the GDS's logical volume to UUID.
(Check UUID by executing the blkid command.)

(*3) Delete the description of rd.driver.post=sfdsksys.

Information

Editing procedure of the grub.cfg file when the initrd file is not updated.

If you have never performed the operation to update the initrd file (such as updating kernels) after mirroring the system disk, you can cancel the setting easily and quickly by performing the following procedure instead of steps 5-2) and 5-3).

1. Check the entryname.

Check the entryname described in the set default entry in the /work/efi/EFI/redhat/grub.cfg file.

```

### BEGIN /etc/grub.d/00_header ###
set pager=1

if [ -s $prefix/grubenv ]; then
  load_env
fi
if [ "${next_entry}" ] ; then
  set default="${next_entry}"
  set next_entry=
  save_env next_entry
  set boot_once=true
else
  set default="GDS_sysvol_entry"
fi

```

2. Edit the grub.cfg file.

```
# vim /work/efi/EFI/redhat/grub.cfg
```

Change the kernel image setting after setting the system disk mirroring defined in the grub.cfg file to a setting which does not mirror the system disk. The entry checked in step 1 displayed is the kernel image setting after mirroring the system disk among entries in the grub.cfg file. Examples are shown below.

- Example before edit

```

menuentry 'GDS_sysvol_entry' --class red --class gnu-linux --class gnu --class os --
unrestricted {
    (*1)
    load_video
    set gfxpayload=keep
    insmod gzio
    insmod part_gpt
    insmod xfs
    insmod regexp
    regexp -s device '(hd[^,]+)'
    set root="$device,gpt2"
    linuxefi /vmlinuz-3.10.0-123.el7.x86_64 root=/dev/sfdsk/gdssys2
    (*2)
    ro vconsole.keymap=jp106 crashkernel=auto
    vconsole.font=latacyrheb-sun16 rhgb quiet LANG=ja_JP.UTF-8
    rd.driver.post=sfdsksys
    (*3)
    initrdefi /initramfs-3.10.0-123.el7.x86_64.img
    (*4)
}

```

- Example after edit

```

menuentry 'GDS_sysvol_entry' --class red --class gnu-linux --class gnu --class os --
unrestricted {
    (*1)
    load_video
    set gfxpayload=keep

    insmod gzio
    insmod part_gpt

    insmod xfs
    insmod regexp
    regexp -s device '(hd[^,]+)'
    set root="$device,gpt2"

```

```

linuxefi /vmlinuz-3.10.0-123.el7.x86_64 root=UUID=ef9cab98-5b1d-4fe2-89e6-4e1548778338
                                                    (*2)

ro vconsole.keymap=jp106 crashkernel=auto
vconsole.font=latacyrheb-sun16 rhgb quiet LANG=ja_JP.UTF-8
initrdefi /initramfs-3.10.0-123.el7.x86_64.img.gdsbak
                                                    (*4)
}

```

- (*1) Entry name checked in step 1
- (*2) Change the root device name from the path of the GDS's logical volume to UUID.
(Check UUID by executing the blkid command.)
- (*3) Delete the description of rd.driver.post=sfdsksys.
- (*4) Add .gdsbak for the final entry of the initrd file.



5-4) Edit the fstab file.

```
# vim /work/root/etc/fstab
```

Edit the mount setting lines of volumes in the root class described in the fstab file according to the partition configuration of the new system disk. Examples are shown below.

- Example before edit

```

/dev/sfdsk/gdssys2    /           ext4  defaults  1  1
/dev/sfdsk/gdssys5    /boot      ext4  defaults  1  2
/dev/sfdsk/gdssys6    /boot/efi  vfat  defaults  0  0
/dev/sfdsk/gdssys3    /etc/opt   ext4  defaults  1  2
/dev/sfdsk/gdssys32   /usr       ext4  defaults  1  2
/dev/sfdsk/gdssys4    /var       ext4  defaults  1  2
/dev/sfdsk/gdssys33   swap       swap  defaults  0  0

```

- Example after edit

```

UUID=ef9cab98-5b1d-4fe2-89e6-4e1548778338    /           ext4  defaults  1  1
UUID=d8355b07-2cae-4604-86d8-88e8fbcc5c54     /boot      ext4  defaults  1  2
UUID=6580-6FA6                                /boot/efi  vfat  defaults  0  0
UUID=965bac91-dfd0-471d-84c4-4fa10592c7c4     /etc/opt   ext4  defaults  1  2
UUID=86ef50f9-aebd-4cf2-9244-24755655ffba     /usr       ext4  defaults  1  2
UUID=eec4b0e6-d9f7-4e3d-aee4-403f824d3f82     /var       ext4  defaults  1  2
UUID=d340f796-dbcf-4e17-9876-0cd57b9f7722     swap       swap  defaults  0  0

```

5-5) Edit the dracut.conf file. If the patch for PH15096 is applied, this procedure is not required.

```
# vim /work/root/etc/dracut.conf
```

Delete the description of the GDS driver from the dracut.conf file.

- Example before edit

```

...
# install local /etc/lvm/lvm.conf
lvmconf="yes"
add_drivers+=" sfdsksys"
                    (*1)

```

- Example after edit

```

...
# install local /etc/lvm/lvm.conf

```

```
lvmconf="yes"
add_drivers+=" "
(*2)
```

(*1) Description of the GDS driver

(*2) Leave one space for a single-byte character and delete sfdsksys.

5-6) Edit the /etc/default/grub file.

Change the setting of GRUB_DEFAULT in the /etc/default/grub file.

Delete rd.driver.post=sfdsksys from the file as well.

```
# vim /work/root/etc/default/grub
```

- Example before edit

```
GRUB_TIMEOUT=5
GRUB_DISTRIBUTOR="$(sed 's, release .*$,,g' /etc/system-release)"
GRUB_DEFAULT='GDS sysvol entry'
GRUB_DISABLE_SUBMENU=true
GRUB_TERMINAL_OUTPUT="console"
GRUB_CMDLINE_LINUX="vconsole.keymap=jp106 crashkernel=auto
                    vconsole.font=latarcyrheb-sun16 rhgb quiet
                    rd.driver.post=sfdsksys"
GRUB_DISABLE_RECOVERY="true"
```

- Example after edit

```
GRUB_TIMEOUT=5
GRUB_DISTRIBUTOR="$(sed 's, release .*$,,g' /etc/system-release)"
GRUB_DEFAULT='Red Hat Enterprise Linux Server, with Linux 3.10.0-123.el7.x86_64'
(*1)

GRUB_DISABLE_SUBMENU=true
GRUB_TERMINAL_OUTPUT="console"
GRUB_CMDLINE_LINUX="vconsole.keymap=jp106 crashkernel=auto
                    vconsole.font=latarcyrheb-sun16 rhgb quiet
                    "
GRUB_DISABLE_RECOVERY="true"
```

(*1) Kernel version that is included in the initrd file name checked in step 5a-2-2)

5-7) Remove the following directories if they exist:

- /work/root/usr/lib/dracut/modules.d/90k_sfdsksys

```
# rm -rf /work/root/usr/lib/dracut/modules.d/90k_sfdsksys
```

- /work/root/etc/opt/FJVSdx/bkup/grub

```
# rm -rf /work/root/etc/opt/FJVSdx/bkup/grub
```

5-8) Change the permission of the root_mknod.sh file.

```
# chmod 400 /work/root/etc/opt/FJVSdx/RootMirror/root_mknod.sh
```

5-9) Delete the information of GDS configuration databases set in the EFI environment variables.

```
# cd /sys/firmware/efi/vars/
# cat sysvol-db2-*/raw_var > del_var
```

5-10) Unmount the EFI system partition and the root partition.

```
# cd /
# umount /work/efi
# umount /work/root
# umount /work/boot
# rmdir /work/efi
# rmdir /work/root
# rmdir /work/boot
# rmdir /work
```

5-11) Restart the system.

```
# shutdown -r now
```

5-12) After the system is restarted, execute the following command to update the grub.cfg file.

```
# grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg
```

6) Add the new boot disk to a boot option of the EFI boot manager (BIOS for PRIMEQUEST 3000).
The procedure differs depending on the OS version.

6-1) Select [Boot Maintenance Manager] in the boot option selection screen of the EFI boot manager (Boot tab of BIOS for PRIMEQUEST 3000).

6-2) Select [Boot Option] in the Boot Maintenance Manager screen.

6-3) Select [Add Boot Option] in the Boot Options screen.

6-4) When the Add Boot Option screen is displayed, add the new boot disk to the boot option.

1. Select the device corresponding to /dev/sda from the device list.

For the method to recognize the device, see the manual of EFI.

2. After the contents in the root directory of the selected disk are displayed, select "EFI."

3. After the contents in the EFI directory are displayed, select "redhat."

4. After the contents in the redhat directory are displayed, select "shim.efi" (for RHEL7) or "shimx64.efi" (for RHEL8).

5. Select "Input the description."

6. Input a character string used for a label.

For example, input "New_sda" here.

7. Select "Commit Changes and Exit."

8. For PRIMEQUEST 3000, input [Esc] and exit [Boot Options].

9. Input [Esc] to exit [Boot Maintenance Manager].

6-5) For PRIMEQUEST 3000, select [Boot Override] on the Exit tab of BIOS.

For other than PRIMEQUEST 3000, select [Boot Manager] in the boot option selection screen of the EFI boot manager.

6-6) Select the label name input in 6. of step 6-4) in the option selection screen (Boot Override screen for PRIMEQUEST 3000) and boot the system.

7) Delete the SDX disk items from the boot option of the EFI boot manager.



See

For details on the efibootmgr command, see the manual of efibootmgr command.

7-1) Check the current setting of the EFI boot manager.

```
# efibootmgr
BootCurrent: 0000
Timeout: 5 seconds
BootOrder: 0001,0002,0000
Boot0000* sda
Boot0001* RootDisk1
Boot0002* RootDisk2
```

In this example, the SDX disk items are RootDisk1 of the number 0001 and RootDisk2 of the number 0002.

7-2) Delete the SDX disk items from the boot option.

```
# efibootmgr -B -b 0001
BootCurrent: 0000
Timeout: 5 seconds
BootOrder: 0002,0000
Boot0000* sda
Boot0002* RootDisk2
# efibootmgr -B -b 0002
BootCurrent: 0000
Timeout: 5 seconds
BootOrder: 0000
Boot0000* sda
```

8) Reconfigure the system disk mirroring.



See

For details on the setting methods, see "6.5 System Disk Mirroring Settings [EFI]."

(5) 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 "7.4.6 Restoring (When System Cannot Be Booted) [EFI]."

(6) System cannot be booted. (Removing errata kernel)

Explanation

After an errata kernel has been removed in an environment where the system disk is mirrored, and if the grub.cfg file is not updated by using the grub2-mkconfig command before the system is restarted, the system may not be started.



See

For notes on removing an errata kernel, see "7.18 Removing Errata Kernel."

Resolution

1. To start the system, select the boot entry of "GDS sysvol entry for Red Hat Enterprise Linux Server, ..." format from the boot option selection screen of the EFI boot manager.
2. After the system is started, execute the grub2-mkconfig command to update the grub.cfg file.

```
# grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg
```

3. If no boot entry exists in step 1 above, follow "(4) System cannot be booted. (Failure of all boot disks)" to restore the system. The disk is not necessary to be replaced at this time.

D.1.6 Errors in the Mirroring Among Servers

When one of the following errors occurs, take the actions described below for each error.

- (1) The netmirror slice that configures the netmirror volume is in INVALID state.
- (2) The cluster application becomes Faulted or Inconsistent after the node is restarted.
- (3) The cluster application becomes Faulted or Inconsistent when the network error occurs.
- (4) The slice becomes INVALID or the operation stops after restarting the other node when the network error occurs.
- (5) The resource of the class where the netmirror volume exists becomes OFF-FAIL.
- (6) The class is closed on the last started node at the startup of both nodes.
- (7) The node is restarted while hot swap of disks is proceeding.
- (8) The message regarding an unmount is output when both nodes stop.

(1) The netmirror slice that configures the netmirror volume is in INVALID state.

Explanation

The slice of the netmirror volume becomes INVALID due to the following reasons:

(Cause a)

- Node stop
- Error in the network that is used for the mirroring among servers
- Disk error

(Cause b)

In an Azure environment or a NIFCLOUD environment, the path for the by-id file was specified, and then the iSCSI target was created.

Resolution

The resolution procedure is illustrated below for each of the two causes a and b.

a) For (Cause a)

1. Start a node if the node is stopped.

After the node is started, synchronization copying is performed automatically. When the synchronization copying is completed without an error, the slice is restored to be ACTIVE. In this case, no more resolution is necessary.

If the slice is not be restored, take the following procedure.

2. Check the status of the network that is used for the mirroring among servers.

If the network has an error, restore the network.

After the network is restored, synchronization copying is started automatically within 30 seconds. Wait until all the synchronization copying is completed. In other words, wait until the COPY netmirror slice no longer exists. When the synchronization copying is completed without an error, the slice is restored to be ACTIVE. In this case, no more resolution is necessary.

If the slice is not restored, take the following procedure.

3. Check the class status. If the class is not started, restore it.

Execute the following command on both nodes to check the class status.

```
# /etc/opt/FJSVsdx/bin/sdxdcdown
```

If the class information is not displayed, the class is not started.

In this case, see the note "Shutting down the node" of the section "(2) Class cannot be started when booting the system." in "D.1.4 Class Status Abnormality", stop the node in question, and then restart it.

If the information is still not displayed after the node is restarted, see the section "(2) Class cannot be started when booting the system." in "D.1.4 Class Status Abnormality" to restore.

4. Check the disk status.
5. If the disk has a failure and the class is closed at the same time, restore the closed class.

5-1) Check if the class is closed or not on both nodes.

Execute the following command on both nodes. If "yes" is displayed in the DOWN field, the class is closed.

```
# /etc/opt/FJSVsdx/bin/sdxdcdown
```

5-2) Restart the node where the class is closed.

5-3) If the cluster application where the closed class is registered has been started on any other node except this node, cancel the start lock of the netmirror volume on this node.

```
# sdxattr -v -c class_name -v volume_name -a lock=off
```

6. If the disk has a failure, swap the disk.

For how to swap the disk, see "7.3.3 Swapping Disks of Netmirror Group."

However, if the error status is equal to the conditions described in the following troubleshootings, see "Resolution" of each troubleshooting to restore the error status:

- "(2) The copy destination slice was made INVALID due to an I/O error generated on the copy source slice during synchronization copying." in "D.1.1 Slice Status Abnormality"
- "(7) I/O error occurs although mirror volume is in ACTIVE status." in "D.1.3 Volume Status Abnormality"

Take the procedure described in "7.3.3 Swapping Disks of Netmirror Group" first before executing "restore physical disk" (sdxswap -I command) according to "Resolution" described in the above sections.

7. If the disk does not have an error, execute the sdxinfo -D command on both nodes, and then check if the physical disk name is properly displayed in the DEVNAM field.

When the asterisk (*) is displayed in the DEVNAM field on one of the nodes, follow the procedure described in "(4) The slice becomes INVALID or the operation stops after restarting the other node when the network error occurs." in "D.1.6 Errors in the Mirroring Among Servers" to restore the error status.

If the asterisk (*) is displayed in DEVNAM field on both nodes, follow the procedure described in "(6) The class is closed on the last started node at the startup of both nodes." in "D.1.6 Errors in the Mirroring Among Servers" to restore the error status.

8. If the error status is not restored after taking steps 1 to 7, check that the following errors do not occur and then execute the resolution.
 - When an error and restoration occurs repeatedly in the network or in the node, or when I/O error and network recovery occurs consecutively due to off and on failures.

Check the output result such as the system log to check the status of the network and the node. If an error is detected, restore this error and then to perform synchronization copying of the netmirror volume, execute the following command on any one node.

```
# sdxcopy -B -c class_name -v volume_name
```

- When the data between the disks of both nodes are unable to be determined if new or old.
For details on conditions for this state and restoration method, see "[7.16.8 Restoration when the Latest Disk Cannot Be Selected Automatically](#)."

9. If the slice is not restored after taking the steps 1 to 8, collect the investigation material and contact field engineers.

b) For (Cause b)

Take the following steps to remove the configuration, specify the path for the by-partuuid file, and then create the iSCSI target again.

1. Restore the disk.

See "[7.3.3 Swapping Disks of Netmirror Group](#)."

You can choose which of the procedures described in "[7.3.3.1 Hot Swap](#)" and "[7.3.3.2 Cold Swap](#)" to perform.

Note that the following steps are not required because you do not need to swap the disk.

- "7. Swap the disk." in "[7.3.3.1 Hot Swap](#)"
- "5. Swap the disk." in "[7.3.3.2 Cold Swap](#)"

2. Back up the volume data if necessary.

3. Remove the configuration of GDS and the iSCSI settings used for the mirroring among servers.

For details, see "[Chapter 9 Removing Configuration](#)."

4. Create a partition in the disk where the gpt disk label is set, specify the path for the by-partuuid file for that partition, and then create the iSCSI target.

For details, see "[4.8.3 Creating iSCSI Target](#)."

5. Set up GDS.

For details, see "[Chapter 6 Settings](#)."

6. Restore the volume data if necessary.

(2) The cluster application becomes Faulted or Inconsistent after the node is restarted.

Explanation

If one slice is in ACTIVE or STOP state and the other slice is in other than these two states among two slices that belong to the netmirror volume, this error occurs when restarting the node where the former slice exists.

At this time, the slice that contains the valid data does not exist in the netmirror volume. In this case, the class resource on the node, which has not been restarted, becomes OFF-FAIL, and the cluster application becomes Faulted or Inconsistent.

When the node that is not restarted meets the following conditions, the cluster application becomes Inconsistent.

- The synchronization copying is being performed on the standby node.
- There is a copy destination disk on the standby node.

When the node does not meet the above conditions, the cluster application becomes Faulted.

Resolution

1. Start the stopped node if it is stopped.
2. Check that the class is not closed on the node which has not been restarted.

Execute the following command on the node which has not been restarted. If "yes" is displayed in the DOWN field, the class is closed.

```
# /etc/opt/FJVSdx/bin/sdxdown
```

3. Restore the class if it is closed in step 2.

If the class is not closed, this procedure is not required.

3-1. If the disk is failed, replace the disk for restoration.

3-2. Check that both nodes are started, and then restore the closed class on the node which has not been restarted.

Execute the following command.

```
# sdxfix -C -c class_name
```

3-3. Cancel the "Lock volume" of the netmirror volume.

Execute the following command on both nodes to check the "Lock volume" attribute.

```
# sdxinfo -V -c class_name -e long
```

The "Lock volume" attribute is displayed in the LOCK field.

If the "Lock volume" attribute is "on", execute the following command on the node to cancel the "Lock volume."

```
# sdxattr -V -c class_name -v volume_name -a lock=off
```

3-4. Perform the synchronization copying if INVALID slice exists in the netmirror volume.

Execute the following command on any one node.

```
# sdxcopy -B -c class_name -v volume_name
```

Note

When the time differs depending on each node in the cluster system, the synchronization copying may not be performed. To recover this error, adjust the time on each node to be consistent with each other and then restart both nodes.

4. If the cluster application is in Faulted or Inconsistent state, clear the Faulted or Inconsistent state.

For how to clear the Faulted or Inconsistent state, see "PRIMECLUSTER Installation and Administration Guide."

5. Check the class resource state. If it is OFF-FAIL, restore it.

- If the started cluster application does not exist

Restart the cluster application on the node which has not been restarted.

If you want to start the cluster application on the restarted node, restart the node which has not been restarted.

- If the only one started cluster application exists

Restart the standby node.

- If two or more cluster applications have been started

1) If the operating node and the standby node exist on one node, switch the cluster application so that only the node which has not been restarted becomes the operating node.

2) Switch back the cluster application switched in step 1) if necessary.

(3) The cluster application becomes Faulted or Inconsistent when the network error occurs.

Explanation

- When all the following conditions are met, the cluster application becomes Faulted.
 - Other than ACTIVE slice exists in the operating node.
 - The network used for the mirroring among servers is disconnected.

At this time, the ACTIVE slice that can be accessed from the operating node does not exist.

In this case, the cluster resource in the operating node becomes OFF-FAIL, and the cluster application becomes Faulted. Due to this, the cluster application is switched.

- When all the following conditions are met, the cluster application becomes Inconsistent.
 - The synchronization copying has been performed on the standby node.
 - The copy destination disk exists in the standby node.
 - The network used for the mirroring among servers is disconnected.

At this time, the ACTIVE slice that can be accessed in copying processing does not exist.

In this case, the cluster resource in the standby node becomes OFF-FAIL, and the cluster application becomes Inconsistent.

Resolution

1. Restore the network.

After the network is restored, the resynchronization copying is automatically performed.

When the class is closed, the copying process may fail.

2. Check the closed class.

Execute the following command on both nodes. The class has been closed if "yes" is displayed in its DOWN field.

```
# /etc/opt/FJSVsdx/bin/sdxdcdown
```

3. Restore the class if it is closed in step 2.

If the class is not closed, this procedure is not required.

3-1. If the disk is failed, replace the disk for restoration.

3-2. If the resynchronization copying is performed, cancel the copying process.

Execute the following command on the node where the class is not closed to check the slice status of the netmirror volume.

If the class is closed on both nodes, this procedure is not required.

```
# sdxinfo -S -c closed_class_name
```

When the COPY state slice exists in the netmirror volume, the resynchronization copying is ongoing.

In this case, execute the following command on the node where the class is not closed to cancel the copying process.

```
# sdxcopy -C -c closed_class_name -v volume_name
```

3-3. Restore the closed class.

Execute the following command on the node where the class is closed.

```
# sdxfix -C -c closed_class_name
```

While restoring the closed class, the synchronization copying is performed.

The synchronization copying of the whole volume is performed even when the both slices of the netmirror volume are already synchronized.

3-4. If the synchronization copying is not performed in step 3-3, perform the synchronization copying.

Execute the following command on any one node.

```
# sdxcopy -B -c closed_class_name -v volume_name
```

Note

When the time differs depending on each node in the cluster system, the synchronization copying may not be performed in step 3-3 and 3-4. To recover this error, adjust the time on each node to be consistent with each other and then restart both nodes.

4. Restore the state of the cluster application on the standby node.

For how to clear the Faulted state, see "PRIMECLUSTER Installation and Administration Guide."

To restore the Inconsistent state, take the same procedure as to clear the Faulted state.

5. Check the class resource state. If it is OFF-FAIL, restore it.

- If the only one cluster application exists

Restart the standby node.

- If multiple cluster applications exist.

1) If the operating node and the standby node exist on one node, switch the cluster application so that the operating node and the standby node do not exist on the same node.

2) If the resynchronization copying is performed, and the copy source disk exists on the standby node at the same time, wait until the resynchronization copying is completed.

3) Restart the standby node.

(4) The slice becomes INVALID or the operation stops after restarting the other node when the network error occurs.

Explanation

If the node is restarted when the network error occur, either of the following errors occur.

- A slice of the operational node becomes INVALID when starting the restarted node.
- After starting the restarted node, the synchronization copying is automatically performed after the network is restored. When the copying process is completed, the slice of the operational node becomes INVALID.
- After starting the restarted node, the synchronization copying is automatically performed after the network is restored. However, the synchronization copying fails.
- The operation stops, and the operation is not restarted after starting the restarted node.
- The operation stops, and the operation is restarted after starting the restarted node. However, the synchronization copying fails after the network is restored.

Resolution

1. If the network used for the mirroring among servers has an error, restore this network.
2. Check the device status.

Execute the `sdxinfo -D` command on both nodes.

For disks that are connected with the netmirror group, check and record the following information.

The information can be obtained from the output of the `sdxinfo -D` command. The recorded information is used in step 7 and step 8 later.

- Check if an asterisk (*) is displayed in the DEVNAM field (physical disk name) of a connected disk. If an asterisk is displayed, record the value in the NAME field (SDX disk name) of the disk. In the example below, the value is d2.
- Record the node name if an asterisk (*) is displayed in the DEVNAM field (physical disk name).

Execution example

```
# sdxinfo -D
OBJ      NAME      TYPE      CLASS      GROUP      DEVNAM      DEVBLKS      DEVCONNECT      STATUS
```

```

-----
disk  d1      netmirror c11    mdg1   sda    1015808  node1:node2  ENABLE
disk  d2      netmirror c11    mdg1   *      1015808  *            ENABLE
-----

```

- Restore the iSCSI device information.

Execute the following command on the node where an asterisk (*) has been displayed in the DEVNAM field of the disk, which is connected with the netmirror group, in step 2 above.

```
# /etc/opt/FJSVsdx/bin/sdxiscsi_ctl -F -e init
```

- Check if the class is closed on both nodes. Restore the class if it is closed.

For details, see "[D.1.4 Class Status Abnormality](#)."

- Check the status of "Lock volume" of the netmirror volume.

Execute the following command on both nodes.

```
# sdxinfo -v -c class_name -e long
```

The "Lock volume" attribute is displayed in the LOCK field.

- If the "Lock volume" attribute is "on" on any one of the nodes, take the following procedures to cancel the "Lock volume."

- When the "Lock volume" attribute is "on" on all nodes

Execute the following command on any one of the nodes in the class scope.

```
# /etc/opt/FJSVsdx/bin/sdxnetdisk -S -c class_name
```

- When the "Lock volume" attribute is "on" on one node

Execute the following command on the node where the "Lock volume" attribute is "on."

```
# sdxattr -v -c class_name -v volume_name -a lock=off
```

Make sure that the "Lock volume" attribute is "off" on both nodes after executing the command.

If the "Lock volume" attribute is not "off", perform the "Restoration Procedure" 1. to 3. of "c." in "[7.16.8 Restoration when the Latest Disk Cannot Be Selected Automatically](#)", and then take the above procedure again.

- Delete the iSCSI device information on both nodes.

```
# rm /var/opt/FJSVsdx/log/.sdxnetmirror_disable.db
# rm /var/opt/FJSVsdx/log/.sdxnetmirror_timestamp
```

- Check the slice status.

Execute the following command on any one node. Check the value in the STATUS field of the slice if the value in the DISK field (SDX disk name) of the slice is equal to the SDX disk name recorded in step 2.

The slice status is INVALID in the execution example below.

Execution example

```
# sdxinfo -s
OBJ    CLASS  GROUP  DISK   VOLUME  STATUS
-----
slice  c11    mdg1   d1     m1      STOP
slice  c11    mdg1   d2     m1      INVALID
```

- If the slice status is ACTIVE or STOP when it is checked in step 8, take the following procedure.

9-1) If the slice status is INVALID when it is checked in step 7 and this slice belongs to the disk of node which is not recorded in step 2, stop RMS.

Execute the following command on any node.

```
# hvshut -a
```


9-2) Restart the node recorded in step 2.

9-3) If RMS is stopped in step 9-1), start RMS. Execute the following command on any node.

```
# hvcm -a
```

10. If the INVALID slice exists in the netmirror volume, restore the slice.

Execute the following commands on any one node.

Use the -d option to specify the disk that contains the INVALID slice. If the volume is not started on any node, the synchronization copying is performed when the volume is started.

```
# sdxswap -O -c class_name -d disk_name
# sdxswap -I -c class_name -d disk_name
```

11. Restore the state of the cluster application.

For how to clear the Faulted state, see "PRIMECLUSTER Installation and Administration Guide."

To restore the Inconsistent state, take the same procedure as to clear the Faulted state.

12. Check the cluster resource state. If it is OFF-FAIL, restore it.

- If the started cluster application does not exist

Restart the cluster application on the node which has not been restarted.

If you want to start the cluster application on the restarted node, restart the node which has not been restarted.

- If the only one started cluster application exists

Restart the standby node.

- If two or more cluster applications have been started

1) If the operating node and the standby node exist on one node, switch the cluster application so that only the node which has not been restarted becomes the operating node.

2) Switch back the cluster application switched in step 1) if necessary.

13. When the service is stopped, restart the service if necessary.

(5) The resource of the class where the netmirror volume exists becomes OFF-FAIL.

Explanation

When the node cannot access the netmirror volume due to an error, the class resource in the node becomes OFF-FAIL.

Resolution

1. Under the following circumstances, restart the node where the OFF-FAIL class resource exists:

- Both nodes were restarted just before this error occurred.
- All the disks in at least one netmirror group are in either of the following status:
 - The disk of the node where the OFF-FAIL class resource exists is in SWAP state or DISABLE state.
 - On the node where the OFF-FAIL class resource exists, an asterisk (*) is displayed in DEVNAM field of the sdxinfo -D command output.

2. Restore the error state of the GDS object according to the troubleshooting described in the following sections:

(1) to (4), and (6) in "[D.1.6 Errors in the Mirroring Among Servers](#)"

"[D.1.4 Class Status Abnormality](#)"

"[D.1.2 Disk Status Abnormality](#)"

The class resource is restored successfully once its state is restored back to normal status (ON or OFF-STOP). When the OFF-FAIL class resource exists, go on to step 3. or later.

3. If the OFF-FAIL class resources exist on both nodes, take the following steps to restore their status.
 - 3-1) Select the node that is to be restarted.
 - 3-2) Start the cluster application that matches the following condition on the node that is not to be restarted.
Condition: The cluster resource is in the OFF-FAIL state on the node that is not be restarted.
 - 3-3) If only one ACTIVE slice or only one STOP slice exists in the netmirror volume, and this slice is the slice existing on the disk of the node that is selected in step 3-1), restore the INVALID slice or the NOUSE slice in the netmirror volume.
 - 3-4) Restart the node selected in step 3-1).
4. If the OFF-FAIL class resource exists on only one of the nodes, take the following steps to restore its state.
Restore the class resource state according to any one of the conditions described in step 4a) to 4c).
 - 4a) If all the cluster applications stop
Take any one of the following steps to restore the state.
 - On the node, start the cluster application that contains the OFF-FAIL class resource.
 - Restart the node on which the class resource is in OFF-FAIL state.
 - 4b) If the node where the OFF-FAIL class resource exists is the operating node
On the node, start the cluster application that contains the OFF-FAIL class resource.
 - 4c) If the node where the OFF-FAIL class resource exists is not the operating node
Take the following steps to restore the state.
 - 4c-1) If only one ACTIVE slice or only one STOP slice exists in the netmirror volume, and this slice is the slice existing on the disk of the node where the OFF-FAIL class resource exists, restore the INVALID slice or the NOUSE slice in the netmirror volume.
 - 4c-2) Restart the node where the OFF-FAIL class resource exists.

(6) The class is closed on the last started node at the startup of both nodes.

Explanation

At the startup of both nodes, if an error has occurred in the network that is used for the mirroring among servers, the class may be closed on the last started node.

At this time, the class resource may become OFF-FAIL also on the first started node. In this case, the cluster application cannot be started.

Resolution

Take different procedures depending on if the cluster application is started or not, and if the network can be restored soon or not.

- (a) If the cluster application is started
 1. Stop the node where the class is closed.
 2. Restore the network that is used for the mirroring among servers. After that, take the "Restoration Procedure" described in "[7.16.7 Restoration to 2-node Operation after Operation on Only One Node.](#)"
- (b) If the cluster application is stopped and the network can be restored soon as well
 1. Restore the network that is used for the mirroring among servers.
 2. Restart both nodes.
- (c) If the cluster application is stopped but the network cannot be restored soon
 1. Select one of the following procedures:
 - Restart the operation after restoring the network.
The operation status can be restored after restoring the network and restarting both nodes.
The data is not restored at this time. Instead, the operation is restarted by using the latest data stored in the disk of the node where the class is closed.

- Restore the data by the backup data to start the operation.
In this case, go on to step 2. or later.

2. Stop the node where the class is closed.

3. Restore the INVALID state of the slice.

Execute the following command. For -d option, specify the disk name of the disk that is connected with the node among all the disks belonging to the netmirror group.

```
# sdxfix -v -c class_name -v netmirror_volume_name -d disk_name -e force -x NoRdchk
```

4. Start the netmirror volume.

```
# sdxvolume -N -c class_name -v netmirror_volume_name -e unlock
```

5. Restore the data of the netmirror volume.

6. Stop the netmirror volume.

```
# sdxvolume -F -c class_name -v netmirror_volume_name
```

7. If the cluster application is in Faulted state, clear the Faulted state.

For how to clear the Faulted state, see "PRIMECLUSTER Installation and Administration Guide."

8. Start the cluster application forcibly.

For how to forcibly start the cluster application, see "PRIMECLUSTER Installation and Administration Guide."

9. Restore the network that is used for the mirroring among servers. After that, take the "Restoration Procedure" described in ["7.16.7 Restoration to 2-node Operation after Operation on Only One Node."](#)

(7) The node is restarted while hot swap of disks is proceeding.

Explanation

When the node is restarted while performing the operation in ["7.3.3.1 Hot Swap"](#) for the disks used for the mirroring among servers, take the actions described below.

Resolution

The resolution depends on whether the restarted node is a node where disks are swapped or not.

Resolution when a node where disks are swapped is restarted

When the node where disks are swapped is restarted, take the actions described below.

The step order described below follows the same steps as ["7.3.3.1 Hot Swap."](#)

Timing of restarting	Resolutions
Step 1 proceeding	After the node is restarted, check the slice status of a disk to be swapped. When the status is NOUSE, execute step 2 and the following steps. (*1) When the status is other than NOUSE, execute step 1 and the following steps. (*1) After step 12 is completed, restart the node where disks were swapped, then execute step 13
Step 1 completed, or from step 2 to 6 proceeding	Execute step 2 and the following steps. (*1) After step 12 is completed, restart the node where disks were swapped, then execute step 13.
Step 6 completed, or step 7 proceeding	After the node is restarted, when a disk swapping is not yet completed, execute step 2 and the following steps.(*1) When a disk swapping is completed, execute step 8 and the following steps. After step 12 is completed, restart the node where disks were swapped, then execute step 13.

Timing of restarting	Resolutions
Step 7 completed, or step 8.9 proceeding	Resume the process that you were executing before the node was restarted. After step 12 is completed, restart the node where disks were swapped, then execute step 13.
Step 9 completed, or from step 10 to 12 proceeding	Execute step 10 and the following steps. After step 12 is completed, restart the node where disks were swapped, then execute step 13.
Step 12 completed, or step 13, 14 proceeding	Execute step 13 and the following steps.
Step 14 completed, or step 15 proceeding	After the node is restarted, check the slice status of a disk to be swapped. When the status is NOUSE, execute step 15. When the status is other than NOUSE, all steps are completed.

(*1) When the device where "Swap Physical Disk" (or execute the `sdxswap -O` command) was executed was not created, step 2, 4, 5, and 6 are not necessary.

Resolution for a node where disks were not swapped is restarted

When the node where disks are not swapped is restarted, disks which can be normally accessed disappear, then the volume fails to be accessed, which belongs to the netmirror group where the disk is swapped.

Take the actions for resolution described below.

1. Execute steps in "Resolution" except restoring a class resource (step 5), in "(2) The cluster application becomes Faulted or Inconsistent after the node is restarted." of "D.1.6 Errors in the Mirroring Among Servers." However, do not execute a disk swapping. (step3-1).
2. Resume the procedure for accessing the volume as necessary.
3. Swap disks.

The steps for swapping disks depend on the timing when the node where disks are not swapped is restarted.

The step order described below follows the same steps as "7.3.3.1 Hot Swap."

Timing of restarting	Steps for disk swapping
Step 1 proceeding	After the node is restarted, check the slice status of a disk to be swapped. When the status is NOUSE, execute step 2 and the following steps. When the status is other than NOUSE, execute step 1 and the following steps.
Step 1 completed, or from step 2 to 4 proceeding	Execute step 2, then resume the process that you were executing before the node was restarted.
Step 4 completed, or from step 5 to 10 proceeding	Resume the process that you were executing before the node was restarted.
Step 10 completed, or from step 11 to 14 proceeding	Resume the process that you were executing before the node was restarted.
Step 14 completed, or step 15 proceeding	After the node is restarted, check the slice status of a disk to be swapped. When the status is NOUSE, execute step 15. When the status is other than NOUSE, all steps are completed.

4. Execute the class resource recovery (step 5) in "Resolution" in "(2) The cluster application becomes Faulted or Inconsistent after the node is restarted." of "D.1.6 Errors in the Mirroring Among Servers."

(8) The message regarding an unmount is output when both nodes stop.

Explanation

When both nodes stop, regarding Fsystem resource of the file system on the netmirror volume, the following RMS wizard (RMSWT) message may be output to the `/var/log/messages` file.

NOTICE: umount *mount_point* failed with error code 1

WARNING: The file system *mount_point* was not unmounted.

Resolution

Even if this message is output, the system is not affected. There is no need to work around.

D.1.7 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 the actions 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.](#)
- (4) [In the process of system disk setting, the volume name is not displayed correctly in the log information field. \[EFI\]](#)
- (5) [Insufficient memory error is displayed at Java console in the environment with a number of SDX objects..](#)
- (6) [The error message is displayed when starting 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 'httpip' 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 "PRIMECLUSTER 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 was changed during system operation.

Resolution

Perform the following steps 1. through 2.

1. Use the parted(8) 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/log/messages file and such.
2. Execute [Update Physical Disk Information] in the [Operation] menu.

(4) In the process of system disk setting, the volume name is not displayed correctly in the log information field. [EFI]

Explanation

In the process of system disk setting, a message displayed in the information field in the GDS Management View main screen may not show the volume name correctly.

Resolution

This error has no effect on the system. There is no need to work around.

(5) Insufficient memory error is displayed at Java console in the environment with a number of SDX objects.

Explanation

The memory of Java VM may become insufficient when using the GDS Management View in the environment with a number of SDX objects.

- A number of disks are registered in GDS.
- A number of volumes are created.
- There are many scope nodes of the shared class.

When applying the above-mentioned cases, change the size of the memory allocation pool of Java VM.

Resolution

Estimate the size to be added according to the following description and change the setting of Java VM.

Management server

Add 2 MB per 1000 of the value of "(the number of disks + the number of volumes) x the number of scope nodes."

Client

Add 5 MB per 1000 of the value of "(the number of disks + the number of volumes) x the number of scope nodes."

The value less than 1000 is rounded up.

Example

When registering 500 disks and creating 400 volumes to the shared disk of two nodes

Estimates of size to be added are as follows:

$$(500 + 400) \times 2 = 1800$$

The size is 2000, since the value less than 1000 is rounded up.

- For the management server

For the management server, 2 MB is added per 1000. Based on the calculating formula $2000 / 1000 \times 2$, add 4 MB.

- For clients

For clients, 5 MB is added per 1000. Based on the calculating formula $2000 / 1000 \times 5$, add 10 MB.

Based on the above-mentioned estimates, change the setting of Java VM.

How to change management server (Web-Based Admin View)

1. Display the shell prompt with the root administration.
2. Check the current setting value with the `wvGetparam` command.

```
# /etc/opt/FJSVwvbs/etc/bin/wvGetparam s-java-option
sys:s-java-option      -mx128m
```

3. Add the calculated value (4 MB) to the confirmed result and set it with the `wvSetparam` command.

```
# /etc/opt/FJSVwvbs/etc/bin/wvSetparam s-java-option -mx132m
sys:s-java-option -mx132m
```

4. Check that no client is connected and restart the management server.
For how to restart the management server, see "PRIMECLUSTER Web-Based Admin View Operation Guide."

How to change clients (Windows)

1. Select [Java] from [Control Panel] of the [Start] menu and click it to display [Java Control Panel].
2. Select the [Java] tab from [Java Control Panel] and click the <View> button to display the [Java Runtime Environment Settings] screen.
3. Check the value of [Runtime Parameters] and add it to the calculated value (10 MB), and then set the total value.
If the default value is -Xmx80m, set -Xmx90m.
If the default value is not set, assuming that the default value is -Xmx64m, add the calculated value.
4. Click the <OK> button of the [Java Runtime Environment Settings] screen.
5. Click the <OK> button of [Java Control Panel].

The setting works from the session which uses browsers started after performing the above-mentioned procedures.



(6) The error message is displayed when starting GDS Management View.

Explanation

When clicking the "Global Disk Services" button on the top menu of Web-Based Admin View, the following message is output. After that, when clicking the "Confirm" button of the message, GDS Management View is started.

```
5002
The following exception was received:
java.rmi.ServerException: RemoteException occurred in server thread; nested exception is:
    java.rmi.RemoteException: Socket server NG
```

The above message is displayed on the Web-Based Admin View management server when IP address for the local host name can not be got.

The host name is the name that is displayed when executing the hostname command or the uname -n command on the management server.

Even if the event occurs, it has no effect on the operation of GDS Management View.

Resolution

In the /etc/hosts file on the Web-Based Admin View management server, describe a pair of the host name of the management server and the correspondent IP address.

D.1.8 Proxy Object Abnormality

For proxy object related errors, in one of the following circumstances, take the actions as indicated for the relevant situation.

- (1) The Advanced Copy function cannot be used in master-proxy copying.
- (2) Alternative boot environment setup fails resulting in an error numbered 60305. [EFI]

(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 ETERNUS Disk storage system, the GDS soft copy function is performing the copy operation.

```
# sdxinfo -s -e long -o pv1
OBJ  NAME CLASS  GROUP  DISK  VOLUME  JRM  MODE  STATUS  COPY  CURBLKS  COPYBLKS  DLY  CPTYPE  CPSOURCE
-----
slice *   Class1 Group1 Disk1 pv1    *   *   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 QuickOPC function 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.8 Using Advanced Copy Function in 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 event 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.

(2) Alternative boot environment setup fails resulting in an error numbered 60305. [EFI]

Explanation

In a situation where the master group in the root class includes a volume that is not declared in the `/etc/fstab` file as a file system or a swap area, if the proxy group is specified and the `sdxproxy Root` command is executed, the following error message is output and the command fails.

```
sdxproxy: ERROR: Proxy_Volume_Name: corresponding volume not specified in /etc/fstab
```

If the current boot is from the proxy group, an event that the "master" and the "proxy" above are exchanged occurs. In this event, exchange the "master" and the "proxy" as shown in [Resolution] below and take action as indicated.

Resolution

In the examples shown below, the following configuration is assumed.

- The root class is `Class1`.
- The master group is `mg1`.
- Master volumes in `mg1` are `mv1`, `mv2`, and `mv3`.
- Among `mv1`, `mv2`, and `mv3`, `mv2` is not declared in `/etc/fstab`.
- The proxy group is `pg1`.
- Proxy volumes in `pg1` are `pv1`, `pv2`, and `pv3`.
- `mg1` and `pg1`, `mv1` and `pv1`, `mv2` and `pv2`, and `mv3` and `pv3` are related respectively.
- The root volume is `mv1`.
- The alternative boot volume is `pv1`.

With this configuration, the following error occurs.

```
# sdxproxy Root -c Class1 -p pg1
sdxproxy: ERROR: pv2: corresponding volume not specified in /etc/fstab
```

Choose one of the methods shown as (Resolution a) and (Resolution b) to take action.

(Resolution a)

Among proxy volumes in the proxy group, specify only volumes that meet the following conditions and execute the `sdxproxy Root` command.

- The proxy volume of the master volume that is declared in the `/etc/fstab` file.
- The master volume of the proxy volume that is declared in the `/etc/fstab` file.
- A proxy volume related to the master volume of the proxy volume that is declared in the `/etc/fstab` file.

Example)

```
# sdxproxy Root -c Class1 -p pv1,pv3
```

(Resolution b)

Following the procedures below, add descriptions of all volumes in the master group to the `/etc/fstab` file and execute the `sdxproxy Root` command and then delete unnecessary descriptions from `fstab`.

1) Among master volumes in the master group, add descriptions of volumes not currently declared in `/etc/fstab` to this file.

Example) Add `mv2` to `/etc/fstab`.

2) Create an alternative boot environment.

Example)

```
# sdxproxy Root -c Class1 -p pg1
```

3) Delete the descriptions of volumes added in step 1) from /etc/fstab.

4) Mount the alternative root volume on a temporary mount point (such as /work) and delete a description of the alternative volume of the volume added in step 1) from /work/etc/fstab.

Example) Mount pv1 on /work and delete a description of pv2 from /work/etc/fstab.

5) Unmount the alternative root volume.

Example) Unmount pv1.

D.1.9 Errors in Dell EMC Storage Units

For errors in Dell EMC storage units, in one of the following circumstances, take the actions 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 pv1
OBJ  NAME  CLASS  GROUP  DISK  VOLUME  JRM  MODE  STATUS  COPY  CURBLKS  COPYBLKS  DLY  CPTYPE  CPSOURCE
-----
slice *      Class1 Group1 Disk1 pv1    *    *    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 canceled 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 canceled and SRDF becomes no longer available.

For details, see "3.19 Using Dell EMC TimeFinder or Dell EMC SRDF in 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	emcpoweri	17596416	node1:node2	ENABLE
disk	Disk2	mirror	Class1	mg1	emcpowerj	17682084	node1:node2	ENABLE

```
# sdxinfo -D -o pg1
```

OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	DEVCONNECT	STATUS
disk	Disk3	mirror	Class1	pg1	emcpowerp	17596416	node1:node2	ENABLE

In this example, the physical disks emcpoweri and emcpowerj have been connected with the master group mg1, and the physical disks emcpowerp 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	Cap
Physical	Sym SA :P DA :IT Config	Attribute	Sts (MB)
...			
/dev/emcpoweri	000 03A:0 02B:C0 RDF1	Grp'd	RW 8592
/dev/emcpowerj	001 03A:0 02B:D0 Unprotected	N/Grp'd	RW 8632
/dev/emcpowerk	002 03A:0 01B:D0 RDF2	Grp'd	RW 8632
...			

Symmetrix ID: 000282600920

Device Name	Directors	Device	Cap
Physical	Sym SA :P DA :IT Config	Attribute	Sts (MB)
...			
/dev/emcpowerp	005 03A:0 02B:C0 RDF2	Grp'd	RW 8592
/dev/emcpowerq	006 03A:0 01B:D0 RDF1	Grp'd	RW 8632
/dev/emcpowerr	007 03A:0 01A:D0 RDF1	Grp'd	RW 8632
...			

In this example, as to the physical disks emcpoweri and emcpowerj, the Symmetrix ID is 000183600262 and the Symmetrix device names are 000 and 001 respectively. As for the physical disk emcpowerp, the Symmetrix ID is 000282600920 and the Symmetrix device name is 005.



See

For details on the symprd 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 the 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
Sym      RDF  -----
Dev  RDev Typ:G SA RA LNK Mode Dom ACp  R1 Inv  R2 Inv  RDF  S T A T E S
-----
          -----
...
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
Sym      RDF  -----
Dev  RDev Typ:G SA RA LNK Mode Dom ACp  R1 Inv  R2 Inv  RDF  S T A T E S
-----
          -----
...
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).



See

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 symprd command results in step 1-2) are indicating that the physical disk names of the Symmetrix devices 006 and 002 are emcpowerq and emcpowerk respectively. The proxy configuration associated with SRDF can be created by registering emcpowerq and emcpowerk 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
```

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 /var/opt/FJVSdx/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.

D.1.10 Cluster System Related Error

For cluster system related errors, in one of the following circumstances, take the actions 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(8) 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.
- (5) The error message "ERROR: class: cannot operate shared objects, ..." is output, and the shared class class cannot be created; or the error message "ERROR: cluster communication failure" is output, and automatic registration of resources is failed.
- (6) The disk space of a file system on a shared disk is "Full" (100%).

(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

Re-create the local class in the cluster system to use according to the following steps.

- 1) In the CF main window of Cluster Admin, execute [Stop CF] in the [Tools] menu to stop CF.
- 2) Back up volume data if necessary.
- 3) Delete the class.
- 4) In the CF main window of Cluster Admin, execute [Load driver] to start CF.
- 5) Re-create the class and volumes deleted in step 3).
- 6) Restore the volume data backed up in step 2) as needed.



.....
For information on how to operate the CF main window of Cluster Admin, see "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide."
.....

(2) The PRIMECLUSTER CF `clinitreset(8)` 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 and 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 "9.3.1 Using GDS Management View."

- For using commands, see "[Appendix B 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 by:

- 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.

- Execute the following command, if the class is registered with the resource used for the cluster application.

```
# /opt/SMAW/SMAWRrms/bin/hvgdsetup -a class_name
...
Do you want to continue with these processes ? [yes/no] 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 by:

- 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.

- Execute the following command.

```
# /opt/SMAW/SMAWRrms/bin/hvgdsetup -d class_name
...
```

```
Do you want to continue with these processes ? [yes/no] y
...
Do you need to start volumes in the specified disk class ? [yes/no]
n
```

(5) The error message "ERROR: class: cannot operate shared objects, ..." is output, and the shared class class cannot be created; or the error message "ERROR: cluster communication failure" is output, and automatic registration of resources is failed.

Explanation

This phenomenon occurs when the package of GDS is installed before installing other products like PRIMECLUSTER HA Server or PRIMECLUSTER Enterprise Edition and setting up a cluster system.

To determine whether or not your case applies to this case, check the installation dates and times of the FJSVsdx-bas and FJSVclapi packages. If the installation date and time of the FJSVsdx-bas package is earlier than that of the FJSVclapi package, this is the cause of the current trouble.



Example

Results of checking installation dates and times

- Installation date and time of the FJSVsdx-bas package

```
# rpm -qi FJSVsdx-bas
Name           : FJSVsdx-bas           Relocations: (not relocatable)
...
Install Date:  Tue Jan 26 15:27:22 2010      Build Host: xxxxxxxx
...
```

- Installation date and time of the FJSVclapi package

```
# rpm -qi FJSVclapi
Name           : FJSVclapi           Relocations: (not relocatable)
...
Install Date:  Wed Jan 27 19:08:13 2010      Build Host: xxxxxxxx
...
```

In this case, the installation date and time of the FJSVsdx-bas package is earlier than that of the FJSVclapi package. From this you can conclude that the package of GDS has already installed before installing PRIMECLUSTER HA Server, PRIMECLUSTER Enterprise Edition or other products.

Resolution

Install the FJSVsdx-bas package again, overwriting the existing one.

1. Insert CD2, used for installing PRIMECLUSTER, into the CD-ROM drive and mount it. In the following, the CD mount point is defined as <CDROM_DIR>.

```
# mount /media/cdrom
```

2. Install the FJSVsdx-bas package again, overwriting the existing one.

```
# cd <CDROM_DIR>/Linux/pkg
# rpm -Uvh -force <package name>
```

<Package name> corresponds to each distribution. Specify the name according to the correspondence chart shown below.

Distribution	Package name
RHEL7(Intel64)	FJSVsdx-bas.rhel7_x86_64.rpm
RHEL8(Intel64)	FJSVsdx-bas.rhel8_x86_64.rpm

3. Reboot the system.

```
# /sbin/shutdown -r now
```

(6) The disk space of a file system on a shared disk is "Full" (100%).

Explanation

The disk space may become "Full" (100%) while using the switchover file system created on a shared class volume.

Resolution

The recovery procedure is shown below.

1. Check the volume

On a node which is not the target for the recovery, confirm that the volume containing the target file is stopped.

Execute the following command on a node other than the recovery target.

```
# sdxinfo -V -c class
```

Example) When the class name is "c0" and the volume name is "v0"

```
# sdxinfo -V -c c0
OBJ   NAME   CLASS  GROUP  SKIP  JRM  1STBLK  LASTBLK  BLOCKS  STATUS
-----
...
volume v0      c0      g0      off  on   131072  163839  32768  STOP
...
```

Make sure that STATUS is STOP at the line which NAME is "v0."

2. Start the volume

On the target node for the recovery, execute the following command.

```
# sdxvolume -N -c class -v volume
```

Example) When the class name is "c0" and the volume name is "v0"

```
# sdxvolume -N -c c0 -v v0
```

3. Mount the file system

On the target node for the recovery, execute the following command.

Example) When the class name is "c0", the volume name is "v0", the file system type is "ext4", and the mount point is "/mnt"

```
# mount -t ext4 /dev/sfdsk/c0/dsk/v0 /mnt
```

4. Delete unnecessary files

On the target node for the recovery, delete unnecessary data under <mount_point>.

5. Unmount the file system

On the target node for the recovery, execute the following command.

Example) When the mount point is "/mnt"

```
# umount /mnt
```

6. Stop the volume

On the target node for the recovery, execute the following command.

```
# sdxvolume -F -c class -v volume
```

Example) When the class name is "c0" and the GDS volume name is "v0"

```
# sdxvolume -F -c c0 -v v0
```

7. Clear the "Faulted" state of the cluster application

Execute the following command on all nodes which compose the cluster.

```
# hvutil -c userApplication_name
```

Example) When the cluster application name is "app1"

```
# hvutil -c app1
```

8. Start the cluster application

Execute the following command on an active node.

```
# hvswitch userApplication_name SysNode
```

Example) When the cluster application of node 1 is "app1"

```
# hvswitch app1 node1RMS
```



See

For the hvutil and hvswitch commands, see the hvutil(1M) and hvswitch(1M) manual pages.

D.1.11 File System Error

For file system related errors, in one of the following circumstances, take the actions as indicated for the relevant situation.

- (1) Restoring the file system created on the local volume
- (2) Restoring the file system created on the shared volume
- (3) Restoring the file system created on the root volume

(1) Restoring the file system created on the local volume

Explanation

This is the procedure to restore the file system created on the volume of the local class by using the fsck(8) command.

Take the same procedure when using the xfs_repair(8) command to restore the xfs file system.

Resolution

If the file system is located on the local volume, how to execute the fsck(8) command is equal to when restoring the file system created on the physical slice, with following three exceptions.

- Start the volume before executing the fsck(8) command.
- When executing the fsck(8) command, specify the device special file (/dev/sfdsk/class_name/dsk/volume_name) of the GDS volume.
- The fsck(8) command cannot be executed in single user mode.

The following example shows the procedure when the class name is Class1, the volume name is Volume1, the mount point is /mount_point, and the file system type is ext4.

1. Stop applications that use the target file system.
2. Unmount the target file system if it is mounted.

```
# umount /mount_point
```

3. Start the target volume.

If the volume is already started, skip this step.

```
# sdxvolume -N -c Class1 -v Volume1
```

4. Execute the fsck(8) command.

When executing the fsck(8) command, specify the device special file of the target volume.

```
# fsck -t ext4 /dev/sfdsk/Class1/dsk/Volume1
```



For more information such as fsck(8) command options, see the manual page of fsck(8).

(2) Restoring the file system created on the shared volume

Explanation

This is the procedure to restore the file system created on the volume of the shared class by using the fsck(8) command.

Take the same procedure when using the xfs_repair(8) command to restore the xfs file system.

Resolution

If the file system is located on the shared volume, how to execute the fsck(8) command is equal to when restoring the file system created on the physical slice, with following three exceptions.

- Before executing the fsck(8) command, stop the target volume on any nodes except the node where the fsck(8) command is executed. Also, start the target volume on the node where the fsck(8) command is executed.
- When executing the fsck(8) command, specify the device special file (/dev/sfdsk/class_name/dsk/volume_name) of the GDS volume.
- The fsck(8) command cannot be executed in single user mode.

The following example shows the procedure when the class name is Class1, the volume name is Volume1, the mount point is /mount_point, and the file system type is ext4.

Take any one of the following two procedures suitable for each case:

- a. The class is not registered to the cluster application.
- b. The class is registered to the cluster application.

a. The class is not registered to the cluster application.

1. On all the nodes in the class scope, stop the applications that use the target file system.
2. On all the nodes in the class scope, unmount the target file system.

```
# umount /mount_point
```

3. On all the nodes in the class scope, stop the target volume.
Log in to any one of the nodes in the class scope, and execute the `sdxvolume -F -e allnodes` command.

```
# sdxvolume -F -c Class1 -v Volume1 -e allnodes
```

4. Start the target volume on any one of the nodes in the class scope.

```
# sdxvolume -N -c Class1 -v Volume1
```

5. Execute the `fsck(8)` command on the node where the volume has been started in step 4.
When executing the `fsck(8)` command, specify the device special file of the target volume.

```
# fsck -t ext4 /dev/sfdsk/Class1/dsk/Volume1
```



For more information such as `fsck(8)` command options, see the manual page of `fsck(8)`.

b. The class is registered to the cluster application.

1. Stop the applications that use the target class.



For how to stop the cluster application, see "PRIMECLUSTER Installation and Administration Guide."

2. Start the target volume on any one of the nodes in the class scope.

```
# sdxvolume -N -c Class1 -v Volume1
```

3. Execute the `fsck(8)` command on the node where the volume has been started in step 2.
When executing the `fsck(8)` command, specify the device special file of the target volume.

```
# fsck -t ext4 /dev/sfdsk/Class1/dsk/Volume1
```



For more information such as `fsck(8)` command options, see the manual page of `fsck(8)`.

4. Stop the target volume on the node where the volume has been started in step 2.

```
# sdxvolume -F -c Class1 -v Volume1
```

(3) Restoring the file system created on the root volume

Explanation

This is the procedure to restore the file system created on the volume of the root class by using the fsck(8) command.

Take the same procedure when using the xfs_repair(8) command to restore the xfs file system.

Resolution

a. Any file systems except / (root), /usr, and /var

For file systems except / (root), /usr, and /var, how to restore the file system is equal to when restoring the file system created on the local class volume.

For details on the procedure, see "[\(1\) Restoring the file system created on the local volume.](#)"

b. / (root), /usr, and /var file systems

The / (root), /usr, and /var file systems cannot be unmounted during system operation. In this case, to execute the fsck(8) command, boot the system from the CD-ROM drive. Boot the system from the CD-ROM drive, execute the fsck(8) command for one of the mirrored physical slices, and then copy the data from that physical slice to another physical slice.

When taking this procedure to restore the file system created on the volume of the root class, check the following items when the system disk is registered to the root class, and write down them for later reference.

- Physical disk information of the system volume
- Slice number of the system volume

For details, see "[6.5.4 Checking Physical Disk Information and Slice Numbers.](#)"

1. Turn on the server, and insert the installation CD into the CD-ROM drive.
2. From boot devices displayed in the boot option selection screen of the EFI boot manager, select the CD-ROM device, and boot the system in rescue mode. For details, see the OS manual.



Do not mount a file system when booting in rescue mode in an environment where the system disk is mirrored.

3. Check the physical slice name of the target file system to be restored among the / (root), /usr, and /var file systems.

For the grep command arguments, specify the physical disk information confirmed in "[6.5.4 Checking Physical Disk Information and Slice Numbers.](#)"

```
# ls -l /sys/block/sd* | grep 0000:06:02:0 | grep 0:0:0
                                (*1)                (*1)
lrwxrwxrwx 1 root root 0 Jun  1  2005 /sys/block/sda ->\
                                                (*2)
../../../../devices/pci0000:02/0000:02:1f.0/0000:06:02.0/host2/\
target0:0:0/0:0:0/0:0:0/block/sda

# ls -l /sys/block/sd* | grep 0000:06:02:0 | grep 0:2:0
                                (*1)                (*1)
lrwxrwxrwx 1 root root 0 Jun  1  2005 /sys/block/sdb ->\
                                                (*2)
../../../../devices/pci0000:02/0000:02:1f.0/0000:06:02.0/host2/\
target0:0:2/0:0:2/0:2:0/block/sdb
```

(*1) Physical disk information

(*2) Physical disk name

In this example, the physical disk names are sda and sdb.

By combining the physical disk name and the slice numbers confirmed in "[6.5.4 Checking Physical Disk Information and Slice Numbers.](#)" you can get the physical slice names.

In this example, the physical slice names are as follows.

Use	Physical slice name	
/	sda1	sdb1
/var	sda2	sdb2
/usr	sda3	sdb3

- Execute the fsck(8) command for one of the physical slices to restore the file system.

In this example, execute the fsck(8) command for the slice of the root file system (/) checked in step 3.

```
# fsck -t ext4 /dev/sda1
```

If the file system cannot be restored, execute the fsck(8) command for another physical slice.



For more information such as fsck(8) command options, see the manual page of fsck(8).

- From the physical slice where the file system is restored in step 4., copy the data to another physical slice.

```
# dd if=/dev/sda1 of=/dev/sdb1 bs=1M
```



If the mirroring multiplicity is n, among n mirrored slices, copy data to all "n-1" slices other than the slice restored by the fsck(8) command in step 4.



Specify a correct option referring to the dd(1) command manual.

- Exit the rescue mode, and start the system.

```
# exit
```

D.1.12 Shadow Object Errors

For shadow object related errors, in one of the following circumstances, take the 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.9 Shadow Volume](#)," and re-create a shadow volume.

D.1.13 Disk Unit Error



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/log/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 "[D.1.1 Slice Status Abnormality](#)."



For the procedures for swapping disks, see "[1.2.3 Hot Swap](#)," "[7.3.1.2 Operation Procedure](#)," and "[B.1.8 sdxswap - Swap Disk](#)."

D.1.14 OS Messages

For OS messages, in one of the following circumstances, take the actions as indicated for the relevant situation.

- (1) Messages are output at node startup.
- (2) A SCSI ioctl related message is output at node startup or command execution.
- (3) Stack information is output at node startup and command execution.
- (4) A devlabel message is output at node startup or during collection of investigation material.
- (5) A message about the `sfdsksys` module is output when the errata kernel is applied.
- (6) Messages are output at node startup, and GDS cannot be started.
- (7) A message is output at node startup, and the system cannot be started.
- (8) A message about the ext4 file system is output at node stop.
- (9) Messages of `systemd` or `systemd-udev` are output when the system is started or the operations of GDS are executed.
- (10) Messages of `multipathd` are output when creating volumes.
- (11) The message about SELinux is output.
- (12) The device name of GDS volume is displayed in the `sfdskX` format, therefore the message for which volume is unknown.
- (13) A message is output from the standby node when the node is restarted, and the cluster cannot be started.
- (14) When the stop operation of the GDS volume is executed, an error message is output, and the stop operation of the GDS volume fails.

(1) Messages are output at node startup.

Explanation

At node startup, the following message may be output.

sfdsksys: no version for "sfdskd_event_exist" found: kernel tainted.

Some or all of the following messages may be output to the console and the syslog during system startup. *XXXXXXXXXX* indicates 40 digit hexadecimal.

kernel: Request for unknown module key 'FUJITSU Software: Fujitsu BIOS DB FJMW Certificate: *XXXXXXXXXX*' err -11

kernel: sfdsk_lib: loading out-of-tree module taints kernel.

kernel: sfdsk_lib: module verification failed: signature and/or required key missing - tainting kernel

Resolution

Even if this message is output, the system is not affected. There is no need to work around.

(2) A SCSI ioctl related message is output at node startup or command execution.

Explanation

The message "program sdxservd is using a deprecated SCSI ioctl, please convert it to SG_IO" may be output when:

- The node is started.
- A disk is registered with a class.
- Physical disk swap (sdxswap -O command) is performed.
- Physical disk restoration (sdxswap -I command) is performed.
- Investigation material collection (pclsnap or sdxsnap.sh command) is performed.

Resolution

Even if this message is output, the system is not affected. There is no need to work around.

(3) Stack information is output at node startup and command execution.

Explanation

Stack information starting with the "Call Trace:" message may be output on the console and the system log when:

- A volume is removed (sdxvolume -R command).
- A slice is attached again (sdxslice -R command).
- A system is started.

Resolution

Even if this message is output, the system is not affected. There is no need to work around.

(4) A devlabel message is output at node startup or during collection of investigation material.

Explanation

At node startup or during collection of investigation material (when executing pclsnap, sdxsnap.sh, or FJQSS), the following message may be output.

devlabel: devlabel's temporary ignore list /etc/sysconfig/devlabel.d/ignore_list has been emptied due to a change in device configuration.

devlabel: The device /dev/sdXX is being put in devlabel's temporary ignore list /etc/sysconfig/devlabel.d/ignore_list to avoid errors.

devlabel: The device *old_device* is now known as *new_device*.`/etc/opt/FJVSvdx/.devlabel/class/sdx_dev_999_999` now points to the new name.

devlabel: The symlink `/etc/opt/FJVSvdx/.devlabel/class/sdx_dev_999_999-> device` is being ignored in `/etc/sysconfig/devlabel` because the correct device cannot be found.

devlabel: devlabel service started/restarted

Resolution

Even if this message is output, the system is not affected. There is no need to work around.

(5) A message about the sfdsksys module is output when the errata kernel is applied.

Explanation

If an errata kernel is applied in an environment where the system disk is mirrored, the following OS error message may be output on the console.

```
Failed to install module sfdsksys
```

Resolution

Execute the following command to check if the installed errata kernel version is displayed.

```
# rpm -q kernel
```

When the installed errata kernel version is displayed correctly, the errata kernel is properly applied. No corrective action is required.

If the installed errata kernel version is not displayed, contact field engineers.

(6) Messages are output at node startup, and GDS cannot be started.

Explanation

If Secure Boot is enabled in the environment that is not capable of Secure Boot, the following messages may be output to the console or the system log. In this case, GDS may not be started. *XXXXXXXX* indicates 40 digit hexadecimal.

```
kernel: Request for unknown module key 'FUJITSU Software: Fujitsu BIOS DB FJMW Certificate: XXXXXXXX' err -11
bash: modprobe: ERROR: could not insert 'sfdsk_lib': Required key not available
bash: modprobe: ERROR: could not insert 'sfdsksys': Required key not available
bash: modprobe: ERROR: could not insert 'sfdsk': Required key not available
systemd: Failed to start PRIMECLUSTER Global Disk Services (fjvsdx) service.
```

Resolution

If Secure Boot is enabled in the environment that is not capable of Secure Boot, disable it.

(7) A message is output at node startup, and the system cannot be started.

Explanation

If system disk mirroring is set in the environment that is not capable of Secure Boot and Secure Boot is enabled, the following message may be output to the console or the system log. In this case, the system may not be started.

```
bash: modprobe: ERROR: could not insert 'sfdsksys': Required key not available
```

Resolution

If Secure Boot is enabled in the environment that is not capable of Secure Boot, disable it.

(8) A message about the ext4 file system is output at node stop.

Explanation

At node stop, the following message may be output to the console.

```
end_request: I/O error, dev dev, sector NNNNNNNNEXT4-fs error (device device):  
ext4_find_entry: reading directory #inode offset N
```

dev is the device name of the system disk. *N* is a number. *device* is the mount device of the file system that includes the /etc/opt directory. *inode* is the inode number of the directory where an error occurred.

Resolution

Execute the following command to check the inode number of the /etc/opt/FJVSdx directory.

```
# ls -ldi /etc/opt/FJVSdx  
41391133 drwxr-xr-x. 4 root root 4096 Oct  2 14:03 /etc/opt/FJVSdx  
inode number
```

If the inode number is consistent with *inode*, the output of this message means that the sdxmond daemon of GDS is referring a file in the /etc/opt/FJVSdx directory at node stop. In this case, the system is not affected by this status. No resolution is required.

If the inode number is not consistent with *inode*, this message is unlikely to be considered as a message regarding GDS error. Contact field engineers.

(9) Messages of systemd or systemd-udev are output when the system is started or the operations of GDS are executed.

Explanation

When the system is started or the operations of GDS are executed, the following messages of systemd or systemd-udev are output to the console and the /var/log/messages file.

(a) systemd-udev: error: device: Device or resource busy

device is the physical disk of the disk that is registered in the shared class or the local class (example: /dev/sdf).

This message is output in RHEL7.2 or later system when:

- A disk is registered to the shared class or the local class.
- The system is started.

(b) systemd-udev: error: device: Device or resource busy

device is the partition on the physical disk of the disk that is registered in the shared class or the local class (example: /dev/sdf2).

This message is output in RHEL7.2 or later system when:

- A volume of the shared class or the local class is created.
- The system is started.

(c) systemd-udev: error: device: No such device or address

device is the device of the GDS volume (example: /dev/sfdsk32, /dev/sfdsk/gdssys2).

This message is output in RHEL7.2 or later system when:

- A volume of any class type is created.
- The system is started.

(d) systemd-udev: inotify_add_watch(N, device, M) failed: No such file or directory

device is the partition for a volume of the root class (example: /dev/sdd5).

N and *M* indicate numerical values.

This message is output when:

- "System Disk Settings" is executed.
- "Unmirror System Disk" is executed.
- "Swap Physical Disk" is executed for the system disk.
- "Restore Physical Disk" is executed for the system disk.

(e) systemd: Device device appeared twice with different sysfs paths device1 and device2

device, *device1*, and *device2* are the text strings that indicate the partitions.

device1 and *device2* are the partitions for the volume, or the volume devices.

[Example 1]

```
systemd: Device dev-disk-by\x2duuid-0028c72b\x2dad94\x2d4b27\x2d912b\x2dbe389de83cd2.device
appeared twice with different sysfs paths
/sys/devices/pci0000:00/0000:00:10.0/host2/target2:0:0/2:0:0:0/block/sda/sda2 and
/sys/devices/pci0000:00/0000:00:10.0/host2/target2:0:1/2:0:1:0/block/sdb/sdb2
```

[Example 2]

```
systemd: Device dev-disk-by\x2duuid-75C6\x2d77DC.device appeared twice with different sysfs paths
/sys/devices/virtual/block/sfdsk!gdssys6 and
/sys/devices/pci0000:00/0000:00:15.0/0000:03:00.0/host2/target2:0:1/2:0:1:0/block/sdb/sdb1
```

This message is output in RHEL7.2 or later system when:

- The system is started.

Resolution

Even if this message is output, the system is not affected. No resolution is required.

(10) Messages of multipathd are output when creating volumes.

Explanation

When creating volumes, the following messages may be output to the console and the system log. *X* is numerical value.

- For RHEL7
 - multipathd: sfdsk*X*: add path (uevent)
 - multipathd: sfdsk*X*: failed to get path uid
 - multipathd: uevent trigger error
- For RHEL8
 - multipathd: sfdsk*X*: triggering change event to reinitialize
 - multipathd: sfdsk*X*: failed to get udev uid: Invalid argument
 - multipathd: failed to get path uid
 - multipathd: uevent trigger error

Resolution

Even if this message is output, the system is not affected. No resolution is required.

For inhibiting the message output, add the following description in the black list of the DM-MP setting file "/etc/multipath.conf."

```
devnode "^sfdsk[0-9]+$"
```

A description example:

```
blacklist {
    devnode "^sfdsk[0-9]+$"
}
```

(11) The message about SELinux is output.

Explanation

If the mirroring among servers is used, the following message about SELinux may be output after a node is restarted or a network error is restored.

```
setroubleshoot: failed to retrieve rpm info for /var/opt/FJVSdx/log/file_name
```

```
setroubleshoot: Plugin Exception restorecon_source
```

```
setroubleshoot: SELinux is preventing /usr/sbin/iscsiadm from read access on the file /var/opt/FJVSdx/log/file_name. For complete SELinux messages. run sealert -l ID
```

```
python: SELinux is preventing /usr/sbin/iscsiadm from read access on the file /var/opt/FJVSdx/log/file_name.
```

Resolution

Even if this message is output, the system is not affected. No resolution is required.

(12) The device name of GDS volume is displayed in the sfdskX format, therefore the message for which volume is unknown.

Explanation

The device name of GDS volume may be displayed in the sfdskX format for the OS message.

Example) kernel: XFS (sfdsk33): Internal error ...

Resolution

The GDS volume corresponding to the sfdskX format device displayed for the OS message can be checked by the following procedure.

1. Check the major and minor numbers of sfdskX format device.

In the following example, the sfdsk33 major number is 487 and minor number is 33.

Example) How to check the sfdsk33 major and minor numbers

```
# ls -l /dev/sfdsk33
brw-rw----. 1 root disk 487, 33 Feb 21 16:32 /dev/sfdsk33
```

2. Search for the volume coincident with the major and minor numbers checked in step 1.

In the following example, the volume in which the major number is 487 and the minor number is 33 is /dev/sfdsk/class0001/dsk/volume2. Therefore, the volume corresponding to sfdsk33 is prove to be /dev/sfdsk/class0001/dsk/volume2.

Example)

```
# ls -l /dev/sfdsk/*/dsk
/dev/sfdsk/class0001/dsk:
total 0
brw----- . 1 root root 487, 32 Feb 21 16:31 volume1
brw----- . 1 root root 487, 33 Feb 21 16:32 volume2
brw----- . 1 root root 487, 34 Feb 21 16:32 volume3
...
```

(13) A message is output from the standby node when the node is restarted, and the cluster cannot be started.

Explanation

When you reboot a node in a VMWare environment that uses I/O fencing, the cluster on the standby node may not start properly with the following message in the system log on the standby node.

kernel: sd X:X:X:X: reservation conflict

kernel: WARNING: sfdsk: write error on disk:

Resolution

Make sure that you have configured the I/O fencing feature of the GDS.

For details of the settings, refer to "H.2.1.1 Installation and Configuration of Related Software" described in "PRIMECLUSTER Installation and Administration Guide."

(14) When the stop operation of the GDS volume is executed, an error message is output, and the stop operation of the GDS volume fails.

Explanation

In a RHEL 8.1 or later environment, if Device Mapper Multipath (DM-MP) is registered in the GDS class, the following message is output to the console when a GDS volume is stopped, and stopping a GDS volume may fail.

```
ERROR: volume: device busy on node node
```

In the above case, one of the following messages is output to the system log. X is a number.

```
multipathd: sfdskX: triggering change event to reinitialize
```

```
multipathd: sfdskX: failed to get udev uid: Invalid argument
```

```
multipathd: failed to get path uid
```

```
multipathd: uevent trigger error
```

```
multipathd: sfdskX: unusable path (wild) - checker failed
```

Resolution

The following procedure must be performed to avoid failure.

This procedure can be performed during the service.

1. Add the following description in the black list of the DM-MP setting file "/etc/multipath.conf".

[A description example]

```
blacklist {
    devnode "^sfdsk[0-9]+$"
}
```

2. Reload the multipathd daemon.

```
# systemctl reload multipathd.service
```

D.1.15 Other Errors

If the following error occurs, take the action described below.

- (1) [The fjsvsdx.service becomes "failed" at node startup.](#)

(1) The fjsvsdx.service becomes "failed" at node startup.

Explanation

If Secure Boot is enabled in the environment that is not capable of Secure Boot, the fjsvsdx.service may become "failed" at node startup.

Resolution

If Secure Boot is enabled in the environment that is not capable of Secure Boot, disable it.

D.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 field engineers.

- Investigation material collected by the `pclsnap` command, `sdxsnap.sh`, or FJQSS.
 - Use FJQSS if you want to collect the investigation material with other middleware products that support FJQSS.
 - If investigation material can be collected by neither the `pclsnap` command nor FJQSS for any reason such as customer requirements, collect the only essential investigation material by `sdxsnap.sh`.
 - In cases other than the above, collect investigation material by the `pclsnap` command.
- Crash dump

For troubles related to the Management View, collect investigation material of Web-Based Admin View. For the method of collecting investigation material, see "PRIMECLUSTER Web-Based Admin View Operation Guide."



- 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.
- If the node in which the failure occurred can be restarted, please collect the crash dump manually before you restart the node. The crash dump is useful if the failure originates from the OS.
- When collecting investigation material, among the by-id files (symbolic link files created under the `/dev/disk/by-id` directory) managed by the `udev` function of Linux, the by-id files of disks which are not registered in classes may not exist temporarily because OS creates the by-id files again triggered by the GDS processes for obtaining the device information.
When access to the by-id files of disks which are not registered in classes fails, retry access.
Alternatively, when collecting investigation material, do not access to the by-id files of disks which are not registered in classes.

D.2.1 Collecting Investigation Material with `pclsnap` Command

The `pclsnap` command, a Linux system information collection tool, is provided in the `FJSVpclsnap` package. When a problem occurs with GDS, this command allows you to collect investigation material to identify the cause.

How to execute the `pclsnap` command is as follows.

1. Log in with superuser authority.
2. Execute the `pclsnap` command as follows.

```
# /opt/FJSVpclsnap/bin/pclsnap -h output
```

or

```
# /opt/FJSVpclsnap/bin/pclsnap -a output
```

output indicates either the special file name (example: `/dev/st0`) of the output media which stores the investigation material collected with the `pclsnap` command, or the output file name.

D.2.2 Collecting Investigation Material with `sdxsnap.sh` Script

You can only collect the essential investigation material of GDS with the `sdxsnap.sh` script. If investigation material cannot be collected by the `pclsnap` command and FJQSS for any reason such as customer requirements, use this script.

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 the `snap.tar` file to field engineers.

D.2.3 Collecting Investigation Material with FJQSS

By using FJQSS (material collecting tool), the same information can be collected when using `sdxsnap.sh -s`.

FJQSS is automatically installed if this software is installed.

How to collect the material by using FJQSS is as follows.

For details, see the manual of FJQSS.

1. Log in with superuser authority.
2. Execute the `fjqss_collect` command.

```
# /opt/FJSVqst1/fjqss_collect outdir
```

For `outdir`, specify the path of the output directory of the material.

If `outdir` is omitted, the material is output in the `/var/tmp`.

3. On the screen to select the product, select PRIMECLUSTER GD.
For the cluster system, when selecting the cluster software of PRIMECLUSTER, all the investigation materials including GDS can be collected at once.

After the materials are collected, the output directory of the investigation material is displayed.

Send the material collected in the displayed output directory to field engineers.

Appendix E Frequently Asked Questions (FAQ)

In this section, frequently asked questions (FAQ) have been compiled.

E.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: Can NAS be mirrored?
- Q: Can system disk mirroring be performed in a multi-boot environment? [EFI]

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 "[3.1.1 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 "[2.3.12 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 "[7.3.1.2 Operation Procedure](#)" and "[B.1.8 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 EFI. For details, see "[D.1.5 System Disk Abnormality \[EFI\]](#)."

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 "[C.2 Driver Messages](#)"
- HALT, ERROR, or WARNING level messages and Internal error messages shown in "[C.3 Daemon Messages](#)"

For details, see "[Appendix C 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.
sdxhealthcheckd	/etc/opt/FJVSdx/bin/sdxhealthcheckd	Monitors GDS daemons and driver hangup (only when GDS I/O Monitor Option is installed) [IOmonitor]

GDS: Global Disk Services

If any resident daemon other than sdxmond and sdxhealthcheckd 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 daemon, osfunc=osfunc,  
errno=errno
```

If sdxmond and sdxhealthcheckd terminates abnormally, the OS SYSTEMD(1) process automatically restarts the daemon.

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 "[C.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.

Q: Can NAS be mirrored?

A

NAS cannot be mirrored.

The disk device managed by GDS is the one that can be seen from the server where GDS operates just like the internal disk.

NAS such as ETERNUS NR1000F series is not supported because it can be seen as a file server from the server where GDS operates.

Q: Can system disk mirroring be performed in a multi-boot environment? [EFI]

A

No. System disk mirroring is not supported in multi-boot environments.

E.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? \[EFI\]](#)
- [Q: How do I make a backup of a disk \(other than a system disk\) that is managed with GDS?](#)
- [Q: I cannot execute the command parted\(8\) on a physical disk managed by GDS \(e.g. sda\). How can I solve this problem?](#)

Q: How do I make a backup of a system disk that is managed with GDS? [EFI]

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 "[7.4 Backing Up and Restoring System Disk](#)" and "[7.5 Backing Up and Restoring System Disk through Alternative Boot Environment \[EFI\]](#)."

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 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](#)" "[7.9 Backup \(by OPC\)](#)" and "[B.1.5 sdxslice - Slice Operations](#)."

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](#)," "[3.15 Online Backup and Restoration Using Proxy Volume](#)," "[7.7 Backup \(by Slice Detachment\)](#)," and "[7.8 Backup \(by Synchronization\)](#)."

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 "[3.20 Backing Up and Restoring through External Server](#)."

Q: I cannot execute the command parted(8) on a physical disk managed by GDS (e.g. sda). 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.1 Restraining Access to Physical Special File](#)."

If all you need to do with the commands parted(8) is to refer to the disk contents, you may do so by booting from the CD-ROM, 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 Booting from a CD-ROM Device](#)."

E.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 cannot boot the system. How do I restore it? \[EFI\]](#)
- [Q: When an I/O error status is created in the test, how do I shorten the restoration time?](#)
- [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 "[B.1.10 sdxcopy - Synchronization Copying Operation](#)," and "[B.1.6 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/fstab`, as necessary.
8. Restore the backup data you collected in procedure 2.

Q: I cannot boot the system. How do I restore it? [EFI]

A

See "[D.1.5 System Disk Abnormality \[EFI\]](#)"

Q: When an I/O error status is created in the test, how do I shorten the restoration time?

A

The `sdxfix -D` command is used for reading check for the restoration target disk, the execution time may be long corresponding to the disk capacity.

If you specify the `-x NoRdchk` option, the reading check will be skipped and shorten the execution time. The `-x NoRdchk` option is only available when the reading check is obviously not needed such as an I/O error status is created. For details, see "[B.1.9 sdxfix - Restore Failed Object](#)."

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 "[1.3.5.4 Lock Volume](#)" and "[B.1.7 sdxattr - Change 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.

Appendix F Systemd Services and Startup Daemons in GDS

This appendix provides explanations on systemd services and startup daemons for GDS.

F.1 Explanation Formats

Scripts and daemons are explained with the following formats:

Name of Unit

Name of Unit.

Function

Function of Unit.

Effect if stopped

Effect if Unit is stopped.

Dependence with other Units

Requires

Prerequisite Units needed by this Unit. If the Units listed here fail to start, this Unit will not be started.

Wants

Prerequisite Units needed by this Unit. If the Units listed here fail to start, this Unit will be started.

Before

Other Units started after this Unit.

After

Other Units started before this Unit.

Startup daemon

Daemon started by Unit.

Remarks

Remarks.

F.2 systemd Service Lists

fjvsdx.service

Function

Basic part of GDS.

Effect if stopped

GDS functions cannot be used.

Dependence with other Units

Requires

None.

Wants

None.

Before

fjsvclctrl.service
fjsvsdx2.service

After

iscsi.service
iscsi-shutdown.service
target.service (*1)

Startup daemon

/usr/sbin/sdxservd
/usr/sbin/sdxlogd
/usr/sbin/sdxexd

Remarks

(*1) Has dependency with target.service only when the mirroring among servers is used.

fjsvsdx2.service

Function

Basic part of GDS.

Effect if stopped

GDS functions cannot be used.

Dependence with other Units

Requires

None.

Wants

None.

Before

fjsvsdxmon.service

After

fjsvsdx.service
fjsvclctrl.service

Startup daemon

None.

Remarks

None.

fjsvsdxmon.service

Function

Monitoring GDS.

Effect if stopped

GDS cannot be restarted when it ends abnormally.

Dependence with other Units

Requires

None.

Wants

None.

Before

None.

After

fjvsdx2.service

Startup daemon

/usr/sbin/sdxmond

Remarks

None.

fjvsdxhealth.service

Function

Hangup monitor of GDS daemon and driver.

Effect if stopped

When GDS daemon or driver hangs-up, system panic does not occur.

Dependence with other Units

Requires

None.

Wants

None.

Before

fjvsdx2.service

After

fjvsdx.service
fjvsctrl.service

Startup daemon

/etc/opt/FJSVsdx/bin/sdxhealthcheckd

Remarks

If GDS I/O Monitor Option has been installed, this daemon will be started. Determine if the GDS I/O Monitor Option is installed or not with the existence of FJSVsdx-iom package.

Appendix G Release Information

G.1 New Features

This section describes primary new features of this software.

Some of the new features are restricted. For restrictions, resolutions, and restriction removal schedules, see "Restrictions" in the software release guide.

G.1.1 New Features in 4.3A00

No.	Products	Feature	Description	For details
1	GDS, GDS Snapshot	Support for local and shared classes on a guest OS in a virtual environment	Local and shared classes can be created on the guest OS in a virtual environment.	"A.3.5 Guidelines for GDS Operation on the Virtual Machine Function"
2	GDS Snapshot	Two generation proxies for system volumes	Two proxy volumes can be related to one system volume.	
3	GDS	Shortening the restoration time of the I/O error status of a disk	When clearing the I/O error status of a disk with the <code>sdxfix -D</code> command, the processing time can be shortened by specifying the <code>-x NoRdchk</code> option.	- "When an I/O error status is created in the test, how do I shorten the restoration time?" in "G.3 Operation" - "D.9 sdxfix - Restore Failed Object"
4	GDS, GDS Snapshot	DM-MP is supported	In local and shared classes, <code>mpath</code> devices of DM-MP can be managed (Patch is required).	- "A.2.41 DM-MP (Device Mapper Multipath)" - "A.2.2 Restraining Access to Physical Special File"

GDS: Global Disk Services

GDS Snapshot: Global Disk Services Snapshot

G.1.2 New Features in 4.3A10

No.	Products	Feature	Description	For details
1	GDS, GDS Snapshot	Platform support extension	The following platforms are newly supported: - RHEL6 (x86) - RHEL6 (Intel64)	
2	GDS	1 TB or larger disks are supported	In local and shared classes, 1 TB or larger disks can be managed (Patch is required).	- "A.2.41 DM-MP (Device Mapper Multipath)" - "A.2.2 Restraining Access to Physical Special File"

GDS: Global Disk Services

GDS Snapshot: Global Disk Services Snapshot

G.1.3 New Features in 4.3A20

No.	Products	Feature	Description	For details
1	GDS, GDS Snapshot	The IPv6 environment is supported	GDS and GDS snapshot can be used in the IPv6 environment.	

GDS: Global Disk Services

GDS Snapshot: Global Disk Services Snapshot

G.1.4 New Features in 4.3A30

No.	Products	Feature	Description	For details
1	GDS	System disk mirroring on a VMware guest	On a VMware guest, system disk can be mirrored in UEFI boot environments with RHEL6 (Intel64) or later.	- "1.2.1.1 System Disk Mirroring [EFI]" - "K.1 Guidelines for GDS Operation in the Virtual Environment"
2	GDS, GDS Snapshot	Virtual disk (VMDK) on a VMware guest is supported	On a VMware guest, virtual disk (VMDK) can be managed in a root class and a local class.	- "K.1 Guidelines for GDS Operation in the Virtual Environment" - "K.4 Use of GDS in a VMware Environment"
3	GDS	virtio-SCSI device is supported	On a KVM guest, virtio-SCSI device can be managed in a local class and a shared class.	- "K.1 Guidelines for GDS Operation in the Virtual Environment " - "K.3 Use of GDS in a KVM Environment [4.3A10 or later] "
4	GDS, GDS Snapshot	Cancellation of restriction in local and shared classes of GPT label	The following functions are supported in local and shared classes of GPT label. - Mirroring - Concatenation - Striping - sdxconfig command - Changing the physical slice attribute value of volume - GDS Snapshot	"A.2.5 Disk Label"
5	GDS, GDS Snapshot	Mirroring and snapshot for 1 TB or larger volumes	Mirroring and snapshot creation are available for 1 TB or larger volumes.	"A.2.7 Volume Size"
6	GDS Snapshot	Snapshot creation with QuickOPC of ETERNUS Disk storage system	Snapshot can be created with the Advanced Copy function of ETERNUS Disk storage system such as QuickOPC.	- "1.5.3 Instant Snapshot by OPC" - "5.1.3.3 Backup (by OPC)" - "5.3.2.3 Backup (by OPC)" - "6.4 Online Backup and Instant Restore through Proxy Volume " - "D.15 sdxproxy - Proxy Object Operations"
7	GDS Snapshot	Snapshot creation with ETERNUS Disk	Snapshot can be created with ETERNUS Disk storage system	"A.2.17 Using Advanced Copy Function in Proxy Configuration"

No.	Products	Feature	Description	For details
		storage system REC function	Remote Advanced Copy function (REC) in other disk cases.	
8	GDS Snapshot	Snapshot creation with ETERNUS Disk storage system Advanced Copy function in the DM-MP environments	Snapshot can be created with ETERNUS Disk storage system Advanced Copy function in DM-MP environment.	- "A.2.17 Using Advanced Copy Function in Proxy Configuration" - "A.2.40 DM-MP (Device Mapper Multipath)"
9	GDS	Improvement on Change Scope screen of GDS Management View	When the following names; a node identifier of PRIMECLUSTER (CF node name) and a host name (mip) of Web-Based Admin View displayed on the GDS Configuration Tree field in the Main screen in GDS Management View, are different, both a node identifier name and a hostname are displayed in the [Node Name] of the [Change Scope] screen.	"5.2.2.2 Cluster System Class Configuration"

GDS: Global Disk Services

GDS Snapshot: Global Disk Services Snapshot

G.1.5 New Features in 4.3A40

No.	Products	Feature	Description	For details
1	GDS	RHEL7 is supported	GDS and GDS Snapshot can be used in RHEL7(Intel64) environment.	
2	GDS	Mirroring among servers	In a RHEL6 environment, the local disk (such as the internal disk) of each node in the cluster system can be mirrored via the network. In the cluster application of the standby operation, the mirroring among servers can be used to take ver the data between the nodes.	"1.2.1.4 Mirroring among Servers" "3.9 Mirroring among Servers" "4.8 Disk Setting for Performing Mirroring among Servers" "6.1.4 To Use Mirroring among Servers" "7.3.3 Swapping Disks of Netmirror Group" "7.16 Mirroring Among Servers" "D.1.6 Errors in the Mirroring Among Servers"
3	GDS	FJQSS support	By FJQSS (material collecting tool), the investigation material required to investigate GDS can be collected.	"D.2 Collecting Investigation Material" "D.2.3 Collecting Investigation Material with FJQSS"

GDS: Global Disk Services

GDS Snapshot: Global Disk Services Snapshot

G.1.6 New Features in 4.4A00

No.	Products	Feature	Description	For details
1	GDS	Extending the supported configuration for the mirroring among servers	<p>The mirroring among servers can be used in the following environments.</p> <ul style="list-style-type: none"> - RHEL7 environment (RHEL7.1 or later) - VMware environment - Cloud environment - Cluster system connecting to external disk units <p>The following disks can be used in the mirroring among servers.</p> <ul style="list-style-type: none"> - Disks that configure RAID by RAID card - External disk units in a virtual environment (such as ETERNUS) 	<p>"3.9.2 Disk"</p> <p>"3.9.4 Configuration and Settings of the Cluster System"</p> <p>"3.21 Virtual Environment"</p>
2	GDS	Hot swap of disks used in the configuration of mirroring among servers	Hot swap of disks used in the configuration of mirroring among servers is available. (swap without stopping the system)	<p>"7.3.3.1 Hot swap (RHEL6)"</p> <p>"7.3.3.1 Hot Swap"</p>
3	GDS	Supporting a virtual disk (VMDK) as a shared disk for a cluster system in a VMware environment	A virtual disk (VMDK) as a shared disk for a cluster system in a VMware environment is available.	"3.21 Virtual Environment"

GDS: Global Disk Services

G.1.7 New Features in 4.5A00

No.	Products	Feature	Description	For details
1	GDS	PRIMEQUEST 3000 Series support	GDS is available with PRIMEQUEST 3000 Series	-
2	GDS	Disk support for VMware vSAN	Disk of VMware vSAN can be used to register to class.	"3.21 Virtual Environment"
3	GDS I/O Monitor Option	I/O Response Time Assurance [IOmonitor]	Assurance of I/O response within a fixed time.	<p>"1.7 GDS I/O Monitor Option Features"</p> <p>"1.8 I/O Response Time Assurance [IOmonitor]"</p> <p>"2.1.1 Disk Class"</p> <p>"3.23 I/O Response Time Assurance [IOmonitor]"</p> <p>"6.12 I/O Response Time Assurance [IOmonitor]"</p> <p>"7.15 I/O Response Time Assurance [IOmonitor]"</p> <p>"8.3 Changing the Settings of I/O Response Time Assurance [IOmonitor]"</p> <p>"8.4 Deleting the Settings of the I/O Response Time"</p>

No.	Products	Feature	Description	For details
				Assurance [IOmonitor] "B.1.2 sdxdisk - Disk Operations" "B.1.6 sdxinfo - Display Object Configuration and Status Information" "B.1.7 sdxattr - Change Objects Attributes"
4	GDS	Java Web Start support for GDS Management View	The screen can be started using Java Web Start command without using the Web browser.	"5.1.1 Starting Web-Based Admin View Operation Menu"

GDS: Global Disk Services

GDS I/O Monitor Option: Global Disk Services I/O Monitor Option

G.1.8 New Features in 4.5A10

There is no new feature in the version 4.5A10 of this software.

G.1.9 New Features in 4.6A00

No.	Products	Feature	Description	For details
1	GDS	Enhanced availability of the mirroring among servers function	Even when a node becomes the LEFTCLUSTER state due to a failure of servers or interconnects, I/O to the netmirror volume can be continued.	7.16.3 Operation when the LEFTCLUSTER Node Exists

GDS: Global Disk Services

G.1.10 New Features in 4.6A10

There is no new feature in the version 4.6A10 of this software.

G.1.11 New Features in 4.6A20

There is no new feature in the version 4.6A20 of this software.

G.2 Manual Changes

This section describes primary changes in this manual.

G.2.1 4.3A00 Manual Changes

No.	Product	Manual code	Description	Chapter/Section
1	GDS	J2UZ-7243-01ENZ0(A) J2UZ-7243-01ENZ2(A)	Added a note for mirroring the system disk in a SAN boot environment.	"1.2.1.1 System Disk Mirroring [PRIMEQUEST] "
2	GDS	J2UZ-7243-01ENZ0(A) J2UZ-7243-01ENZ2(A)	Added /opt to file systems that cannot be used as volumes in local and shared classes.	"A.2.1 Managing System Disks"
3	GDS	J2UZ-7243-01ENZ0(A) J2UZ-7243-01ENZ2(A)	Added a condition "(SUSE Linux Enterprise Server 10 or later)" in "Export to NFS Clients (SUSE	"A.2.32 NFS Mount"

No.	Product	Manual code	Description	Chapter/Section
			Linux Enterprise Server 10 or later)."	
4	GDS	J2UZ-7243-01ENZ0(A) J2UZ-7243-01ENZ2(A)	Added "Guidelines for GDS Operation on the Virtual Machine Function."	"Appendix A General Notes"
5	GDS	J2UZ-7243-01ENZ0(A) J2UZ-7243-01ENZ2(A)	Now explicitly describes in [<i>Resolution</i>] of the daemon warning message 44028 that restoration work is not necessary after the daemon <i>daemon</i> ends abnormally.	"E.3.3 Warning Messages (44000-44099)"
6	GDS	J2UZ-7243-01ENZ0(A) J2UZ-7243-01ENZ2(A)	Removed a description of the noauto option setting related to a specification of file system checking on OS startup in the sixth field of /etc/fstab.	"A.2.29 File System Auto Mount"
7	GDS	J2UZ-7243-01ENZ0(A) J2UZ-7243-01ENZ2(A)	Added the setting for exporting file systems created on local class volumes to NFS clients.	"A.2.29 File System Auto Mount"
8	GDS	J2UZ-7243-01ENZ0(A) J2UZ-7243-01ENZ2(A)	Changed the boot priority from "73 or higher" to "54 or higher."	"A.2.31 Volume's Block Special File Access Permission"
9	GDS	J2UZ-7243-01ENZ0(A) J2UZ-7243-01ENZ2(A)	Changed the variable name <i>pdevtree</i> used in message descriptions to <i>pdev</i> and removed <i>pdevinst</i> . Changed the following messages that use these variable names: 22000, 22001, 22002, 22003, 22004, 22005, 22006, 22007, 22008, 22009, 22010, 22011, 24004, 24005, 24007, 24009, and 24010.	"E.2 Driver Messages"

GDS: Global Disk Services

G.2.2 4.3A10 Manual Changes

No.	Product	Manual code	Description	Chapter/Section
1	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added "Software Release Guide PRIMECLUSTER(TM) GDS" and "Software Release Guide PRIMECLUSTER(TM) GDS Snapshot" to "Related documentation."	"Preface"
2	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Changed the description of the PRIMEQUEST server in "Symbol."	"Preface"
3	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Changed the description of Java and related matters in "Trademarks."	"Preface"

No.	Product	Manual code	Description	Chapter/Section
4	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added the KVM host OS in the environment in which system disk mirroring is supported.	"1.2.1.1 System Disk Mirroring [EFI]"
5	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added reference information on operating systems that support root classes in the PRIMEQUEST 1000 Series.	"1.2.1.1 System Disk Mirroring [EFI]"
6	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added description of shared disk mirroring which can be performed in the virtual machine environment.	"1.2.1.3 Shared Disk Mirroring"
7	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added Note on conditions to be fulfilled for the hot spare function to work as specified.	"1.2.2 Hot Spare"
8	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added Note on "System Disk Management."	"1.3.2 Centralized Disk Management"
9	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Corrected the number of available partitions.	"1.3.6.1 Logical Partitioning"
10	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Corrected the number of logical devices divided by GDS.	"1.3.6.1 Logical Partitioning"
11	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Changed the ETERNUS support models described in Note on "Instant Snapshot by OPC."	"1.5.3 Instant Snapshot by OPC"
12	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Changed the file name described in "Configuring an alternative boot environment."	"1.5.6 Creating Alternative Boot Environment [EFI]"
13	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Moved articles on "Disk Class," "SDX Disk," "Disk Group," "Logical Volume," and "Logical Slice" under "2.1 SDX Object."	"Chapter 2 Structure"
14	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added the figure on the interrelation of SDX objects.	"2.1 SDX Object"
15	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added description of operating systems that support root classes in the PRIMEQUEST 1000 Series to "Root [EFI]" under "Type" of "Attributes."	"2.1.1 Disk Class"
16	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added description of specifying the disk type to "Root [EFI]" under "Type" of "Attributes."	"2.1.1 Disk Class"
17	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added description of purposes of using root classes to "Root [EFI]" under "Type" of "Attributes."	"2.1.1 Disk Class"
18	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Corrected the number of available partitions.	"2.1.4 Logical Volume"
19	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Corrected the number of single disks and groups that can be partitioned as volumes.	"2.1.4 Logical Volume"

No.	Product	Manual code	Description	Chapter/Section
20	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Corrected the mistake in Note on "Abnormality Detected with GDS Management."	"4.3 Icon Types and Object Status"
21	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added description of time required for processing after clicking <Create> in "5. Confirming system disk configuration."	"5.2.1 System Disk Settings [PRIMEQUEST] "
22	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added environmental conditions for mirroring system disks.	"5.2.1 System Disk Settings [PRIMEQUEST] "
23	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added description of Note on "When performing system disk mirroring on PRIMEQUEST."	"5.2.1 System Disk Settings [PRIMEQUEST] "
24	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added the screenshot output for RHEL6 to "Setting System Disk Mirroring Complete" in "6. Completing system disk configuration."	"5.2.1 System Disk Settings [PRIMEQUEST] "
25	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added description of operation after completing the system disk configuration to "6. Completing system disk configuration."	"5.2.1 System Disk Settings [PRIMEQUEST] "
26	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added the reference section for checking system disk settings.	"5.2.1 System Disk Settings [PRIMEQUEST] "
27	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added the information on how to register physical disks which have different sizes with the class.	"5.2.2.1 Class Configuration"
28	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added reference information on the number of physical disks that can be selected at a time to "2. Selecting physical disk."	"5.2.2.1 Class Configuration"
29	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added Note on "Size displayed in the <Unused> field" to "2. Selecting unused volume."	"5.2.2.4 Volume Configuration"
30	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added description of volume sizes that can be created and on display of [Disk Size] to Note in "b) Setting [Volume Size]" of "3. Setting volume attributes."	"5.2.2.4 Volume Configuration"
31	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added "1. Check the device name change" to "Restore Physical Disk."	"5.3.4 Disk Swap"
32	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added "5. Updating physical disk information" to "3. Changing configuration" in "Changing Class Configuration."	"5.4.1 Class Configuration"

No.	Product	Manual code	Description	Chapter/Section
33	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added Note on "Size displayed in the <Unused> field" to step 2 in "3. Removing volume."	"5.5.2 Removing Volume"
34	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added Note on "Defining /etc/fstab during Unmirroring System Disks."	"5.5.5 Unmirroring the System Disk [PRIMEQUEST]"
35	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added Note on the disk displayed in the [Mirror Disk] field to "1. Confirming system disk configuration."	"5.5.5 Unmirroring the System Disk [PRIMEQUEST]"
36	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Changed the command examples in "1) Check the root class name and the system volume names." and "5) Check information on the physical disks composing the system volumes."	"6.1.1 Checking Physical Disk Information and Slice Numbers"
37	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Changed the command examples in "a5) Check the backup target physical slice name."	"6.1.2 Backing Up"
38	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added description of variation of the command to be used due to the type of a file system to be backed up in "a6) Back up data of a file system to a tape medium." and "b3) Check the volume of the file system to be backed up."	"6.1.2 Backing Up"
39	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Changed the command examples in "6) Check the restore destination physical slice name."	"6.1.3 Restoring (When System Can Be Booted)"
40	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Changed the command examples in "3) Check the restore destination physical slice names."	"6.1.4 Restoring (When System Cannot Be Booted)"
41	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added "12) Restarting the system."	"6.5.3 Restoring from Backup Disks"
42	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Corrected the mistake in "15) Restoring the object configuration of the class."	"6.5.3 Restoring from Backup Disks"
43	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Clarified the necessity to reboot the node after changing the class type and expanding the class scope.	"6.7.2 Restoring"
44	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Changed the number of volumes that can be created within a single cluster.	"A.1.5 Number of Volumes"
45	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Changed the title of A.2.27 to "Editing Configuration Tables Created by sdxconfig Backup Command."	"A.2 Important Points"
46	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added "A.2.34 System Reconfiguration."	"A.2 Important Points"

No.	Product	Manual code	Description	Chapter/Section
47	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added "A.2.35 Recovery of FULL Shared Disk."	"A.2 Important Points"
48	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added Note on servers that can be used for system disk management.	"A.2.1 Managing System Disks"
49	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added description of displaying disk sizes by the sdxinfo command to Note on "Cylinder Size."	"A.2.5 Disk Size"
50	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added conditions essential for working the hot spare function as specified to Note.	"A.2.7 Hot Spare"
51	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Changed the conditions for not being able to register a disk in the root class.	"A.2.9 Keep Disk [EFI]"
52	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added description of cases of implementing normal resynchronization even if the volume's JRM mode is enabled to Note on "JRM for Volumes."	"A.2.12 Just Resynchronization Mechanism (JRM)"
53	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added description of "Swapping Physical Disks when the Proxy Volume Is a Mirror Volume."	"A.2.14 Swapping Physical Disks"
54	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added description of changing over a root class to a cluster system.	"A.2.24 Migration from Single Nodes to Cluster System"
55	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added description of changing the character device special file of the raw device.	"A.2.30 raw Device Binding"
56	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added Note on using a volume or a temporarily detached slice as a raw device in RHEL5 or later.	"A.2.30 raw Device Binding"
57	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added description of KVM.	"A.3.5 Guidelines for GDS Operation on the Virtual Machine Function"
58	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added "D.14 sdxdevinfo - Display device information."	"Appendix D Command Reference"
59	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added the section "D.22 Checking System Disk Settings Using Commands."	"Appendix D Command Reference"
60	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added a note on checks to be made in the cluster domain before executing the command in the primary option "-R."	"D.1 sdxclass - Class Operations"
61	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added notes on checks to be made in the cluster domain before executing the command in the primary options "-M" and "-R."	"D.2 sdxdisk - Disk Operations"
62	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added reference information on the number of devices that can be	"D.2 sdxdisk - Disk Operations"

No.	Product	Manual code	Description	Chapter/Section
			registered at a time to the primary option "-M."	
63	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Changed the description of the sub option "root [EFI]."	"D.2 sdxdisk - Disk Operations"
64	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Corrected the number of volumes that can be created if <i>class</i> is a local class or a shared class in the primary option "-M."	"D.4 sdxvolume - Volume Operations"
65	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added reference information on the number of volumes that can be specified at a time to the primary options "-F" and "-N."	"D.4 sdxvolume - Volume Operations"
66	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added reference information on the number of slices that can be specified at a time to the primary options "-F," "-N," and "-T."	"D.5 sdxslice - Slice Operations"
67	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added description of variation of time required for display due to the configuration of the objects to "DESCRIPTION."	"D.6 sdxinfo - Display Object Configuration and Status Information"
68	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added Note on I/O error events in the root class to "Disk Information" of "DISPLAYED INFORMATION."	"D.6 sdxinfo - Display Object Configuration and Status Information"
69	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added Note on viewing the COPY status of slices in proxy volumes to the sub option "-o object."	"D.6 sdxinfo - Display Object Configuration and Status Information"
70	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Deleted description of "name= <i>classname</i> " in sub options.	"D.7 sdxattr - Change Objects Attributes"
71	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added the sub option "-x NoRdchk" specified to restore the object status without reading the disk area.	"D.9 sdxfix - Restore Failed Object"
72	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added reference information on the number of volumes that can be specified at a time to the primary options.	"D.10 sdxcopy - Synchronization Copying Operation"
73	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Changed file names used for the primary options "-M" and "-R."	"D.11 sdxroot - Root File System Mirroring Definition and Cancellation [EFI]"
74	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added a note on checks to be made in the cluster domain before executing the command to the primary option "Restore."	"D.13 sdxconfig - Object Configuration Operations"
75	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Changed file names used for the primary option "Root."	"D.15 sdxproxy - Proxy Object Operations"

No.	Product	Manual code	Description	Chapter/Section
76	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added "Resolution" to the daemon warning message 44020.	"E.3.3 Warning Messages (44000-44099)"
77	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added description on the case in which there is a device name change to Explanation and Resolution in the daemon internal error message 48000.	"E.3.5 Internal Error Messages (48000)"
78	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Corrected the number of groups that can be created within <i>class</i> if <i>class</i> is a local class or a shared class in Explanation of the command error message 60071.	"E.4.1 Error Messages (60000-60399)"
79	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added description on the case in which there is a device name change to Resolution in the command error messages 60176 and 60181.	"E.4.1 Error Messages (60000-60399)"
80	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added 60367 to command error messages.	"E.4.1 Error Messages (60000-60399)"
81	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Changed file names of the command information messages 64047 and 64079.	"E.4.3 Information Messages (64000-64099)"
82	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Changed a file name of the management view message 5211.	"E.5.3 Information Messages (5200-5299)"
83	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added description of I/O error events caused by SCSI timeouts to "Cause c" of "(1) Mirror slice configuring the mirror volume is in INVALID status."	"F.1.1 Slice Status Abnormality"
84	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Changed procedures and command examples in "B. For the swap area [EFI]" 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"
85	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added Note on I/O error events in root classes to "(3) Disk is in I/O error status."	"F.1.2 Disk Status Abnormality"
86	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added Notes to the steps c8) and d3) in "(5) Proxy volume is in INVALID status."	"F.1.3 Volume Status Abnormality"
87	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added Note on I/O error events in root classes to "(9) An I/O error occurs on a stripe volume or a volume in a concatenation group."	"F.1.3 Volume Status Abnormality"
88	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Clarified the necessity to cancel the I/O error status of the E field in steps (8) and (9).	"F.1.3 Volume Status Abnormality"

No.	Product	Manual code	Description	Chapter/Section
89	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added description of "(2) Class cannot be started when booting the system" as the resolution for class status errors.	"F.1.4 Class Status Abnormality"
90	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added description of "(5) Insufficient memory error is displayed at Java console in the environment with a number of SDX objects."	"F.1.6 GDS Management View Abnormality"
91	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Corrected the mistake in "(1) SRDF is not used in master-to-proxy synchronization copying."	"F.1.8 EMC Symmetrix Abnormality"
92	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Changed description of Resolution in "(1) The error message "ERROR: <i>class</i> : cannot operate in cluster environment,..." is output, and the operation cannot be conducted on the class <i>class</i> ."	"F.1.9 Cluster System Related Error"
93	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added description of "(5) The error message "ERROR: <i>class</i> : cannot operate shared objects,..." is output, and the shared class <i>class</i> cannot be created; or the error message "ERROR: cluster communication failure" is output, and automatic registration of resources is failed."	"F.1.9 Cluster System Related Error"
94	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added description of necessary materials for contacting field engineers.	"F.2 Collecting Investigation Material"
95	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added Note on timing of collecting the crash dump.	"F.2 Collecting Investigation Material"
96	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added notes on checks to be made in the cluster domain before executing the command.	"H.5.1clautoconfig(8)-Execute resource registration"
97	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Changed the reference manual names for restrictions, resolutions, and restriction removal schedules.	"I.1 New Features"
98	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Added description of operating systems that support root classes in the PRIMEQUEST 1000 Series to "root class [EFI]."	"Glossary"
99	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Deleted description of "root disk [Itanium]."	"Glossary"
100	GDS	J2UZ-7243-02ENZ0(00) J2UZ-7243-02ZEN2(00)	Deleted description of "JBOD."	"Glossary"
101	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Modified the description of the environment where the system	- "1.2.1.1 System Disk Mirroring [EFI]"

No.	Product	Manual code	Description	Chapter/Section
			disk management and mirroring are possible.	- "1.3.2 Centralized Disk Management" - "A.2.1 Managing System Disks"
102	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Deleted Note about conditions where specifications by the Hot Spare mode work correctly.	"1.2.2 Hot Spare"
103	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Changed the description of the number of available partitions in physical disks.	- "1.3.6.1 Logical Partitioning" - "2.1.4 Logical Volume"
104	GDS Snapshot	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Deleted model names of disk units with the OPC function.	- "1.5.3 Instant Snapshot by OPC" - "A.2.17 Instant Snapshot by OPC"
105	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Changed the description about the attributes of class type "Root."	- "2.1.1 Disk Class" - "D.2 sdxdisk - Disk Operations"
106	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added references to Important Points of "Reference."	- "2.1.1 Disk Class" - "2.1.2 SDX Disk" - "2.1.4 Logical Volume"
107	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Changed the description of Status "DISABLE."	"2.1.2 SDX Disk"
108	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Changed the description about conditions of free space for the system disk and keep disk.	- "5.2.1 System Disk Settings [PRIMEQUEST]" - "A.2.9 Keep Disk [EFI]"
109	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added or changed Note "Information collection and environment configuration before and after setting the system disk."	- "5.2.1 System Disk Settings [PRIMEQUEST]" - "D.11 sdxroot - Root File System Mirroring Definition and Cancellation [EFI]"
110	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Deleted Note "Checking Physical Disk Information and Slice Numbers."	- "5.2.1 System Disk Settings [PRIMEQUEST]" - "D.11 sdxroot - Root File System Mirroring Definition and Cancellation [EFI]"
111	GDS Snapshot	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added the description of the maximum number of volumes in a group.	"5.2.4.1 Join"
112	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Moved steps 1) to 4) of "1. Check the device name change" in "Restore Physical Disk" to "Swapping Internal Disks"	- "5.3.4 Disk Swap" - "A.2.14 Swapping Physical Disks"

No.	Product	Manual code	Description	Chapter/Section
			Registered with Root Classes or Local Classes [RHEL6]" in "A. 2.14 Swapping Physical Disks." Then, the description was changed to direct you to see A. 2.14 for steps.	
113	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Deleted "5.Updating physical disk information" in "Changing Class Configuration."	"5.4.1 Class Configuration"
114	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added Note "In Cluster Systems."	"5.5.4 Removing Class"
115	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added the example when the type attribute is changed to Information.	- "6.1.2 Backing Up" - "6.1.3 Restoring (When System Can Be Booted)"
116	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added Information when the type attribute is not changed.	- "6.1.2 Backing Up" - "6.1.3 Restoring (When System Can Be Booted)"
117	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Changed the example for checking the backup target or restore destination of a physical disk name. (only for RHEL6)	- "6.1.2 Backing Up" - "6.1.3 Restoring (When System Can Be Booted)" - "6.1.4 Restoring (When System Cannot Be Booted)"
118	GDS Snapshot	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Deleted the description of ROPC.	- "6.4 Online Backup and Restoration Using Proxy Volume" - "6.4.3 Instant Restore" - "A.1.8 Proxy Configuration Preconditions" - "A.2.16 Using Advanced Copy Function in Proxy Configuration" - "D.6 sdxinfo - Display Object Configuration and Status Information" - "F.1.7 Proxy Object Abnormality" - "Glossary"
119	GDS	1.1	Changed the example of the sdxinfo -G -e long command.	- "6.6.1.6 Backing Up" - "6.6.1.7 Restoring" - "F.1.3 Volume Status Abnormality"
120	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Corrected misdescriptions of step 4-2).	"6.7.2 Restoring"
121	GDS, GDS Snapshot	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Described the server and storage migration.	"Appendix"

No.	Product	Manual code	Description	Chapter/Section
122	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Described the disaster recovery.	"Appendix"
123	GDS Snapshot	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added the description regarding the number of volumes when the proxy operation is performed per group.	"2.3.5 Number of Volumes"
124	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added the description of a proxy operation of the root class to General Preconditions.	"2.3.11 Proxy Configuration Preconditions"
125	GDS Snapshot	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added the description for the root class.	"2.3.12 Number of Proxy Volumes"
126	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added the following sections: - A.2.35 Operating When There is a Disk in DISABLE Status or There is a Class not Displayed with the sdxinfo Command - A.2.36 Adding and Removing Disks [4.3A00] - A.2.37 Use of GDS in a Xen Environment - A.2.38 Use of GDS in a KVM Environment [4.3A10 or later] - A.2.39 Use of GDS in a VMware Environment	"A.2 Important Points"
127	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Moved the description of "A.2.35 Recovery of FULL Shared Disk" to "(6) The disk space of a file system on a shared disk is "Full" (100%)" in "F.1.9 Cluster System Related Error."	- "A.2 Important Points" - "D.1.9 Cluster System Related Error"
128	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added virtual disks on a KVM guest to the list of physical special files.	"A.1 Restraining Access to Physical Special File"
129	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added descriptions of the number of disks within the class.	"3.8 Hot Spare"
130	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Changed the descriptions of "Hot Spare for Hierarchized Mirror Groups."	"3.8 Hot Spare"
131	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added descriptions of the processing time of Synchronization Copying Invoked by Hot Spare to Note.	"3.8 Hot Spare"
132	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added the setting procedure for system disk mirroring.	"A.2.8 System Disk Mirroring [PRIMEQUEST]"
133	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Changed the description of "Swapping Physical Disk when the Proxy Volume Is in a Mirroring Configuration."	"A.2.14 Swapping Physical Disks"

No.	Product	Manual code	Description	Chapter/Section
134	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added the description of a VMware environment.	"A.3.5 Guidelines for GDS Operation in the Virtual Environment"
135	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Moved the description of the maximum number of objects which can be operated at once from PRIMARY OPTIONS "-M" to SUB OPTIONS "-d <i>device=disk[:type]</i> [, <i>device=disk[:type]</i> ,...]" (when using -M)."	"B.1.2 sdxdisk - Disk operations"
136	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added virtual disks on a KVM guest to the format of a physical disk name.	- "B.1.2 sdxdisk - Disk operations" - "B.1.6 sdxinfo - Display object configuration and status information" - "B.1.13 sdxconfig - Object configuration operations" - "B.2.2 sdxshadowdisk - Shadow disk operations"
137	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Moved the description of the maximum number of objects which can be operated at once from PRIMARY OPTIONS "-F" and "-N" to SUB OPTIONS "-v <i>volume</i> ,... (when using -F,-N)."	"B.1.4 sdxvolume - Volume operations"
138	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Moved the description of the maximum number of objects which can be operated at once from PRIMARY OPTIONS "-F", "-N", and "-T" to SUB OPTIONS "-s <i>slice</i> ,... (when using -F,-N,-T)."	"B.1.5 sdxslice - Slice operations"
139	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Changed the description of SUB OPTIONS "-x NoRdchk."	"B.1.9 sdxfix - Restore a failed object"
140	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Moved the description of the maximum number of objects which can be operated at once from Information of PRIMARY OPTIONS to SUB OPTIONS "-v <i>volume</i> ,..."	"B.1.10 sdxcopy - Synchronization copying operation"
141	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added a description of a cluster system to "Restore" in PRIMARY OPTIONS.	"B.1.13 sdxconfig - Object configuration operations"
142	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added Note about restoration of the device numbers, owners, and access permissions to "Restore" in PRIMARY OPTIONS.	"B.1.13 sdxconfig - Object configuration operations"

No.	Product	Manual code	Description	Chapter/Section
143	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added a description to SUB OPTIONS "-e chkps (when using Restore)."	"B.1.13 sdxconfig - Object configuration operations"
144	GDS Snapshot	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added a description when Dell EMC SRDF is used to SUB OPTIONS "-e force (when using Break)."	"B.2.1 sdxproxy - Proxy object operations"
145	GDS Snapshot	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added a description of the maximum number of objects which can be operated at once to SUB OPTIONS "-p <i>proxy</i> ,..." (when using Part, Rejoin, RejoinRestore, Restore, Root or Update)."	"B.2.1 sdxproxy - Proxy object operations"
146	GDS Snapshot	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added a description of the maximum number of objects which can be operated at once to SUB OPTIONS "-d <i>device=disk[:type]</i> [, <i>device=disk[:type]</i> ,...]" (when using -M)."	"B.2.2 sdxshadowdisk - Shadow disk operations"
147	GDS Snapshot	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Changed the description of the number of volumes which can be created in "-M" in PRIMARY OPTIONS.	"B.2.4 sdxshadowvolume - Shadow volume operations"
148	GDS Snapshot	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added a description of the maximum number of objects which can be operated at once to SUB OPTIONS "-v <i>volume</i> ,..." (when using -F,-N)."	"B.2.4 sdxshadowvolume - Shadow volume operations"
149	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Specified that [EFI] is the target.	"6.4.4 Checking System Disk Settings"
150	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Changed Explanation and Resolution of the message. (44015)	"C.3.3 Warning Messages (44000 - 44099)"
151	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Changed Explanation of the error message. (60071)	"C.4.1 Error Messages (60000-60399)"
152	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Changed Explanation and Resolution of the message. (60089)	"C.4.1 Error Messages (60000-60399)"
153	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Changed Resolution of the error message. (60181)	"C.4.1 Error Messages (60000-60399)"
154	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Changed the example of the sdxfix -D command stated in (8) and (9) in Resolution b.	"D.1.3 Volume Status Abnormality"
155	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added "(4) System cannot be booted. (Failure of all boot disks)"	"D.1.5 System Disk Abnormality [EFI]"

No.	Product	Manual code	Description	Chapter/Section
156	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added "Q: When an I/O error status is created in the test, how do I shorten the restoration time?"	"E.3 Operation"
157	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added Note when the device name of the shared disk unit is different depending on a node in step 3).	"6.2.4 Shared Disk Unit Resource Registration Procedure"
158	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added "K.1.4 New Features in 4.3A10."	"F.1 New Features"
159	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Added the following features: - Two generation proxies for system volumes - Shortening the restoring time of the I/O error status of a disk	"F.1.1 New Features in 4.3A00"
160	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Deleted "K.3.5.1 Support of Red Hat Enterprise Linux 6 (for Intel64)."	"F.3.2 Specification Changes in 4.3A10"
161	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Changed the description of the term "root class [EFI]."	"Glossary"
162	GDS	J2UZ-7243-02ENZ0(01) J2UZ-7243-02ZEN2(01)	Deleted the following terms: - ETERNUS2000 series - ETERNUS3000 series - ETERNUS4000 series - ETERNUS6000 series - ETERNUS8000 series	"Glossary"

GDS: Global Disk Services

GDS Snapshot: Global Disk Services Snapshot

G.2.3 4.3A20 Manual Changes

No.	Product	Manual code	Description	Chapter/Section
1	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Changed the number of partitions which can be used for physical disks.	- "1.3.6.1 Logical Partitioning" - "2.1.4 Logical Volume"
2	GDS Snapshot	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Described that cooperation with REC, Dell EMC TimeFinder, and Dell EMC SRDF are not supported as Note.	- "1.5.2 Snapshot Function with No Load to Server/SAN" - "A.2.17 Using Advanced Copy Function in Proxy Configuration" - "A.2.20 Using EMC TimeFinder or EMC SRDF in Proxy Configuration"
3	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Added the description of the maximum number of the stripe width.	- "2.1.3 Disk Group" - "2.2.2.3 Shadow Group" - "D.2 sdxdisk - Disk operations" - "D.3 sdxgroup - Group Operations"

No.	Product	Manual code	Description	Chapter/Section
				- "D.16 sdxshadowdisk - Shadow Disk Operations" - "D.17 sdxshadowgroup - Shadow Group Operations"
4	GDS Snapshot	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Described that shadow objects are not supported as Note.	- "2.2.2 Shadow Object" - "A.2.27 Shadow Volume"
5	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Added Note "Rebooting the System after System Disk Settings are Completed."	"5.2.1 System Disk Settings [PRIMEQUEST]"
6	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Added Note "Node Name."	"5.2.2.2 Cluster System Class Configuration"
7	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Described that starting a volume is necessary before creating or deleting a file system.	- "5.2.3 File System Configuration" - "5.5.1 Removing File System"
8	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Described that the class name of class attributes cannot be changed.	"5.4.1 Class Configuration"
9	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Added the description of changes in the special file path name by changing the class name or volume name.	- "5.4.3 Volume Configuration" - "D.7 sdxattr - Change Objects Attributes"
10	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Described that stopping the operation used the volume and stopping the volume are necessary before removing a volume.	"5.5.2 Removing Volume"
11	GDS Snapshot	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Described about the number of disks which can be connected to a proxy group of a root class.	"A.1.3 Number of Disks"
12	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Added the following descriptions: - "A.2.5 Disk Label" - "A.2.41 Excluded Device List" - "A.2.42 DM-MP (Device Mapper Multipath) [4.3A10 or later]" - "A.2.43 Root Class Operation [EFI]"	"A.2 Important Points"
13	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Added the note about backing up and restoring object configurations and change the title to "Backing Up and Restoring Object Configuration (sdxconfig)."	"A.2 Important Points"
14	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Added the mpath device of DM-MP as a physical special file.	"A.2.2 Restraining Access to Physical Special File"
15	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Added "1 TB or larger disks [4.3A10 or later]."	"A.2.6 Disk Size"
16	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Added Note "JRM for Proxies for a Root Class [EFI]."	"A.2.13 Just Resynchronization Mechanism (JRM)"

No.	Product	Manual code	Description	Chapter/Section
17	GDS Snapshot	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Added Note "Multipath Software."	"A.2.17 Using Advanced Copy Function in Proxy Configuration"
18	GDS Snapshot	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Changed the description of the multipath software.	"A.2.19 To Use EMC Symmetrix"
19	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Added the description of setting up the configuration parameter when using Dell EMC PowerPath.	"A.2.19 To Use EMC Symmetrix"
20	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Changed Information "Disks Described in the Excluded Device List."	"A.2.19 To Use EMC Symmetrix"
21	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Added the note about a scope in a cluster system with three or more nodes.	"A.2.24 Operation in Cluster System"
22	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Changed the setting procedure when using a virtual disk as a shared disk on a KVM guest.	"A.2.39 Use of GDS in a KVM Environment [4.3A10 or later]"
23	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Added the mpath device of DM-MP to the format of a physical disk name.	- "D.2 sdxdisk - Disk Operations" - "D.6 sdxinfo - Display Object Configuration and Status Information" - "D.13 sdxconfig - Object Configuration Operations " - "E.3 Daemon Messages" - "E.4 Command Messages"
24	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Added SUB OPTIONS "-e label."	"D.6 sdxinfo - Display Object Configuration and Status Information"
25	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Added LABEL to Class information.	"D.6 sdxinfo - Display Object Configuration and Status Information"
26	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Added a message.(22013)	"E.2.1 Warning Messages (22000-22099)"
27	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Added the virtual disk on a KVM guest to the format of a physical disk name.	- "E.3 Daemon Messages" - "E.4 Command Messages"
28	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Added the mpath device of DM-MP and virtual disk on a KVM guest to the format of a physical slice name.	- "E.3 Daemon Messages" - "E.4 Command Messages"
29	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Added Messages. - 46115 - 60375 - 60379 - 60383	- "E.3.4 Information Messages (46000-46199)" - "E.4.1 Error Messages (60000-60399)"

No.	Product	Manual code	Description	Chapter/Section
30	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Described the method of collecting investigation material of the Management View.	"F.2 Collecting Investigation Material"
31	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Described that DM-MP can be used by applying patches.	"K.1.3 New Features in 4.3A00"
32	GDS	J2UZ-7243-03ENZ0(00) J2UZ-7243-03ENZ2(00)	Added the term "mpath device."	"Glossary"
33	GDS	J2UZ-7243-03ENZ0(01) J2UZ-7243-03ENZ2(01)	Changed the required patch number when using 1 TB or larger disks in GDS 4.3A10.	- "A.2.6 Disk Size" - "E.4.1 Error Messages (60000-60399)"
34	GDS	J2UZ-7243-03ENZ0(01) J2UZ-7243-03ENZ2(01)	Changed the required patch numbers for RHEL6.	"A.2.42 DM-MP (Device Mapper Multipath)"
35	GDS	J2UZ-7243-03ENZ0(01) J2UZ-7243-03ENZ2(01)	Changed the description of a device name in the shared disk definition file.	- "A.2.42 DM-MP (Device Mapper Multipath)" - "H.4 Shared Disk Unit Resource Registration Procedure"
36	GDS	J2UZ-7243-03ENZ0(01) J2UZ-7243-03ENZ2(01)	Added the following messages: - 44039 - 44040	"E.3.3 Warning Messages (44000-44099)"
37	GDS	J2UZ-7243-03ENZ0(01) J2UZ-7243-03ENZ2(01)	Modified the procedure for PRIMEQUEST 1000 Series described in "Resolution" in "(4) System cannot be booted. (Failure of all boot disks).	"F.1.5 System Disk Abnormality [PRIMEQUEST]"

GDS: Global Disk Services

GDS Snapshot: Global Disk Services Snapshot

G.2.4 4.3A30 Manual Changes

No.	Product	Manual code	Description	Chapter/Section
1	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Changed the description for the environment where the system disk mirroring is supported.	"1.2.1.1 System Disk Mirroring [EFI] "
2	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Changed the Note "System Disk Management.	"1.3.2 Centralized Disk Management"
3	GDS, GDS Snapshot	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Specified that the target version is 4.3A20 or earlier. - Snapshot creation by using ETERNUS Disk storage system REC function is not supported. - File System Configuration screen - In DM-MP environment, snapshot creation by using ETERNUS Disk storage system Advanced Copy function is not supported.	- "1.5.2 Snapshot Function with No Load to Server/SAN" - "4.2.2 Settings" - "5.2.3 File System Configuration [4.3A20 or earlier] " - "5.5.1 Removing a File System [4.3A20 or earlier] " - "A.2.17 Using Advanced Copy Function in Proxy Configuration"

No.	Product	Manual code	Description	Chapter/Section
				- "A.2.40 DM-MP (Device Mapper Multipath)"
4	GDS Snapshot	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Added the description of QuickOPC.	"1.5.3 Instant Snapshot by OPC" - "6.4 Online Backup and Instant Restore through Proxy Volume" - "6.4.2 Online Backup (Snapshot by OPC)" - "6.4.3 Instant Restore" - "A.1.8 Proxy Configuration Preconditions" - "A.2.17 Using Advanced Copy Function in Proxy Configuration" - "A.2.18 Instant Snapshot by OPC" - "D.15 sdxproxy - Proxy Object Operations" - "F.1.7 Proxy Object Abnormality"
5	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Changed the environment conditions for using the root class.	- "2.1.1 Disk Class" - "Glossary"
6	GDS Snapshot	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Added the Note "When using QuickOPC function (4.3A30 or later).	- "5.1.3.3 Backup (by OPC)" - "5.3.2.3 Backup (by OPC)"
7	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Changed the description of the Note "Node Name."	"5.2.2.2 Cluster System Class Configuration"
8	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Added the necessary procedure when file systems are created in a volume.	"5.5.2 Removing Volume"
9	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Added the following sections: - "A.2.38 Hot Adding a Disk [RHEL6]" - "A.2.42 Changing system run level" - "A.2.44 kdump Settings [RHEL6]"	"A.2 Important Points"
10	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Changed the environment conditions for managing system disks.	"A.2.1 Managing System Disks"
11	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Described the notes for LVM configuration.	"A.2.1 Managing System Disks"
12	GDS, GDS Snapshot	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Described the target version for "Functions unavailable in local and shared classes of GPT label.	"A.2.5 Disk Label"
13	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Changed Note "Large Volumes [4.3A20 or earlier]"	"A.2.7 Volume Size"

No.	Product	Manual code	Description	Chapter/Section
14	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Added the procedure for migrating a class configuration to other systems.	"A.2.28 Backing Up and Restoring Object Configuration (sdxconfig)"
15	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Changed the title of "A.2.35" from "System Reconfiguration" to "OS Re-installation." Also, added the description for recovery methods when an operating system is reinstalled without deleting GDS settings.	"A.2.35 OS Re-installation"
16	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Added the Note "Device Name Change."	"A.2.39 Excluded Device List"
17	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Modified the patch numbers for GDS 4.3A00.	"A.2.40 DM-MP (Device Mapper Multipath)"
18	GDS, GDS Snapshot	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Added the Note "When using GDS Snapshot" in "CURBLKS" of "Slice information."	"D.6 sdxinfo - Display Object Configuration and Status Information"
19	GDS, GDS Snapshot	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Added the description for "QOPC [4.3A30 or later]" in "CPTYPE" and "FUNC" of "Slice information."	"D.6 sdxinfo - Display Object Configuration and Status Information"
20	GDS, GDS Snapshot	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Added "execute" and "track" in "CPSTAT" of "Slice information."	"D.6 sdxinfo - Display Object Configuration and Status Information"
21	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Modified the sub option "-e chkps (when using Restore)."	"D.13 sdxconfig - Object Configuration Operations"
22	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Added the sub option "-e mklablel (when using Restore)."	"D.13 sdxconfig - Object Configuration Operations"
23	GDS Snapshot	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Added "-e OPC" and "-e QOPC" in "SUB OPTIONS."	"D.15 sdxproxy - Proxy Object Operations"
24	GDS Snapshot	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Added the description of QuickOPC to the error conditions in "Restore" and "Update" of "PRIMARY OPTIONS."	"D.15 sdxproxy - Proxy Object Operations"
25	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Added the Note "RHEL6 /var/log/messages [4.3A10 or later]."	"Appendix E GDS Messages"
26	GDS, GDS Snapshot	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Added the following messages: 22014, 22015, 42049, 42050, 42051, 46116, 46117, 46118, 46119, 60373, 60374, 60376, 60377, 60378, 60380, 60384, 60385, 60386, 64096	- "E.2.1 Warning Messages (22000-22099)" - "E.3.2 Error Messages (42000-42099)" - "E.3.4 Information Messages (46000-46199)" - "E.4.1 Error Messages (60000-60399)" - "E.4.3 Information Messages (64000-64099)"
27	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Modified the explanation and resolution of the error message. (60379)	- "E.4.1 Error Messages (60000-60399)"

No.	Product	Manual code	Description	Chapter/Section
28	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Added the message output in 4.3A30 or later and its resolution. (5119)	"E.5.2 Warning Messages (5000,5100-5199)"
29	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Added the following section: - "F.1.13 ACM Abnormality [RHEL6] [4.3A30 or later]"	"Appendix F Troubleshooting"
30	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Modified the descriptions of step 5) and 6) in "Explanation" of "(4) System cannot be booted. (Failure of all boot disks)."	"F.1.5 System Disk Abnormality [EFI]"
31	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Modified the description in "Explanation" of "(3) Stack information is output at command execution."	"F.1.12 OS Messages"
32	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Added "Appendix K Usage in the Virtual Environment" and move the following sections: - "A.2.38 Use of GDS in a Xen Environment" - "A.2.39 Use of GDS in a KVM Environment [4.3A10 or later]" - "A.2.40 Use of GDS in a VMware Environment" - "A.3.5 Guidelines for GDS Operation in the Virtual Environment"	"Appendix K Usage in the Virtual Environment"
33	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Modified the following operational environments; mirroring of system disks and virtual disks on a guest, and managing shared disks in a cluster system.	"Appendix K Usage in the Virtual Environment"
34	GDS Snapshot	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Added the description about GDS Snapshot that can be used in virtual environments (Xen environment, KVM environment [4.3A10 or later], and VMware environment.)	"Appendix K Usage in the Virtual Environment"
35	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Added the description of a virtual disk (VMDK) on a VMware guest.	- "Appendix K Usage in the Virtual Environment" - "K.4 Use of GDS in a VMware Environment"
36	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Changed the description of "Adding a virtual disk to a KVM guest."	"K.4 Use of GDS in a VMware Environment"
37	GDS	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Added the description for using virtio-SCSI device.	"K.3 Use of GDS in a KVM Environment [4.3A10 or later]"
38	GDS, GDS Snapshot	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Added the following sections: - "L.1.6 New Features in 4.3A30" - "L.2.6 4.3A30 Manual Changes"	"Appendix L Release Information"

No.	Product	Manual code	Description	Chapter/Section
			- "L.3.7 Specification Changes in 4.3A30"	
39	GDS, GDS Snapshot	J2UZ-7243-04ENZ0(00) J2UZ-7243-04ENZ2(00)	Added the following terms: - "device name change" - "tracking" - "QuickOPC" - "QuickOPC function" - "QuickOPC session"	- "Glossary"
40	GDS	J2UZ-7243-04ENZ0(01) J2UZ-7243-04ENZ2(01)	Described the function addition of a patch to Note "Local and shared classes of GPT label [4.3A10/4.3A20]."	"2.3.8 Disk Label"
41	GDS	J2UZ-7243-04ENZ0(01) J2UZ-7243-04ENZ2(01)	Modified the description of VMware environment.	"3.21 Usage in the Virtual Environment"
42	GDS	J2UZ-7243-04ENZ0(01) J2UZ-7243-04ENZ2(01)	Modified the description of a necessary setting when the virtual disk (VMDK) is used.	"4.9 Use of GDS in a VMware Environment"
43	GDS	J2UZ-7243-04ENZ0(01) J2UZ-7243-04ENZ2(01)	Added descriptions of how to write when adding a shared disk unit in "2. Creating a shared disk definition file."	"6.3.4 Shared Disk Unit Resource Registration Procedure"
44	GDS	J2UZ-7243-04ENZ0(01) J2UZ-7243-04ENZ2(01)	Changed the descriptions of "Stopping the services" and "Resuming the services."	"7.6 Backup (by Offline)" "7.10.2.1 Without Using GDS Snapshot"
45	GDS	J2UZ-7243-04ENZ0(01) J2UZ-7243-04ENZ2(01)	Added the following messages: - 60387 - 60388 - 60389	"C.4.1 Error Messages (60000 - 60399)"
46	GDS	J2UZ-7243-04ENZ0(01) J2UZ-7243-04ENZ2(01)	Added "Q: Can NAS be mirrored?"	"E.1 Operation Design"
47	GDS	J2UZ-7243-04ENZ0(01) J2UZ-7243-04ENZ2(01)	Added "Mirroring and snapshot for 1 TB or larger volumes."	"F.1.4 New Features in 4.3A30"
48	GDS	J2UZ-7243-04ENZ0(02) J2UZ-7243-04ENZ2(02)	Modified the name of "PRIMEQUEST."	"Preface"
49	GDS	J2UZ-7243-04ENZ0(03) J2UZ-7243-04ENZ2(03)	Added "6.2 Using iSCSI Disk [RHEL6 or later]."	"6.2 Using iSCSI Disk [RHEL6 or later]"
50	GDS	J2UZ-7243-04ENZ0(03) J2UZ-7243-04ENZ2(03)	Added the error message 42055.	"C.3.2 Error Messages (42000-42099)"
51	GDS	J2UZ-7243-04ENZ0(03) J2UZ-7243-04ENZ2(03)	Added the error messages 60390 to 60394.	"C.4.1 Error Messages (60000-60399)"
52	GDS	J2UZ-7243-04ENZ0(03) J2UZ-7243-04ENZ2(03)	Added the warning message 62025.	"C.4.2 Warning Messages (62000-62099)"

GDS: Global Disk Services

GDS Snapshot: Global Disk Services Snapshot

G.2.5 4.3A40 Manual Changes

No.	Product	Manual code	Description	Chapter/Section
1	-	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Changed the chapter organization.	All chapters
2	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added the description of the mirroring among servers.	"1.2.1.4 Mirroring among Servers" "3.9 Mirroring among Servers" "4.8 Disk Setting for Performing Mirroring among Servers" "6.1.4 To Use Mirroring among Servers" "6.3.1 What Is Shared Disk Resource Registration?" "A.1 Restraining Access to Physical Special File"
3	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added the description of the system file that is changed in RHEL7 environment.	"1.5.6 Creating Alternative Boot Environment [EFI]" "B.1.11 sdxroot - Root File System Mirroring Definition and Cancellation [EFI]" "B.2.1 sdxproxy - Proxy Object Operations"
4	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added the description of the netmirror type.	"2.1.2 SDX Disk"
5	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added the description of the netmirror group.	"2.1.3 Disk Group" "2.3.3 Number of Disks" "2.3.7 Creating Group Hierarchy" "5.4 Icon Types and Object Status" "6.6.1.2 Other Volume Configuration Settings" "6.6.1.5 Group Configuration" "7.3.3 Swapping Disks of Netmirror Group"
6	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added the description of the netmirror volume.	"2.1.4 Logical Volume"
7	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added the note on the netmirror.	"2.1.5 Logical Slice"
8	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added the description of RHEL7 to "1 TB or larger disks [4.3A10 or later]."	"2.3.9 Disk Size"
9	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added "3.2.3 virtio Block Device on KVM."	"3.2 DM-MP (Device Mapper Multipath) "

No.	Product	Manual code	Description	Chapter/Section
10	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added the description of RHEL7.	"3.5 System Disk Mirroring Preconditions" "3.21 Virtual Environment" "4.6 kdump Settings" "6.5.3 Checking System Disk Settings" "6.5.4 Checking Physical Disk Information and Slice Numbers" "6.7.2 File System Auto Mount" "6.8 raw Device Binding" "6.9 Volume's Block Special File Access Permission" "7.3.1.2 Operation Procedure" "7.4.2 Backing Up [EFI]" "7.4.5 Restoring (When System Can Be Booted) [EFI]" "7.4.6 Restoring (When System Cannot Be Booted) [EFI]" "9.5 OS Re-installation" "A.10 Changing system run level"
11	GDS, GDS Snapshot	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added the description that Xen environment is not supported in 4.3A40 or later.	"3.21 Virtual Environment"
12	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added the note when using the virtio block device as the shared disk.	"4.6 Configuring KVM Environment"
13	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added "6.1.4 To Use Mirroring among Servers."	"6.1 Creating Excluded Device List"
14	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added the description when using the mirroring among servers, and the description for a guest in the virtual environment to the "Device name."	"6.3.4 Shared Disk Unit Resource Registration Procedure"
15	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added "6.6.2.3 Creating Netmirror Volume [4.3A40 or later]."	"6.6.2.3 Creating Netmirror Volume"
16	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added the description of cloning.	"6.11 Cloning"
17	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added the description of xfs.	"7.5.6 Backing Up"
18	GDS, GDS Snapshot	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added the note "Disk in DISABLE status."	"7.5.7 Restoring"
19	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added the descriptions of mirroring among servers.	"7.16 Mirroring Among Servers"

No.	Product	Manual code	Description	Chapter/Section
20	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added the descriptions of removing the errata kernel.	"7.18 Removing Errata Kernel"
21	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added the note "For the cluster system."	"9.4.1.2 Operation Flow"
22	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added the description of the -a type=netmirror option of the sdxdisk -C command.	"B.1.2 sdxdisk - Disk Operations"
23	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added netmirror to TYPE descriptions of the disk information, the group information, and the volume information.	"B.1.6 sdxinfo - Display Object Configuration and Status Information"
24	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added the error messages 42056 to 42059.	"C.3.2 Error Messages (42000-42099)"
25	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Changed "Explanation" of the error message 44039.	"C.3.3 Warning Messages (44000-44099)"
26	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Changed "Explanation" and "Resolution" of the internal error message 48000.	"C.3.5 Internal Error Messages (48000)"
27	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added the explanation of the mirroring among servers to "Resolution" of the error message 60142.	"C.4.1 Error Messages (60000-60499)"
28	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Changed "Explanation" of the error messages 60308 and 60363.	"C.4.1 Error Messages (60000-60499)"
29	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added the error messages 60399 to 60416.	"C.4.1 Error Messages (60000-60499)"
30	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Changed "Explanation" of the error messages 64047 and 64079.	"C.4.3 Information Messages (64000-64099)"
31	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Changed "Explanation" and "Resolution" of the internal error message 68000.	"C.4.5 Internal Error Messages (68000)"
32	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Changed "Explanation" of the error message 5211.	"C.5.3 Information Messages (5200-5299)"
33	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Modified the description of the "A. For /(root), /usr, or /var [EFI]" in "(2) The copy destination slice was made INVALID due to an I/O error generated on the copy source slice during synchronization copying."	"D.1.1 Slice Status Abnormality"
34	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added the procedure for RHEL7 to "(4) System cannot be booted. (Failure of all boot disks)."	"D.1.5 System Disk Abnormality [EFI]"
35	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added "(6) System cannot be booted. (Removing errata kernel) [RHEL7]."	"D.1.5 System Disk Abnormality [EFI]"
36	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added "D.1.6 Errors in the Mirroring Among Servers."	"D.1.6 Errors in the Mirroring Among Servers"
37	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Changed the description of how to restart the management server in "How to change management server"	"D.1.7 GDS Management View Abnormality"

No.	Product	Manual code	Description	Chapter/Section
			(Web-Based Admin View)" in "(5) Insufficient memory error is displayed at Java console in the environment with a number of SDX objects."	
38	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added the description of FJQSS.	"D.2 Collecting Investigation Material" "D.2.3 Collecting Investigation Material with FJQSS"
39	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added the description of OS versions to be the target of the specification change.	"G.3.1.2 Change of I/O Amount Supported for Simultaneous Processing"
40	GDS	J2UZ-7243-05ENZ0(00) J2UZ-7243-05ENZ2(00)	Added the following terms: - "mirroring among servers" - "netmirror group" - "netmirror slice" - "netmirror volume"	"Glossary"
41	GDS	J2UZ-7243-05ENZ0(01) J2UZ-7243-05ENZ2(01)	Changed the description of the number of servers.	"1.2.1.4 Mirroring among Servers"
42	GDS	J2UZ-7243-05ENZ0(01) J2UZ-7243-05ENZ2(01)	Improved the description of the maximum disk numbers for the mirroring among servers.	"3.9.3 Object Configuration" "6.3.1 What Is Shared Disk Resource Registration?" "6.3.4 Shared Disk Unit Resource Registration Procedure"
43	GDS	J2UZ-7243-05ENZ0(01) J2UZ-7243-05ENZ2(01)	Added the description of the maximum number of the disk in the class.	"3.9.3 Object Configuration"
44	GDS	J2UZ-7243-05ENZ0(01) J2UZ-7243-05ENZ2(01)	Added the description of the number of nodes for a cluster system when using the mirroring among servers.	"3.9.4 Configuration and Settings of the Cluster System"
45	GDS	J2UZ-7243-05ENZ0(01) J2UZ-7243-05ENZ2(01)	Added the description of volume expansion.	"3.9.7 Note"
46	GDS	J2UZ-7243-05ENZ0(01) J2UZ-7243-05ENZ2(01)	Added the note when logging out the iSCSI target.	"4.8.4 Establishing iSCSI Session"
47	GDS	J2UZ-7243-05ENZ0(01) J2UZ-7243-05ENZ2(01)	Changed the description of how to set the tuning parameter that is set when using the mirroring among servers.	"4.8.1 Setting Tuning Parameters"
48	GDS	J2UZ-7243-05ENZ0(01) J2UZ-7243-05ENZ2(01)	Added "7.15.7 Stopping RMS."	"7.15.7 Stopping RMS"
49	GDS	J2UZ-7243-05ENZ0(01) J2UZ-7243-05ENZ2(01)	Changed the description of "Restoration procedure."	"9.5 OS Re-installation"
50	GDS	J2UZ-7243-05ENZ0(01) J2UZ-7243-05ENZ2(01)	Added "A.13 Notes on Mirroring among Servers [4.3A40] [RHEL6]".	"A.13 Notes on Mirroring among Servers"
51	GDS	J2UZ-7243-05ENZ0(01) J2UZ-7243-05ENZ2(01)	Added the panic message 20001.	"C.2.1 Panic Message (20001)"

No.	Product	Manual code	Description	Chapter/Section
52	GDS	J2UZ-7243-05ENZ0(01) J2UZ-7243-05ENZ2(01)	Added the error messages 42057 to 42061.	"C.3.2 Error Messages (42000-42099)"
53	GDS	J2UZ-7243-05ENZ0(01) J2UZ-7243-05ENZ2(01)	Deleted the error message 60415. Added the error messages 60419 and 60420.	"C.4.1 Error Messages (60000-60499)"
54	GDS	J2UZ-7243-05ENZ0(01) J2UZ-7243-05ENZ2(01)	Added "(7) System disk mirroring failed to be set [RHEL7]."	"D.1.5 System Disk Abnormality [EFI]"
55	GDS	J2UZ-7243-05ENZ0(01) J2UZ-7243-05ENZ2(01)	Added "(8) System is not started after system disk is unmirrored [RHEL6 or later]."	"D.1.5 System Disk Abnormality [EFI]"
56	GDS	J2UZ-7243-05ENZ0(01) J2UZ-7243-05ENZ2(01)	Modified the description of "Resolution" in "(2) The cluster application becomes Faulted after the node is restarted."	"D.1.6 Errors in the Mirroring Among Servers"
57	GDS	J2UZ-7243-05ENZ0(01) J2UZ-7243-05ENZ2(01)	Modified the description of "Resolution" in "(3) The cluster application becomes Faulted or Inconsistent when the network error occurs."	"D.1.6 Errors in the Mirroring Among Servers"
58	GDS	J2UZ-7243-05ENZ0(01) J2UZ-7243-05ENZ2(01)	Modified the description of "Resolution" in "(4) The slice becomes INVALID or the operation stops after restarting the other node when the network error occurs."	"D.1.6 Errors in the Mirroring Among Servers"
59	GDS	J2UZ-7243-05ENZ0(01) J2UZ-7243-05ENZ2(01)	Added "(5) The class gets closed down and the operation may stop if RMS on a standby node is stopped."	"D.1.6 Errors in the Mirroring Among Servers"
60	GDS	J2UZ-7243-05ENZ0(01) J2UZ-7243-05ENZ2(01)	Added "(5) A message about the sfdsksys module is output when the errata kernel is applied [RHEL7 or later]."	"D.1.14 OS Messages"
61	GDS	J2UZ-7243-05ENZ0(02) J2UZ-7243-05ENZ2(02)	Updated the patch numbers.	"3.2.1 Required patch"
62	GDS	J2UZ-7243-06ENZ0(00) J2UZ-7243-06ENZ2(00)	Added the note when mirroring the system disk in RHEL7 environment.	"3.5 System Disk Mirroring Preconditions"
63	GDS	J2UZ-7243-06ENZ0(00) J2UZ-7243-06ENZ2(00)	Changed the procedure for RHEL7 (Intel64) environment described in "Resolution" of "(4) System cannot be booted. (Failure of all boot disks)."	"D.1.5 System Disk Abnormality [EFI]"
64	GDS	J2UZ-7243-06ENZ0(00) J2UZ-7243-06ENZ2(00)	Added the explanations of OS messages about Secure Boot.	"D.1.14 OS Messages"
65	GDS	J2UZ-7243-06ENZ0(00) J2UZ-7243-06ENZ2(00)	Added the troubleshooting for Secure Boot.	"D.1.15 Other Errors"
66	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added the notes.	"1.2.1.4 Mirroring among Servers"
67	GDS, GDS Snapshot	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Changed the title of each section.	"3.3 Dell EMC Storage Units" "4.10 Settings to Use Dell EMC PowerPath"

No.	Product	Manual code	Description	Chapter/Section
				"6.1.3 To Use Dell EMC Storage Units" "D.1.9 Errors in Dell EMC Storage Units"
68	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Add the note "Sharing the network with the public LAN."	"3.9.1 Network Configuration"
69	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added the note on the disk area that cannot be specified as the target of the mirroring among servers.	"3.9.2 Disk"
70	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added the notes.	"3.9.3 Object Configuration"
71	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Changed the maximum number of the iSCSI devices that are the target of the mirroring among servers.	"3.9.3 Object Configuration" "6.3.1 What Is Shared Disk Resource Registration?" "6.3.4 Shared Disk Unit Resource Registration Procedure"
72	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Changed the section title and added the descriptions of the cluster system settings.	"3.9.4 Configuration and Settings of the Cluster System"
73	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Changed the descriptions why the mirroring among servers is suitable for the operation if the amount of data update is small.	"3.9.5 Data Amount and Data Update Amount"
74	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Add the descriptions of the required time for copying processing of the mirroring among servers.	"3.9.6 Required Time for Copying Processing"
75	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added the description about the environment where the mirroring among servers cannot be used.	"3.9.7 Note"
76	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Changed the setting procedure to first perform "Setting Tuning Parameters."	"4.8 Disk Setting for Performing Mirroring among Servers"
77	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added the note on the number of iSCSI targets.	"4.8 Disk Setting for Performing Mirroring among Servers"
78	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Changed the values of the tuning parameters.	"4.8.1 Setting Tuning Parameters"
79	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Changed the setting.	"4.8.2 Checking and Setting Required Packages"
80	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Changed the parameters specified by the iscsiadm command.	"4.8.4 Establishing iSCSI Session"
81	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added the notes.	"4.8.4 Establishing iSCSI Session"
82	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added the description when performing mirroring per partition.	"6.1.4 To Use Mirroring among Servers"
83	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added the note when selecting the netmirror group to "Setting [JRM]."	"6.6.1.6 Volume Configuration"

No.	Product	Manual code	Description	Chapter/Section
84	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Changed the procedure to swap the disk that is connected to the netmirror group.	"7.3.3 Swapping Disks of Netmirror Group"
85	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added the information "I/O error on the disk of another node."	"7.16.1 I/O Error Due to Node Stop or Network Error"
86	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added the descriptions when stopping and restarting the node.	"7.16.2 Stopping or Restarting the Node"
87	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added the descriptions of operation when the LEFTCLUSTER node exists.	"7.16.3 Operation when the LEFTCLUSTER Node Exists"
88	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added the note when the disk containing the latest data is automatically selected.	"7.16.5 Lock Volume of the Netmirror Volume"
89	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added "Procedure."	"7.16.6 Operation on Only One Node"
90	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Deleted "7.15.7 Stopping RMS."	"7.15.7 Stopping RMS"
91	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added the descriptions when stopping both nodes.	"7.16.9 Stopping both Nodes"
92	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added "7.15.10 Failover of the Cluster Application when the Network is Disconnected."	"7.16.11 Switchover of the Cluster Application when the Network is Disconnected"
93	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added the descriptions of operations that cannot be performed when the stopped node exists.	"7.16.12 Configuration Change when the Node is Stopped"
94	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added the description of the display order of the object to "DESCRIPTION."	"B.1.6 sdxinfo - Display Object Configuration and Status Information"
95	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Changed the description of SKIP.	"B.1.6 sdxinfo - Display Object Configuration and Status Information"
96	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added the -e force option of the sdxfix -V command.	"B.1.9 sdxfix - Restore Failed Object"
97	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Deleted "C.2.1 Panic Message (20001)."	"C.2.1 Panic Message (20001)"
98	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added the description to "Explanation" of the error message 26000.	"C.2.5 Internal Error Messages (26000-26001)"
99	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Deleted the error message 42053.	"C.3.2 Error Messages (42000-42099)"
100	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added the error messages 42062 to 42065.	"C.3.2 Error Messages (42000-42099)"
101	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added the warning message 44041.	"C.3.3 Warning Messages (44000-44099)"
102	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added the information message 46120.	"C.3.4 Information Messages (46000-46199)"
103	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added the description about the mirroring among servers to	"C.4.1 Error Messages (60000-60499)"

No.	Product	Manual code	Description	Chapter/Section
			Explanation and Resolution of the error message 60046.	
104	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Changed "Resolution" of the error message 60144.	"C.4.1 Error Messages (60000-60499)"
105	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Changed "Explanation" of the error message 60396.	"C.4.1 Error Messages (60000-60499)"
106	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added the error messages 60418, and 60422 to 60425.	"C.4.1 Error Messages (60000-60499)"
107	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Changed "Note" of step 3-4) described in "Resolution" of "(4) System cannot be booted. (Failure of all boot disks)."	"D.1.5 System Disk Abnormality [EFI]"
108	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Changed troubleshooting.	"D.1.6 Errors in the Mirroring Among Servers"
109	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added "(5) The resource of the class where the netmirror volume exists becomes OFF-FAIL."	"D.1.6 Errors in the Mirroring Among Servers"
110	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added "(6) The class is closed on the last started node at the startup of both nodes."	"D.1.6 Errors in the Mirroring Among Servers"
111	GDS	J2UZ-7243-07ENZ0(00) J2UZ-7243-07ENZ2(00)	Added "(8) A message about the ext4 file system is output at node stop."	"D.1.14 OS Messages"
112	GDS	J2UZ-7243-07ENZ0(01) J2UZ-7243-07ENZ2(01)	Added the descriptions of the systems and the operations suitable for the mirroring among servers.	"1.2.1.4 Mirroring among Servers"
113	GDS	J2UZ-7243-07ENZ0(01) J2UZ-7243-07ENZ2(01)	Added the descriptions to the note "Sharing the network with the cluster interconnect or the administrative LAN."	"3.9.1 Network Configuration"
114	GDS	J2UZ-7243-07ENZ0(01) J2UZ-7243-07ENZ2(01)	Removed an unnecessary step of restarting the system.	"7.11.3 Restoring from Backup Disks"
115	GDS	J2UZ-7243-07ENZ0(01) J2UZ-7243-07ENZ2(01)	Added the descriptions of the procedure when upgrading Dell EMC PowerPath.	"8.9 Upgrading Procedure for Dell EMC PowerPath"
116	GDS	J2UZ-7243-07ENZ0(01) J2UZ-7243-07ENZ2(01)	Changed the descriptions of the note "When You Cannot Unmirror the System Disk."	"9.4.1.2 Operation Flow"
117	GDS	J2UZ-7243-07ENZ0(01) J2UZ-7243-07ENZ2(01)	Added the descriptions to the note "Definition of devices in the /etc/fstab file."	"9.4.1.2 Operation Flow"
118	GDS	J2UZ-7243-07ENZ0(01) J2UZ-7243-07ENZ2(01)	Added the descriptions to explain that access to the physical special files are not prevented when the disk is registered to the root class.	"A.1 Restraining Access to Physical Special File"
119	GDS	J2UZ-7243-07ENZ0(01) J2UZ-7243-07ENZ2(01)	Added the warning message 22016.	"C.2.3 Warning Messages (22000-22099)"
120	GDS	J2UZ-7243-07ENZ0(01) J2UZ-7243-07ENZ2(01)	Added the descriptions to "Explanation" and "Resolution" of the internal error message 26000.	"C.2.5 Internal Error Messages (26000-26001)"

No.	Product	Manual code	Description	Chapter/Section
121	GDS	J2UZ-7243-07ENZ0(01) J2UZ-7243-07ENZ2(01)	Added the error messages 60426 to 60430.	"C.4.1 Error Messages (60000-60499)"
122	GDS	J2UZ-7243-07ENZ0(01) J2UZ-7243-07ENZ2(01)	Added the warning message 62026.	"C.4.2 Warning Messages (62000-62099)"
123	GDS	J2UZ-7243-07ENZ0(01) J2UZ-7243-07ENZ2(01)	Changed the resolution of "(4) System cannot be booted. (Failure of all boot disks)" for RHEL6 and RHEL7 environments.	"D.1.5 System Disk Abnormality [EFI]"
124	GDS	J2UZ-7243-07ENZ0(01) J2UZ-7243-07ENZ2(01)	Added the description of how the symptom occurs to "Explanation" of "(8) System is not started after system disk is unmirrored. [RHEL6 or later]."	"D.1.5 System Disk Abnormality [EFI]"
125	GDS	J2UZ-7243-07ENZ0(01) J2UZ-7243-07ENZ2(01)	Added the descriptions of the procedure when restoring the file systems on the volume.	"D.1.11 File System Error"
126	GDS	J2UZ-7243-07ENZ0(02) J2UZ-7243-07ENZ2(02)	Added "(9) Messages of systemd or systemd-udev are output when the system is started or the operations of GDS are executed. [RHEL7]."	"D.1.14 OS Messages"

GDS: Global Disk Services

GDS Snapshot: Global Disk Services Snapshot

G.2.6 4.4A00 Manual Changes

No.	Product	Manual code	Description	Chapter/Section
1	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Added the note on a VMware environment.	"1.2.1.4 Mirroring among Servers"
2	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted the note about hot swap.	"1.2.1.4 Mirroring among Servers"
3	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted the note about external disk units.	"1.2.1.4 Mirroring among Servers"
4	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted the description of note about GDS 4.3A00.	"1.3.6.1 Logical Partitioning" "2.1.4 Logical Volume" "2.3.5 Number of Volumes"
5	GDS Snapshot	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted the description of GDS Snapshot 4.3A20 or earlier.	"1.5.2 Snapshot Function with No Load to Server/SAN" "A.8 Using Advanced Copy Function in Proxy Configuration"
6	GDS Snapshot	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted the description of RHEL4 and RHEL5.	"1.5.6 Creating Alternative Boot Environment [EFI]"
7	GDS Snapshot	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Specified that the shadow object is not supported.	"1.6 Shadow Volume" "3.8.7 Shadow Classes" "7.3.4 Swapping Physical Disks Registered with Shadow Classes"
8	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted the note of GDS 4.3A10/4.3A20.	"2.3.8 Disk Label"

No.	Product	Manual code	Description	Chapter/Section
9	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted the description about patches of GDS 4.3A10 from "1 TB or larger disks."	"2.3.9 Disk Size"
10	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted Note "Large Volumes [4.3A20 or earlier]."	"2.3.11 Volume Size"
11	GDS, GDS Snapshot	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted the description of GDS 4.3A10/4.3A20. Added the notes about the number of mpath devices.	"3.2 DM-MP (Device Mapper Multipath) "
12	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Added the description of the mirroring among servers.	"3.2 DM-MP (Device Mapper Multipath) " "4.5 Setting Up DM-MP"
13	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted the description of RHEL4, RHEL5, and PRIMEQUEST 500A/500/400 Series.	"3.5 System Disk Mirroring Preconditions"
14	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Specified that the hot spare function is not supported in the root class.	"3.8.1 Number of Disks within Class" "6.5.2.1.1 Operation Flow" "6.5.2.1.2 Operation Procedure"
15	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Changed the formula to estimate the indication of the required time.	"3.8.11 Required Time for Synchronization Copying Invoked by Hot Spare"
16	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Changed the description for using IP aliasing.	"3.9.1 Network Configuration"
17	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Added the description about an available disk and an unavailable disk for the mirroring among servers.	"3.9.2 Disk"
18	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Added the description about "Forced stop method in a VMware environment."	"3.9.4 Configuration and Settings of the Cluster System"
19	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Changed the description of "Setting shutdown facility."	"3.9.4 Configuration and Settings of the Cluster System"
20	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Changed the description about the environment where the mirroring among servers cannot be used.	"3.9.7 Note"
21	GDS, GDS Snapshot	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Added the description about the mirroring among servers and shared disk (VMDK). Deleted the description of a Xen environment.	"3.21 Virtual Environment"
22	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Added " Inhibiting the unnecessary message output of multipathd."	"4.5 Setting Up DM-MP"
23	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Added the notes when mirroring a partition.	"4.8 Disk Setting for Performing Mirroring among Servers"
24	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted unnecessary setting procedures.	"4.8.1 Setting Tuning Parameters"

No.	Product	Manual code	Description	Chapter/Section
25	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Added the procedure to check that the sg3_utils package is installed in RHEL6 environment.	"4.8.2 Checking and Setting Required Packages"
26	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Added the setting procedures in RHEL7 environment.	"4.8.2 Checking and Setting Required Packages" "4.8.3 Creating iSCSI Target"
27	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Added the description in the restart procedure of the SCSI target daemon (tgt) in RHEL6 environment.	"4.8.3 Creating iSCSI Target"
28	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Changed the setting procedure in "Establishing iSCSI Session."	"4.8.4 Establishing iSCSI Session"
29	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted "4.7 Configuring Xen Environment."	"4.7 Configuring Xen Environment"
30	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Changed the description about setting of the virtual machine configuration parameter: disk.EnableUUID.	"4.7 Configuring VMware Environment"
31	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted "File System Configuration [4.3A20 or earlier]."	"5.3.2 [Settings] Menu"
32	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted the description about patches of GDS 4.3A30.	"6.2 Using iSCSI Disk"
33	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted the description about the restriction of registering the pair of the iSCSI devices.	"6.3.1 What Is Shared Disk Resource Registration?" "6.3.4 Shared Disk Unit Resource Registration Procedure"
34	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted "2) Applying patches for adapting to changes of SCSI host ID [4.3A00]."	"6.5 System Disk Mirroring Settings [EFI]"
35	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted "6.5.2 Applying Patches for Adapting to Changes of SCSI Host ID [4.3A00]."	"6.5.2 Applying Patches for Adapting to Changes of SCSI Host ID [4.3A00]"
36	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted the description of RHEL4 and RHEL5.	"6.5.3 Checking System Disk Settings" "6.5.4 Checking Physical Disk Information and Slice Numbers" "7.4.2 Backing Up [EFI]" "7.4.5 Restoring (When System Can Be Booted) [EFI]" "7.4.6 Restoring (When System Cannot Be Booted) [EFI]" "9.5 OS Re-installation" "B.1.11 sdxroot - Root File System Mirroring Definition and Cancellation [EFI]" "B.2.1 sdxproxy - Proxy Object Operations"

No.	Product	Manual code	Description	Chapter/Section
37	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted the description for 4.3A20 or earlier in NOTE "Node Name."	"6.6.1.4 Cluster System Class Configuration"
38	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted "6.7.1 Setting with GDS Management View [4.3A20 or earlier]."	"6.7.1 Setting with GDS Management View [4.3A20 or earlier]."
39	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Changed the title.	"6.7.1 Creating the File System"
40	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted Information "Disk Swap of PRIMEQUEST 500A/500/400 series [4.3A00]."	"7.3.1.2 Operation Procedure"
41	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Added the procedures to swap a disk connected to the netmirror group and the procedures to swap a disk in a RHEL7.	"7.3.3 Swapping Disks of Netmirror Group"
42	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted "7.3.5.4 In VMware environment."	"7.3.5.4 In VMware environment"
43	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Changed step b3).	"7.4.2 Backing Up [EFI]"
44	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Added the description about how to perform the procedure to be performed in single user mode in the cloud environment.	"7.17 Maintenance in Cloud Environment"
45	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted "8.3 Adding and Removing Disks [4.3A00]."	"8.3 Adding and Removing Disks [4.3A00]"
46	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted "8.4 Hot Adding Disk [RHEL6 or later]."	"8.4 Hot Adding Disk [RHEL6 or later]"
47	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Added the procedures to change IP addresses used for mirroring among servers.	"8.10 Changing IP Addresses Used for Mirroring among Servers"
48	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Added the procedures of adding disks used for mirroring among servers.	"8.11 Adding Disks Used for Mirroring among Servers"
49	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted "9.2.1 Using GDS Management View [4.3A20 or earlier]."	"9.2 Removing File System"
50	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted the description for GDS 4.3A20 or earlier.	"9.3.1.2 Removing Volume"
51	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Changed the note "Definition of devices in the /etc/fstab file."	"9.4.1.2 Operation Flow"
52	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Added the procedures to delete iSCSI settings used for mirroring among servers.	"9.6 Deleting iSCSI Settings Used for Mirroring among Servers"
53	GDS Snapshot	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Specified that the command of the shadow class is not supported.	"B.2.2 sdxshadowdisk - Shadow Disk Operations" "B.2.3 sdxshadowgroup - Shadow Group Operations" "B.2.4 sdxshadowvolume - Shadow Volume Operations"

No.	Product	Manual code	Description	Chapter/Section
54	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Changed EXAMPLES.	"B.3.2 cldelrsc(8)-Delete resource"
55	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Changed the description about the warning message 22016.	"C.2.3 Warning Messages (22000-22099)"
56	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted the description about RHEL5.5 or earlier from Resolution of the error message 60181.	"C.4.1 Error Messages (60000-60499)"
57	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted the description about GDS 4.3A10 from Explanation and Resolution of the error message 60379.	"C.4.1 Error Messages (60000-60499)"
58	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Added the error message 60431, 60432.	"C.4.1 Error Messages (60000-60499)"
59	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted the error messages 5010 and 5011, because they are not output in this version.	"C.5.1 Error Messages (5000-5099)"
60	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted the description about GDS 4.3A20 or earlier of the warning message 5119.	"C.5.2 Warning Messages (5000,5100-5199)"
61	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted the description about RHEL4 and RHEL5 to 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."	"D.1.1 Slice Status Abnormality"
62	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Changed c) of "Resolution" in "(1) Disk is in DISABLE status."	"D.1.2 Disk Status Abnormality"
63	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Changed the procedure to extract the initrd file in Resolution of "(4) System cannot be booted. (Failure of all boot disks)."	"D.1.5 System Disk Abnormality [EFI]"
64	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted the description about RHEL4 and RHEL5 from Resolution of "(4) System cannot be booted. (Failure of all boot disks)."	"D.1.5 System Disk Abnormality [EFI]"
65	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted "(8) System is not started after system disk is unmirrored. [RHEL6 or later]."	"D.1.5 System Disk Abnormality [EFI]"
66	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Added the description when the cluster application becomes Inconsistent in "(2) The cluster application becomes Faulted or Inconsistent after the node is restarted."	"D.1.6 Errors in the Mirroring Among Servers"
67	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Added "(7) The node is restarted while hot swap of disks is proceeding. [RHEL6]."	"D.1.6 Errors in the Mirroring Among Servers"
68	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Added "(8) The node is restarted while hot swap on a disk is proceeding. [RHEL7]"	"D.1.6 Errors in the Mirroring Among Servers"

No.	Product	Manual code	Description	Chapter/Section
69	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Added "(9) The message regarding an unmount is output when both nodes stop."	"D.1.6 Errors in the Mirroring Among Servers"
70	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Added "(6) The error message is displayed when starting GDS Management View."	"D.1.7 GDS Management View Abnormality"
71	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted the description about RHEL4 and RHEL5 from Resolution of "(5) The error message "ERROR: class: cannot operate shared objects, ..." is output, and the shared class cannot be created; or the error message "ERROR: cluster communication failure" is output, and automatic registration of resources is failed."	"D.1.10 Cluster System Related Error"
72	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Deleted the description about RHEL4 and RHEL5 from Resolution of "(3) Restoring the file system created on the root volume."	"D.1.11 File System Error"
73	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Added "(10) Messages of multipathd are output when creating volumes."	"D.1.14 OS Messages"
74	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Changed the description of VMware vCenter Server described in "Note."	"1.2.1.4 Mirroring among Servers"
75	GDS	J2UL-2107-01ENZ0(00) J2UL-2107-01ENZ2(00)	Changed the description of "Forced stop method in a VMware environment."	"3.9.4 Configuration and Settings of the Cluster System"
76	GDS	J2UL-2107-02ENZ0(00) J2UL-2107-02ENZ2(00)	Added RDM in the virtual compatible mode to the explanation for the virtual disk.	"4.7 Configuring VMware Environment"
77	GDS	J2UL-2107-02ENZ0(00) J2UL-2107-02ENZ2(00)	Changed the setting procedures in RHEL7 environment.	"4.8.2 Checking and Setting Required Packages" "4.8.3 Creating iSCSI Target" "7.3.3.1 Hot Swap" "7.3.3.2 Cold Swap" "8.10 Changing IP Addresses Used for Mirroring among Servers" "8.11 Adding Disks Used for Mirroring among Servers" "9.6.1 Delete iSCSI devices" "9.6.3 Deleting iSCSI Target Settings"
78	GDS	J2UL-2107-02ENZ0(00) J2UL-2107-02ENZ2(00)	Changed to the procedure to be used in RHEL7 environment.	"4.8.5 Confirming the Correspondence of iSCSI Device Name of Each Node"
79	GDS	J2UL-2107-02ENZ0(00) J2UL-2107-02ENZ2(00)	Added the procedure for a rescanning of iSCSI session.	"7.3.3.1 Hot Swap (RHEL6)" "7.3.3.2 Cold Swap (RHEL6)" "7.3.3.1 Hot Swap" "7.3.3.2 Cold Swap"

No.	Product	Manual code	Description	Chapter/Section
80	GDS	J2UL-2107-02ENZ0(00) J2UL-2107-02ENZ2(00)	Changed the procedures described in "(7) The node is restarted while hot swap of disks is proceeding. [RHEL6]."	"D.1.6 Errors in the Mirroring Among Servers"
81	GDS	J2UL-2107-02ENZ0(00) J2UL-2107-02ENZ2(00)	Added the steps for the procedures described in "(8) The node is restarted while hot swap of disks is proceeding. [RHEL7]."	"D.1.6 Errors in the Mirroring Among Servers"
82	GDS	J2UL-2107-02ENZ0(01) J2UL-2107-02ENZ2(01)	Added the description for the root class to "Size of a private slice."	"2.3.9 Disk Size"
83	GDS	J2UL-2107-02ENZ0(01) J2UL-2107-02ENZ2(01)	Added the note when mirroring the system disk in RHEL6 environment.	"3.5 System Disk Mirroring Preconditions"
84	GDS	J2UL-2107-02ENZ0(01) J2UL-2107-02ENZ2(01)	Added the description about the virtual disk on KVM guest.	"4.6 Configuring KVM Environment"
85	GDS	J2UL-2107-02ENZ0(01) J2UL-2107-02ENZ2(01)	Changed the procedure to specify the parameter of iSCSI session.	"4.8.4 Establishing iSCSI Session"
86	GDS	J2UL-2107-02ENZ0(01) J2UL-2107-02ENZ2(01)	Changed the description about the description of the Excluded Device List in mirroring per partition.	"6.1.4 To Use Mirroring among Servers"
87	GDS	J2UL-2107-02ENZ0(01) J2UL-2107-02ENZ2(01)	Added the note "Backing up system disk."	"6.5 System Disk Mirroring Settings [EFI]"
88	GDS	J2UL-2107-02ENZ0(01) J2UL-2107-02ENZ2(01)	Changed the procedure to check the possibility of device name change.	"7.3.5.4 Swapping Internal Disks Registered with Root Classes or Local"
89	GDS	J2UL-2107-02ENZ0(01) J2UL-2107-02ENZ2(01)	Added the note when the system disk mirroring is configured again to "Note."	"7.4 Backing Up and Restoring System Disk"
90	GDS	J2UL-2107-02ENZ0(01) J2UL-2107-02ENZ2(01)	Changed the path format of device special file in steps b3) and b4).	"7.4.2 Backing Up [EFI]"
91	GDS	J2UL-2107-02ENZ0(01) J2UL-2107-02ENZ2(01)	Added the description to delete the iSCSI device from the resource database before taking the procedure.	"9.6.1 Delete iSCSI devices"
92	GDS	J2UL-2107-02ENZ0(01) J2UL-2107-02ENZ2(01)	Added error messages 60441 to 60443.	"C.4.1 Error Messages (60000-60499)"
93	GDS	J2UL-2107-02ENZ0(01) J2UL-2107-02ENZ2(01)	Changed the corrective action for disk failure in "(1) Disk is in DISABLE status."	"D.1.2 Disk Status Abnormality"
94	GDS	J2UL-2107-02ENZ0(01) J2UL-2107-02ENZ2(01)	Added "(11) The message about SELinux is output."	"D.1.14 OS Messages"
95	GDS	J2UL-2107-02ENZ0(02) J2UL-2107-02ENZ2(02)	Added a description to "Information."	"7.16.7 Restoration to 2-node Operation after Operation on Only One Node"
96	GDS	J2UL-2107-02ENZ0(02) J2UL-2107-02ENZ2(02)	Changed the description of restoration procedure.	"7.16.8 Restoration when the Latest Disk Cannot Be Selected Automatically"
97	GDS	J2UL-2107-02ENZ0(02) J2UL-2107-02ENZ2(02)	Added the note to "-e force (when using -V)."	"B.1.9 sdxfix - Restore Failed Object"

No.	Product	Manual code	Description	Chapter/Section
98	GDS	J2UL-2107-02ENZ0(02) J2UL-2107-02ENZ2(02)	Specified that internal disk is a target.	"B.1.14 sdxdevinfo - Display Device Information of Internal Disk"
99	GDS	J2UL-2107-02ENZ0(02) J2UL-2107-02ENZ2(02)	Added error messages 42067 to 42070.	"C.3.2 Error Messages (42000-42099)"
100	GDS	J2UL-2107-02ENZ0(02) J2UL-2107-02ENZ2(02)	Changed the description of warning message 44041.	"C.3.3 Warning Messages (44000-44099)"
101	GDS	J2UL-2107-02ENZ0(02) J2UL-2107-02ENZ2(02)	Added error messages 60439 and 60440.	"C.4.1 Error Messages (60000-60499)"
102	GDS	J2UL-2107-02ENZ0(02) J2UL-2107-02ENZ2(02)	Changed the resolution of "(5) The resource of the class where the netmirror volume exists becomes OFF-FAIL."	"D.1.6 Errors in the Mirroring Among Servers"
103	GDS	J2UL-2107-02ENZ0(02) J2UL-2107-02ENZ2(02)	Changed the description of "(11) The message about SELinux is output."	"D.1.14 OS Messages"
104	GDS	J2UL-2107-02ENZ0(03) J2UL-2107-02ENZ2(03)	Added "2.3.10 Group Size."	"2.3.10 Group Size"
105	GDS	J2UL-2107-02ENZ0(03) J2UL-2107-02ENZ2(03)	Added the note "When using GDS on both host and guest in the KVM environment."	"3.21 Virtual Environment"
106	GDS	J2UL-2107-02ENZ0(03) J2UL-2107-02ENZ2(03)	Added a note to the procedure in RHEL6 environment.	"4.8.2 Checking and Setting Required Packages"
107	GDS	J2UL-2107-02ENZ0(03) J2UL-2107-02ENZ2(03)	Added the information "Disk Size."	"6.6.1.6 Volume Configuration"
108	GDS	J2UL-2107-02ENZ0(03) J2UL-2107-02ENZ2(03)	Added a description of file system expansion.	"7.14.2 Volume Expansion Using Commands [EFI]"
109	GDS	J2UL-2107-02ENZ0(03) J2UL-2107-02ENZ2(03)	Added a description of restarting both nodes.	"7.16.2 Stopping or Restarting the Node"
110	GDS	J2UL-2107-02ENZ0(03) J2UL-2107-02ENZ2(03)	Added a description about the status after both nodes are started.	"7.16.5 Lock Volume of the Netmirror Volume"
111	GDS	J2UL-2107-02ENZ0(03) J2UL-2107-02ENZ2(03)	Added a condition of changing volume name to the note "Preconditions for Changing Volume Attributes."	"8.2.1.4 Volume Configuration"
112	GDS	J2UL-2107-02ENZ0(03) J2UL-2107-02ENZ2(03)	Changed the note "Changes in Special File Path Name by Changing Volume Name."	"8.2.1.4 Volume Configuration"
113	GDS	J2UL-2107-02ENZ0(03) J2UL-2107-02ENZ2(03)	Added the note when upgrading scsi-target-utils package.	"8.12 Upgrading scsi-target-utils Package [RHEL6]"
114	GDS	J2UL-2107-02ENZ0(03) J2UL-2107-02ENZ2(03)	Added a description of the effect when restoration procedure 1) is not executed.	"9.5 OS Re-installation"
115	GDS	J2UL-2107-02ENZ0(03) J2UL-2107-02ENZ2(03)	Added a description to the "Resolution" of warning message 22010.	"C.2.3 Warning Messages (22000-22099)"
116	GDS	J2UL-2107-02ENZ0(03) J2UL-2107-02ENZ2(03)	Changed the explanation and resolution of error message 60037.	"C.4.1 Error Messages (60000-60499)"

No.	Product	Manual code	Description	Chapter/Section
117	GDS	J2UL-2107-02ENZ0(03) J2UL-2107-02ENZ2(03)	Changed the explanation and resolution of warning message 5123.	"C.5.2 Warning Messages (5000,5100-5199)"
118	GDS	J2UL-2107-02ENZ0(03) J2UL-2107-02ENZ2(03)	Changed the "Explanation" of "(1) Messages are output at node startup."	"D.1.14 OS Messages"
119	GDS	J2UL-2107-02ENZ0(03) J2UL-2107-02ENZ2(03)	Changed the description of collecting investigation material.	"D.2 Collecting Investigation Material"
120	GDS	J2UL-2107-02ENZ0(04) J2UL-2107-02ENZ2(04)	Added "4.6.4 Setting up virtio Block Device Used for Mirroring among Servers."	"4.6.4 Setting up virtio Block Device Used for Mirroring among Servers"
121	GDS	J2UL-2107-02ENZ0(04) J2UL-2107-02ENZ2(04)	Deleted unnecessary descriptions designated for the KVM guest.	"4.8.3 Creating iSCSI Target" "7.3.3.1 Hot Swap" "7.3.3.2 Cold Swap" "8.11 Adding Disks Used for Mirroring among Servers"
122	GDS	J2UL-2107-02ENZ0(04) J2UL-2107-02ENZ2(04)	Changed descriptions designated for the KVM guest.	"7.3.3.1 Hot Swap (RHEL6)" "7.3.3.2 Cold Swap (RHEL6)"
123	GDS	J2UL-2107-02ENZ0(04) J2UL-2107-02ENZ2(04)	Add a new word "by-id file".	"Glossary"

GDS: Global Disk Services

GDS Snapshot: Global Disk Services Snapshot

G.2.7 4.5A00 Manual Changes

No.	Product	Manual code	Description	Chapter/Section
1	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Deleted the information of RHEL 7.0	"1.2.1.4 Mirroring among Servers" "3.9.7 Note"
2	GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added "1.7 GDS I/O Monitor Option Features."	"1.7 GDS I/O Monitor Option Features"
3	GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added "1.8 I/O Response Time Assurance Operations [IOMonitor]."	"1.8 I/O Response Time Assurance [IOMonitor]"
4	GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added I/O Response Time Assurance, I/O response time, and system panic mode to "Attributes."	"2.1.1 Disk Class"
5	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Deleted the x86 environment description from "Disk label types."	"2.3.8 Disk Label"
6	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Deleted the description of the OS version from "1TB or larger disks."	"2.3.9 Disk Size"
7	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the PRIMEQUEST 3000 Series to "Conditions for server" and deleted PRIMEQUEST 1000 Series.	"3.5 System Disk Mirroring Preconditions"
8	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the disk size condition to "Conditions for system disk configuration and settings."	"3.5 System Disk Mirroring Preconditions"
9	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the note for iSCSI boot environment to "Note."	"3.5 System Disk Mirroring Preconditions"

No.	Product	Manual code	Description	Chapter/Section
10	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Deleted the "Setting shutdown facility."	"3.9.4 Configuration and Settings of the Cluster System"
11	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the information "Red Hat OpenStack Platform."	"3.21 Virtual Environment"
12	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the note "VMware vSAN."	"3.21 Virtual Environment"
13	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the reference for note in "See."	"3.22 Disaster Recovery"
14	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Changed the procedure.	"3.22.1 Building Disaster Recovery System" "3.22.3 Switching to Disaster Recovery System" "3.22.4 Restoration to Operation System from Disaster Recovery System"
15	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the description of the procedure when synchronizing the data of operation system and disaster recovery system.	"3.22.2 Data Synchronization of Operation System and Disaster Recovery System"
16	GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added "3.23 I/O Response Time Assurance [IOMonitor]."	"3.23 I/O Response Time Assurance [IOMonitor]"
17	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the note for reinstalling OS.	"4.1 Notes on Installing OS"
18	GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added "4.4 Installing GDS I/O Monitor Option."	"4.4 Installing GDS I/O Monitor Option"
19	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the description when starting the screen with Java Web Start command.	"4.10.3 Setting up the Web Environment" "5.1.1 Starting Web-Based Admin View Operation Menu" "5.7 Exiting GDS Management View"
20	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Changed the description of host name.	"5.1.1 Starting Web-Based Admin View Operation Menu"
21	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Changed the note "When the Web-Based Admin View does not start."	"5.1.1 Starting Web-Based Admin View Operation Menu"
22	GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the change attribute (I/O Response Time Assurance) to "Class operation."	"5.6 Operation"
23	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Changed the description of "exdevtab.sh" information.	"6.1.3 To Use Dell EMC Storage Units"
24	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the note "For RHEL7.4 or later environment."	"6.5.2.1.2 Operation Procedure"
25	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the description when automatically unmounting file system on exiting OS to information "For Local Classes."	"6.7.2 File System Auto Mount"
26	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Change the description on how to change the access permission.	"6.9 Volume's Block Special File Access Permission"

No.	Product	Manual code	Description	Chapter/Section
27	GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added "6.13 I/O Response Time Assurance [IOmonitor]."	"6.12 I/O Response Time Assurance [IOmonitor]"
28	GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added "6.13.1 Setting I/O Response Time Assurance."	"6.12.1 Setting I/O Response Time Assurance"
29	GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added "6.13.2 Hang-up Monitoring Function."	"6.12.2 Hang-up Monitoring Function"
30	GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the note "If the I/O Response Time Assurance Function is set."	"7.3.1.2 Operation Procedure"
31	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the note "When using the mirroring among servers function."	"7.4 Backing Up and Restoring System Disk"
32	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added "7.4.1 Settings Before Backing Up."	"7.4.1 Settings Before Backing Up"
33	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the settings procedure for before and after backing up.	"7.4.2 Backing Up [EFI]"
34	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added "7.4.3 Settings After Backing Up."	"7.4.3 Settings After Backing Up"
35	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added "7.4.4 Settings Before Restoring."	"7.4.4 Settings Before Restoring"
36	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the setting procedure for after restoration.	"7.4.5 Restoring (When System Can Be Booted) [EFI]" "7.4.6 Restoring (When System Cannot Be Booted) [EFI]"
37	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added "7.4.7 Settings After Restoring."	"7.4.7 Settings After Restoring"
38	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added "See."	"7.13 Backing Up and Restoring Object Configurations"
39	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the necessary procedures in cluster system.	"7.13.2 Restoring"
40	GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added "7.15 I/O Response Time Assurance [IOmonitor]."	"7.15 I/O Response Time Assurance [IOmonitor]"
41	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added "Note."	"7.16.1 I/O Error Due to Node Stop or Network Error"
42	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the note "Recreating shared class registered in a cluster application."	"8.2 Changing Object Configuration"
43	GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added "8.3 Changing the Settings of I/O Response Time Assurance [IOmonitor]."	"8.3 Changing the Settings of I/O Response Time Assurance [IOmonitor]"
44	GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added "8.4 Deleting the Settings of the I/O Response Time Assurance [IOmonitor]"	"8.4 Deleting the Settings of the I/O Response Time Assurance [IOmonitor]"
45	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the reference for note in "See."	"8.6 Server Migration" "8.7.2 Using Copy Function of Storage Units and sdxconfig Command"
46	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the note of I/O Response Time Assurance Function.	"9.3.2 Using Commands" "9.4 System Disk Mirroring Cancellation [EFI]"

No.	Product	Manual code	Description	Chapter/Section
				"A.5 Backing Up and Restoring Object Configuration (sdxcfg)"
47	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Changed the message in restoration procedure 2).	"9.5 OS Re-installation"
48	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the description for necessary procedure when restoring the object configuration of the shared class registered in the cluster application.	"A.5 Backing Up and Restoring Object Configuration (sdxcfg)"
49	GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the command to be used when performing the root class operation.	"A.11 Root Class Operation [EFI]"
50	GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Changed the description of scope sub options.	"B.1.2 sdxdisk - Disk Operations" "B.1.7 sdxcfg - Change Objects Attributes"
51	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the description of -a iotimeout option.	"B.1.2 sdxdisk - Disk Operations" "B.1.7 sdxcfg - Change Objects Attributes"
52	GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added notes about multiple and continuous executions to "DESCRIPTION."	"B.1.6 sdxcfg - Display Object Configuration and Status Information"
53	GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the description of -I option.	"B.1.6 sdxcfg - Display Object Configuration and Status Information"
54	GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the node when using I/O Response Time Assurance Function in the description of -I option.	"B.1.8 sdxcfg - Swap Disk"
55	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added "Note" to "Description."	"B.1.11 sdxcfg - Root File System Mirroring Definition and Cancellation [EFI]"
56	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the description for cluster system in the description of -S primary options.	"B.1.12 sdxcfg - Configuration Parameter Operations"
57	GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the note about setting I/O Response Time Assurance in the description of Restore option.	"B.1.13 sdxcfg - Object Configuration Operations"
58	GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added "C.2.1 Panic Messages 20000 - 20099."	"C.2.1 Panic Messages (20000-20099)"
59	GDS, GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added warning messages 22017 to 22023.	"C.2.3 Warning Messages (22000-22099)"
60	GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added information messages 24015 to 24017.	"C.2.4 Information Messages (24000-24099)"
61	GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added error messages 42066, 42071 to 42077.	"C.3.2 Error Messages (42000-42099)"
62	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the description when an internal error message is output during installation of GDS.	"C.3.5 Internal Error Messages (48000)"

No.	Product	Manual code	Description	Chapter/Section
63	GDS, GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Changed the explanation and resolution of error message 60379.	"C.4.1 Error Messages (60000-60499)"
64	GDS, GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added error messages 60441 to 60446, 60450.	"C.4.1 Error Messages (60000-60499)"
65	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the description of PRIMEQUEST 3000 Series to the "Resolution" in "(4) System cannot be booted. (Failure of all boot disks)."	"D.1.5 System Disk Abnormality [EFI]"
66	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Deleted the information of RHEL 6.0 and description of PRIMEQUEST 1000 Series from "Resolution" in "(4) System cannot be booted. (Failure of all boot disks)."	"D.1.5 System Disk Abnormality [EFI]"
67	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Changed "Resolution" in "(1) The netmirror slice that configures the netmirror volume is in INVALID state."	"D.1.6 Errors in the Mirroring Among Servers"
68	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Deleted the description of RHEL6(x86) in "Resolution" of "(5) The error message "ERROR: class: cannot operate shared objects," is output, and the shared class <i>class</i> cannot be created; or the error message "ERROR: cluster communication failure" is output, and automatic registration of resources is failed."	"D.1.10 Cluster System Related Error"
69	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Changed the startup procedure of cluster application in "Resolution" of "(6) The disk space of a file system on a shared disk is "Full" (100%)."	"D.1.10 Cluster System Related Error"
70	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Changed the description in "(4) A devlabel message is output at node startup or during collection of investigation material."	"D.1.14 OS Messages"
71	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Changed the description in "(6) Messages are output at node startup, and GDS cannot be started. [RHEL7]."	"D.1.14 OS Messages"
72	GDS I/O Monitor Option	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added the description of sdxhealthcheckd daemon in "Q: I want to know GDS resident daemons to monitor resident processes."	"E.1 Operation Design"
73	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added "Appendix F Startup Scripts and Startup Daemons in GDS [RHEL6]"	"Appendix F Startup Scripts and Startup Daemons in GDS [RHEL6]"

No.	Product	Manual code	Description	Chapter/Section
74	GDS	J2UL-2266-01ENZ0(00) J2UL-2266-01ENZ2(00)	Added "Appendix G Systemd Services and Startup Daemons in GDS [RHEL7]"	"Appendix F Systemd Services and Startup Daemons in GDS"
75	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Added the descriptions of RHOSP environment.	"1.2.1.3 Shared Disk Mirroring" "3.9.2 Disk"
76	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Changed the note about Symfoware.	"1.2.1.4 Mirroring among Servers"
77	GDS I/O Monitor Option	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Added cloud and RHOSP to the virtual environments described in the Note.	"1.7 GDS I/O Monitor Option Features"
78	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Added the notes of cloud environment and RHOSP environment.	"3.2 DM-MP (Device Mapper Multipath) "
79	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Added "Capacity Expansion of Disk Registered in the Class."	"3.12 Online Volume Expansion"
80	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Added the descriptions of usage in the K5 environment and RHOSP environment.	"3.21 Virtual Environment"
81	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Added the description "Before deleting the shared class, change the shared class to the local class by the sdxattr -C command." to step 2.	"3.22.2 Data Synchronization of Operation System and Disaster Recovery System"
82	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Added the description "Before deleting the shared class, change the shared class to the local class by the sdxattr -C command." to step 3.	"3.22.4 Restoration to Operation System from Disaster Recovery System"
83	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Added the descriptions of K5 environment and RHOSP environment.	"4.8.1 Setting Tuning Parameters"
84	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Added the Note to "Update Physical Disk Information."	"5.3.3 [Operation] Menu"
85	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Added the Note.	"6.4 Updating Physical Disk Information"
86	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Deleted the note "When using the mirroring among servers function."	"7.4 Backing Up and Restoring System Disk"
87	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Added the explanation for the reason to set the configuration parameter SDXDB_FAIL_NUM=0.	"7.4.1 Settings Before Backing Up"
88	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Added the necessary procedure when restoring the system disk for each node one by one in the mirroring among servers.	"7.4.4 Settings Before Restoring"
89	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Specified the unnecessary procedure when restoring the system disk for each node one by one in the mirroring among servers.	"7.4.7 Settings After Restoring"
90	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Added the description of RHOSP environment to [Maximum value of	"7.16.3 Operation when the LEFTCLUSTER Node Exists"

No.	Product	Manual code	Description	Chapter/Section
			the time to suspend the processes (default)].	
91	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Changed the Note "Changes in Special File Path Name by Changing Volume Name."	"8.2.1.4 Volume Configuration"
92	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Changed the Note "Changes in Special File Path Name by Changing Volume Name."	"B.1.6 sdinfo - Display Object Configuration and Status Information"
93	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Added the conditions necessary to restart the node to the descriptions of PRIMARY OPTIONS Restore.	"B.1.13 sdconfig - Object Configuration Operations"
94	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Added the internal error message 26001.	"C.2.5 Internal Error Messages (26000-26001)"
95	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Added g) to the Explanation and Resolution of the internal error message 48000.	"C.3.5 Internal Error Messages (48000)"
96	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Added the descriptions to Resolution of the error message 60126.	"C.4.1 Error Messages (60000-60499)"
97	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Added the note "Shutting down the node" to Resolution of "(2) Class cannot be started when booting the system."	"D.1.4 Class Status Abnormality"
98	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Changed the procedure of Resolution in "(1) The netmirror slice that configures the netmirror volume is in INVALID state."	"D.1.6 Errors in the Mirroring Among Servers"
99	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Changed the procedure of Resolution in "(4) The slice becomes INVALID or the operation stops after restarting the other node when the network error occurs."	"D.1.6 Errors in the Mirroring Among Servers"
100	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Added "(12) Error messages of modprobe and systemd are output and GDS is not started when the system is started. [RHEL7]."	"D.1.14 OS Messages"
101	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Deleted "D.1.15 ACM Abnormality."	"D.1.15 ACM Abnormality"
102	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Changed the explanation of sdxsnap.sh.	"D.2 Collecting Investigation Material" "D.2.2 Collecting Investigation Material with sdxsnap.sh Script"
103	GDS	J2UL-2266-01ENZ0(01) J2UL-2266-01ENZ2(01)	Added the note about by-id file.	"D.2 Collecting Investigation Material"
104	GDS	J2UL-2266-01ENZ0(02) J2UL-2266-01ENZ2(02)	Added ERROR in "Explanation" table.	"C.2 Driver Messages"
105	GDS	J2UL-2266-01ENZ0(02) J2UL-2266-01ENZ2(02)	Added the error message 21000.	"C.2.2 Error Messages (21000-21099)"

G.2.8 4.5A10 Manual Changes

No.	Product	Manual code	Description	Change/Section
1	GDS	J2UL-2266-02ENZ0(00) J2UL-2266-02ENZ2(00)	Added the description for the case of a class where a netmirror volume is created to "Size of a private slice."	"2.3.9 Disk Size"
2	GDS	J2UL-2266-02ENZ0(00) J2UL-2266-02ENZ2(00)	Added the Note to "7. Register iSCSI initiators of both nodes."	"4.8.3 Creating iSCSI Target"
3	GDS	J2UL-2266-02ENZ0(00) J2UL-2266-02ENZ2(00)	Described the necessity of the settings before registering the disk to the resource database.	"4.10 Settings to Use Dell EMC PowerPath"
4	GDS	J2UL-2266-02ENZ0(00) J2UL-2266-02ENZ2(00)	Deleted "6.10 NFS Mount."	"6.10 NFS Mount"
5	GDS Snapshot	J2UL-2266-02ENZ0(00) J2UL-2266-02ENZ2(00)	Added the procedure necessary for starting in single user mode to "4) Parting the proxy groups."	"7.5.6 Backing Up"
6	GDS Snapshot	J2UL-2266-02ENZ0(00) J2UL-2266-02ENZ2(00)	Added the procedure necessary after parting the proxy groups in single user mode to "6) Verifying the alternative boot environment."	"7.5.6 Backing Up"
7	GDS Snapshot	J2UL-2266-02ENZ0(00) J2UL-2266-02ENZ2(00)	Added the procedure necessary after parting the proxy groups in single user mode to "8) Switching to the alternative boot environment."	"7.5.7 Restoring"
8	GDS	J2UL-2266-02ENZ0(00) J2UL-2266-02ENZ2(00)	Described the conditions that the entire slice becomes the copy target.	"7.16.1 I/O Error Due to Node Stop or Network Error"
9	GDS	J2UL-2266-02ENZ0(00) J2UL-2266-02ENZ2(00)	Changed the description location of driver's major number change method.	"8.8 Changing Driver Major Number"
10	GDS	J2UL-2266-02ENZ0(00) J2UL-2266-02ENZ2(00)	Added the <i>driver</i> variable to "Variable Names" table.	"C.2 Driver Messages"
11	GDS	J2UL-2266-02ENZ0(00) J2UL-2266-02ENZ2(00)	Added the error message 21001.	"C.2.2 Error Messages (21000-21099)"
12	GDS	J2UL-2266-02ENZ0(00) J2UL-2266-02ENZ2(00)	Added "(13) The device name of GDS volume is displayed in the <i>sfdisk.X</i> format, therefore the message for which volume is unknown."	"D.1.14 OS Messages"
13	GDS	J2UL-2266-02ENZ0(01) J2UL-2266-02ENZ2(01)	Changed the "Resolution" of error message 21001.	"C.2.2 Error Messages (21000-21099)"
14	GDS	J2UL-2266-02ENZ0(01) J2UL-2266-02ENZ2(01)	Reflected the procedure for the patch for PH15096 is applied in the "Resolution" in "(4) System cannot be booted. (Failure of all boot disks)."	"D.1.5 System Disk Abnormality [EFI]"
15	GDS	J2UL-2266-02ENZ0(02) J2UL-2266-02ENZ2(02)	Changed the description of devices that cannot be used for the mirroring among servers.	"3.9.2 Disk"

No.	Product	Manual code	Description	Change/Section
16	GDS	J2UL-2266-02ENZ0(02) J2UL-2266-02ENZ2(02)	Changed the procedure for swapping a disk connected to a netmirror group.	7.3.3.1 Hot Swap (RHEL6) 7.3.3.2 Cold Swap (RHEL6) "7.3.3.1 Hot Swap" "7.3.3.2 Cold Swap"
17	GDS	J2UL-2266-02ENZ0(02) J2UL-2266-02ENZ2(02)	Added the error message 60451.	"C.4.1 Error Messages (60000-60499)"

GDS: Global Disk Services

GDS Snapshot: Global Disk Services Snapshot

G.2.9 4.6A00 Manual Changes

No.	Product	Manual code	Description	Change/Section
1	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Deleted RHEL6 from the supported OS and added RHEL8.	All chapters
2	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Changed the name from "FUJITSU Cloud Service K5" to "FUJITSU Cloud Service for OSS."	All chapters
3	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Added the descriptions that striping and concatenation of the disk group are not supported in RHEL8.	All chapters
4	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Deleted the note when the LEFTCLUSTER node exists from "Note."	"1.2.1.4 Mirroring among Servers"
5	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Added the description when a slice is netmirror type to the descriptions of "ACTIVE" and "STOP" status.	"2.1.5 Logical Slice"
6	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Changed the note when mirroring the system disk.	"3.5 System Disk Mirroring Preconditions"
7	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Added the devices with which virtual disk mirroring is enabled on a guest.	"3.21.1 KVM Environment and VMware Environment"
8	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Changed the description of "See."	"3.21.2 Cloud Environment"
9	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Added "Note" and "See" to "Operation mode."	"3.9.4 Configuration and Settings of the Cluster System"
10	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Changed the setting procedure of the tuning parameter.	"4.8.1 Setting Tuning Parameters"
11	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Changed the description of operation.	"7.16.3 Operation when the LEFTCLUSTER Node Exists"
12	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Added "7.16.4 Restoration from the LEFTCLUSTER State."	"7.16.4 Restoration from the LEFTCLUSTER State"
13	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Changed the restoration procedure.	"7.16.8 Restoration when the Latest Disk Cannot Be Selected Automatically"
14	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Changed the description of operations in the configuration of mirroring among servers.	"7.16.12 Configuration Change when the Node is Stopped"

No.	Product	Manual code	Description	Change/Section
15	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Changed the descriptions when maintaining the system in the cloud environment.	"7.17 Maintenance in Cloud Environment"
16	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Changed the procedure to change the IP address used for the mirroring among servers function.	"8.10 Changing IP Addresses Used for Mirroring among Servers"
17	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Added "Note" about the GDS commands that are not described in this command reference.	"Appendix B Command Reference"
18	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Added the error messages 42078 and 42079.	"C.3.2 Error Messages (42000-42099)"
19	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Added the error messages 60447 to 60449, 60452, and 60453.	"C.4.1 Error Messages (60000-60499)"
20	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Changed "Resolution" in "(4) System cannot be booted. (Failure of all boot disks)."	"D.1.5 System Disk Abnormality [EFI]"
21	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Deleted "(7) System disk mirroring failed to be set. [RHEL7]."	"D.1.5 System Disk Abnormality [EFI]"
22	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Changed "Resolution" in "(4) The slice becomes INVALID or the operation stops after restarting the other node when the network error occurs."	"D.1.6 Errors in the Mirroring Among Servers"
23	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Changed "Resolution" in "(5) The error message "ERROR: class: cannot operate shared objects, ..." is output, and the shared class <i>class</i> cannot be created; or the error message "ERROR: cluster communication failure" is output, and automatic registration of resources is failed."	"D.1.10 Cluster System Related Error"
24	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Changed "Explanation" in "(1) Messages are output at node startup."	"D.1.14 OS Messages"
25	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Changed "Explanation" in "(10) Messages of multipathd are output when creating volumes."	"D.1.14 OS Messages"
26	GDS	J2UL-2506-01ENZ0(00) J2UL-2506-01ENZ2(00)	Deleted "(12) Error messages of modprobe and systemd are output and GDS is not started when the system is started. [RHEL7]."	"D.1.14 OS Messages"

GDS: Global Disk Services

G.2.10 4.6A10 Manual Changes

No.	Product	Manual code	Description	Change/Section
1	GDS GDS I/O Monitor Option	J2UL-2506-02ENZ0(00) J2UL-2506-02ENZ2(00)	Changed the name of "FUJITSU Cloud Service for OSS" to "FUJITSU Hybrid IT Service FJcloud-O."	All chapters
2	GDS	J2UL-2506-02ENZ0(00) J2UL-2506-02ENZ2(00)	Changed the note when using Symfoware Server(Native).	"1.2.1.4 Mirroring among Servers"
3	GDS	J2UL-2506-02ENZ0(00) J2UL-2506-02ENZ2(00)	Changed the startup procedure of GDS Management View.	"1.3.1 Operation Management Interface" "4.9.2 Setting up the Client Environment" "5.1.1 Starting Web-Based Admin View Operation Menu" "5.7 Exiting GDS Management View"
4	GDS	J2UL-2506-02ENZ0(00) J2UL-2506-02ENZ2(00)	Deleted "4.6 kdump Settings."	"4.6 kdump Settings"
5	GDS	J2UL-2506-02ENZ0(00) J2UL-2506-02ENZ2(00)	Added the descriptions of the settings in an Azure environment or an environment where the by-id file does not exist.	"4.8.3 Creating iSCSI Target" "7.3.3.1 Hot Swap" "7.3.3.2 Cold Swap" "8.11 Adding Disks Used for Mirroring among Servers"
6	GDS	J2UL-2506-02ENZ0(00) J2UL-2506-02ENZ2(00)	Deleted "4.10.3 Setting up the Web Environment."	"4.10.3 Setting up the Web Environment"
7	GDS	J2UL-2506-02ENZ0(00) J2UL-2506-02ENZ2(00)	Changed the note when exiting the GDS Management View.	"5.7 Exiting GDS Management View"
8	GDS	J2UL-2506-02ENZ0(00) J2UL-2506-02ENZ2(00)	Added "5.8 Uninstalling Java Application."	"5.8 Uninstalling Java Application"
9	GDS	J2UL-2506-02ENZ0(00) J2UL-2506-02ENZ2(00)	Changed the description of the Excluded Device List.	"6.1.1 Excluded Device List"
10	GDS	J2UL-2506-02ENZ0(00) J2UL-2506-02ENZ2(00)	Added the description when creating a shared disk definition file.	"6.3.4 Shared Disk Unit Resource Registration Procedure"
11	GDS	J2UL-2506-02ENZ0(00) J2UL-2506-02ENZ2(00)	Changed the execution examples of the parted command.	"6.5.1 Checking Partition Configuration of System Disk" "9.5 OS Re-installation"
12	GDS	J2UL-2506-02ENZ0(00) J2UL-2506-02ENZ2(00)	Added the point when building the mirroring among servers in the class.	"6.6.1.3 Class Configuration"
13	GDS	J2UL-2506-02ENZ0(00) J2UL-2506-02ENZ2(00)	Added the procedure to restrict and to cancel restricting the automatic startup of the PRIMECLUSTER service in a cloud environment.	"6.11.5 Setting up Copy Destination System" "7.3.3.2 Cold Swap" "8.10 Changing IP Addresses Used for Mirroring among Servers"
14	GDS	J2UL-2506-02ENZ0(00) J2UL-2506-02ENZ2(00)	Added the note about the setting in a cloud environment.	"A.10 Changing system run level"
15	GDS	J2UL-2506-02ENZ0(00) J2UL-2506-02ENZ2(00)	Added "Information" to the description of the -e chkps (when using Restore) sub option.	"B.1.13 sdxconfig - Object Configuration Operations"

No.	Product	Manual code	Description	Change/Section
16	GDS	J2UL-2506-02ENZ0(00) J2UL-2506-02ENZ2(00)	Changed the description of the -e keepid (when using Remove) sub option.	"B.1.13 sdxconfig - Object Configuration Operations"
17	GDS	J2UL-2506-02ENZ0(00) J2UL-2506-02ENZ2(00)	Changed the resolution of "(4) System cannot be booted. (Failure of all boot disks)."	"D.1.5 System Disk Abnormality [EFI]"
18	GDS	J2UL-2506-02ENZ0(01) J2UL-2506-02ENZ2(01)	Added the descriptions of an FJcloud-Baremetal environment.	"3.2 DM-MP (Device Mapper Multipath) " "3.21.2 Cloud Environment" "4.5 Setting Up DM-MP"
19	GDS	J2UL-2506-02ENZ0(01) J2UL-2506-02ENZ2(01)	Changed "Explanation" and "Resolution" in "(1) The netmirror slice that configures the netmirror volume is in INVALID state."	"D.1.6 Errors in the Mirroring Among Servers"
20	GDS	J2UL-2506-02ENZ0(01) J2UL-2506-02ENZ2(01)	Added that an FJcloud-Baremetal environment is not supported in the descriptions of a cloud environment.	"1.2.1.4 Mirroring among Servers" "3.9.2 Disk" "3.21.1 KVM Environment and VMware Environment" "4.8.1 Setting Tuning Parameters" "6.11.5 Setting up Copy Destination System" "7.3.3.2 Cold Swap" "8.10 Changing IP Addresses Used for Mirroring among Servers" "A.10 Changing system run level"
21	GDS	J2UL-2506-02ENZ0(02) J2UL-2506-02ENZ2(02)	Added the descriptions of a NIFCLOUD environment.	"4.8.3 Creating iSCSI Target" "7.3.3.1 Hot Swap" "7.3.3.2 Cold Swap" "7.16.3 Operation when the LEFTCLUSTER Node Exists" "8.11 Adding Disks Used for Mirroring among Servers" "D.1.6 Errors in the Mirroring Among Servers"

GDS: Global Disk Services

GDS I/O Monitor Option: Global Disk Services I/O Monitor Option

G.2.11 4.6A20 Manual Changes

No.	Product	Manual code	Description	Change/Section
1	GDS	J2UL-2506-03ENZ0(00) J2UL-2506-03ENZ2(00)	Changed the description of the mirroring among servers.	"1.2.1.4 Mirroring among Servers" "3.9.4 Configuration and Settings of the Cluster System"
2	GDS	J2UL-2506-03ENZ0(00) J2UL-2506-03ENZ2(00)	Added the note about of the mirroring among servers.	"3.9.5 Data Amount and Data Update Amount"

No.	Product	Manual code	Description	Change/Section
3	GDS	J2UL-2506-03ENZ0(00) J2UL-2506-03ENZ2(00)	Added the note about of the Shared Disk Resource Registration for a single node.	"6.3.1 What Is Shared Disk Resource Registration?"
4	GDS	J2UL-2506-03ENZ0(00) J2UL-2506-03ENZ2(00)	Added the descriptions of the using multipath.	"A.5 Backing Up and Restoring Object Configuration (sdxconfig)"
5	GDS	J2UL-2506-03ENZ0(00) J2UL-2506-03ENZ2(00)	Added "Explanation" of the information messages 46015.	"C.3.4 Information Messages (46000-46199)"
6	GDS	J2UL-2506-03ENZ0(00) J2UL-2506-03ENZ2(00)	Added "Resolution" of the information messages 60346.	"C.4.1 Error Messages (60000-60499)"
7	GDS	J2UL-2506-03ENZ0(00) J2UL-2506-03ENZ2(00)	Added the note when mirroring the system disk.	"D.1.11 File System Error"
8	GDS	J2UL-2506-03ENZ0(00) J2UL-2506-03ENZ2(00)	Added "(13) A message is output from the standby node when the node is restarted, and the cluster cannot be started."	"D.1.14 OS Messages"
9	GDS	J2UL-2506-03ENZ0(00) J2UL-2506-03ENZ2(00)	Added "(14) When the stop operation of the GDS volume is executed, an error message is output, and the stop operation of the GDS volume fails."	"D.1.14 OS Messages"

GDS: Global Disk Services

G.3 Specification Changes Other Than New Features

This section describes specification changes of this software that are not new features.

G.3.1 Specification Changes in 4.3A00

Specification changes other than new features in the version 4.3A00 of this software are as follows.

G.3.1.1 Support for Export to NFS Clients

By defining file systems on volumes in the `/etc/exports` file, export of the file systems to an NFS client is achieved.

G.3.1.2 Change of I/O Amount Supported for Simultaneous Processing

The amount of I/O requests to volumes that can be accepted simultaneously is limited.

As a result, it may take time for making responses to the I/O requests if the processing issues many I/O requests at the same time.

This specification change is only for RHEL5 or later.

For RHEL4, no similar specification change was made and the I/O amount is unlimited.

G.3.2 Specification Changes in 4.3A10

Specification changes other than new features in the version 4.3A10 of this software are as follows.

G.3.2.1 Disk Label Type of a Disk with the Size smaller than 1 TB

When a GPT labeled disk with the size smaller than 1 TB is registered with the new local or shared class, the disk label of the disk and the class are indicated as below:

- For the environment where 1 TB or larger disks are unsupported: MSDOS label

- For the environment where 1 TB or larger disks are supported: GPT label

In local and shared classes of GPT label, some functions are unavailable. To use these functions, change the disk label to MSDOS type, and then register it with a class.



- For disk label types of disks registered with a class, see "[2.3.8 Disk Label](#)."
- For functions unavailable in local and shared classes of GPT label, see "[2.3.8 Disk Label](#)."

G.3.3 Specification Changes in 4.3A20

Specification changes other than new features in the version 4.3A20 of this software are as follows.

G.3.3.1 Disk Label Type of a Disk with the Size smaller than 1 TB

When a GPT labeled disk with the size smaller than 1 TB is registered with the new local or shared class, the disk label of the disk and the class are indicated as below:

- For the environment where 1 TB or larger disks are unsupported: MSDOS label
- For the environment where 1 TB or larger disks are supported: GPT label

In local and shared classes of GPT label, some functions are unavailable. To use these functions, change the disk label to MSDOS type, and then register it with a class.



- For disk label types of disks registered with a class, see "[2.3.8 Disk Label](#)."
- For functions unavailable in local and shared classes of GPT label, see "[2.3.8 Disk Label](#)."

G.3.4 Specification Changes in 4.3A30

Specification changes other than new features in the version 4.3A30 of this software are as follows.

G.3.4.1 Restoring Object Configuration (sdxcfg Restore Command)

The performance is changed when a physical disk to be registered with a class by the sdxcfg Restore command met one of the following conditions.

- When a disk label of the physical disk is not created.
- When a disk label type of the physical disk is different from a disk label type of a class which is described in the configuration file.

When one of the above conditions is met, the performance for restoring an object configuration of a class by the sdxcfg Restore command varies by the GDS version and applied patches.

- For the patch for PG97489 is not applied to the GDS version 4.3A20 or earlier,
or the patch for PG97489 is not applied to the GDS version 4.3A30 (RHEL5)

An object configuration of a class can be restored by the sdxcfg Restore command.

- For the patch for PG97489 is applied to the GDS version 4.3A20 or earlier,
or the patch for PG97489 is applied to the GDS version 4.3A30 (RHEL5),
or the GDS version 4.3A30 (RHEL6),
or the GDS version 4.3A40 or later

The sdxcfg Restore command will fail and the object configuration of the class cannot be restored. At this time, the error message 60380 will be output. For the error message 60380, see "[C.4.1 Error Messages \(60000-60499\)](#)."

By this change, you need to note for restoring an object configuration of a class in the system which uses the different GDS versions. For details, see "[A.5 Backing Up and Restoring Object Configuration \(sdxconfig\)](#)."

G.3.4.2 File System Configuration screen of GDS Management View

The File System Configuration screen of GDS Management View is stopped providing.

G.3.4.3 Changes of the warning message "5119" in the Volume Configuration screen

The warning message "5119" in the Volume Configuration screen is changed because the File System Configuration screen is no longer supported.

- GDS 4.3A20 or earlier

5119 A file system has been configured on the selected volume. Remove the file system from the File System Configuration.

- GDS 4.3A30 or later

5119 A file system has been configured on the selected volume. Remove the file system.

G.3.5 Specification Changes in 4.3A40

Except for newly added functions, the specification of this software remains unchanged in version 4.3A40.

G.3.6 Specification Changes in 4.4A00

Except for newly added functions, the specification of this software remains unchanged in version 4.4A00.

G.3.7 Specification Changes in 4.5A00

Except for newly added functions, the specification of this software remains unchanged in version 4.5A00.

G.3.8 Specification Changes in 4.5A10

Except for newly added functions, the specification of this software remains unchanged in version 4.5A10.

G.3.9 Specification Changes in 4.6A00

Specification changes other than new features in the version 4.6A00 of this software are as follows.

G.3.9.1 Operation after Restoration from LEFTCLUSTER in Mirroring among Servers Configuration

A software specification is improved to continue I/O to the netmirror volume even when a node becomes the LEFTCLUSTER state due to a failure of servers or interconnects.

According to this improvement, the synchronization copying and volume startup after restoration from the LEFTCLUSTER state may be suppressed due to data protection. For details, see "[7.16.4 Restoration from the LEFTCLUSTER State](#)."

G.3.9.2 Unsupported Functions in RHEL8

The following functions of the disk group are not supported in RHEL8.

- Striping
- Concatenation

G.3.10 Specification Changes in 4.6A10

Except for newly added functions, the specification of this software remains unchanged in version 4.6A10.

G.3.11 Specification Changes in 4.6A20

Except for newly added functions, the specification of this software remains unchanged in version 4.6A20.

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 ETERNUS Disk storage system. A generic term for EC (Equivalent Copy) and OPC (One Point Copy). For details, see the ETERNUS Disk storage system manual.

alternative root volume [EFI]

The alternative volume of the root volume.

alternative volume [EFI]

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.

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.

boot disk

A bootable disk with /boot/efi installed.

break

To cancel the master-proxy relationship, and make independent again.

by-id file

One of the disk device identifiers managed by the udev function of OS. A file name of the by-id file is not changed even if a device name is changed. It is the symbolic link file corresponded to the device file in the /dev directory, and exists in the /dev/disk/by-id directory.

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 [EFI]

The volume currently running as the root file system.

current volume [EFI]

The master volume or the proxy volume that is described as a file system or a swap area in the /etc/fstab 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.

device name change

For Linux specification, device names are changed by changing recognized order of devices in starting the system. Device name changes of disks may occur when restarting the system after occurring a disk error or adding and removing a disk.

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.

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 ETERNUS Disk storage system. 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 ETERNUS Disk storage system manual.

EC session

EC request from the host (server) to ETERNUS Disk storage system.

For details, see the ETERNUS Disk storage system 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.

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 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."

Global File Services

A software that manages shared file system.

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 ETERNUS Disk storage system. 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."

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 [EFI]

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 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, 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.

mirroring among servers

To mirror the local disks (such as an internal disk) of more than one server via the network.

mpath device

A virtual device that provides multiple access paths with the device-mapper driver which is the standard OS multipath function of RHEL. 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.

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.

netmirror group

A group whose type is "netmirror." It is the set of disks, which will be mirrored among servers.

netmirror slice

A logical slice that composes a netmirror volume. A netmirror volume is created by mirroring more than one netmirror slice. For each netmirror volume, there will be one logical slice in each disk belonging to the netmirror group.

netmirror volume

A volume that is created within a netmirror group. Its data will be made redundant by mirroring among servers.

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 ETERNUS Disk storage system. 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 ETERNUS Disk storage system manual.

OPC session

OPC request from the host (server) to the ETERNUS Disk storage system.

For details, see the ETERNUS Disk storage system 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. A disk partition. The starting sector number and the number of sectors occupied by the physical slice is registered with the disk label (partition table).

physical special file

A device special file of a physical disk, such as `/dev/sd[a-z]* [1-4]*`.

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.

QuickOPC

Copy processing by QuickOPC function.

QuickOPC function

One of the OPC functions of ETERNUS Disk storage system. All data is copied at the first copy processing and only updated data (difference) is copied from the second time.

For details, see the manual for ETERNUS Disk storage system.

QuickOPC session

QuickOPC request that is from a host (server) to an ETERNUS Disk storage system.

For details, see the manual for ETERNUS Disk storage system.

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 ETERNUS Disk storage system. REC is EC function available between ETERNUS Disk storage system boxes. For details, see the ETERNUS Disk storage system 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, shared disk units 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.

resource registration

The function that registers hardware units with the resource database in a PRIMECLUSTER system. This function is also referred to as automatic resource registration. To perform resource registration, use the PRIMECLUSTER CF clautoconfig command. For details, see "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide" and "[6.3 Shared Disk Resource Registration](#)."

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 [EFI]

A "root" type class. Objects that belong to a root class can be used only on the current node. With the root class, system disks connected to the current node, disks to mirror with the system disk, spare disks, and disks on which snapshots of the system disks are created, can be registered.

For the environment conditions where a root class can be used, see "[3.4 Managing System Disks](#)."

root object [EFI]

An object that belongs to a "root" class.

root volume [EFI]

The volume mounted as the root file system.

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 disk definition file

A file that declares physical connection configuration of nodes and shared disk units in a PRIMECLUSTER system. Specify the full path to a shared disk definition file with the -f option when executing the clautoconfig(8) command, and disk unit information in the shared disk definition file is registered with the PRIMECLUSTER resource database. For details, see "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide" and "[6.3 Shared Disk Resource Registration](#)."

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.
-

slice number

A partition number.

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

Dell 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 Dell 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 the Linux operating system 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, /boot, /boot/efi, or a swap area

system volume [EFI]

The volume currently running as one of the following file systems (or a swap area).

/, /usr, /var, /boot, /boot/efi, 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

Dell EMC's software. Using Symmetrix hardware functions, disk data is copied within a Symmetrix case. For details see the TimeFinder manual.

tracking

One of the statuses of QuickOPC session. The status while updated parts are being recorded.

For details, see the manual for ETERNUS Disk storage system.

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.

Web-Based Admin View

Software for system operation and management using the Web screen.

For details, see "PRIMECLUSTER Web-Based Admin View Operation Guide."