



PRIMECLUSTER™

Global Disk Services

Configuration and Administration Guide 4.2

(Linux)

June 2006

Preface

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

Purpose

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

Target Reader

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

Organization

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

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

Headers	Contents	Operator	Administrator
Chapter 1 Functions	Explains the features of GDS and GDS Snapshot.	YES	YES
Chapter 2 Objects	Explains the objects used by GDS and GDS Snapshot.	YES	YES
Chapter 3 Starting and Exiting GDS Management View	Explains how to start and exit GDS Management and the required browser environment.	YES	YES
Chapter 4 Management View Screen Elements	Explains the contents of Management View.	YES	YES
Chapter 5 Operation	Explains the details about GDS operation such as operation flow, settings, maintenance and management.	YES	YES
Chapter 6 Backing Up and Restoring	Explains how to back up and restore data on disks managed by GDS.	YES	YES
Appendix A General Notes	Explains guidelines, general precautions, and configuration tips necessary for using GDS.	-	YES
Appendix B Log Viewing with Web-Based Admin View	For details, see the supplementary "Web-Based Admin View Operation Guide."	-	YES
Appendix C Web-Based Admin View Operating Environment Setting	For details, see the supplementary "Web-Based Admin View Operation Guide."	YES	YES
Appendix D Command Reference	Explains the commands available in GDS and GDS Snapshot.	-	YES
Appendix E GDS Messages	Explains the contents, possible causes, and resolutions for GDS messages that appear when setting or operating GDS and GDS Snapshot.	YES	YES
Appendix F troubleshooting	Explains resolutions for abnormality of objects and physical disks managed by GDS.	-	YES
Appendix G Frequently Asked Questions (FAQ)	Lists frequently asked questions regarding GDS and GDS Snapshot.	YES	YES
Appendix H Shared Disk Unit Resource Registration	Explains shared disk unit resource registration to be performed first to use GDS in a cluster system.	YES	YES
Glossary	Explains GDS and GDS Snapshot terminology. Please refer to it as necessary.	YES	YES

Related documentation

Please refer to the following documents according to need.

- PRIMECLUSTER Concepts Guide
- PRIMECLUSTER Cluster Foundation (CF) (Linux) 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 (Linux)

Manual Printing

Use the PDF file supplied in the product CD-ROM to print this manual.

Adobe Acrobat Reader Version 4.0 or higher is required to read and print this PDF file.

Symbol

The following conventions are used in this manual:

- [Itanium] indicates that the description is for the Itanium Processor Family (IPF) server platform.
- [Linux2.6] indicates that the description is for Linux kernel version 2.6.
- [Linux2.4] indicates that the description is for Linux kernel version 2.4.



- **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

- Itanium Processor Family is abbreviated as IPF.

June 2006

Trademarks

Netscape and Netscape Navigator are trademarks of Netscape Communications Corporation.

Microsoft, Windows, Windows NT, Windows Me, Internet Explorer, and Windows 2000 are registered trademarks of Microsoft Corporation in the United States and other countries.

Linux is a registered trademark or a trademark of Linus Torvalds in the United States and other countries.

All Red Hat, RPM, and Red Hat based trademarks and logos are registered trademarks of Red Hat, Inc. in the United States and other countries.

Java and all Java based-marks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

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

SAN Manager is a trademark of EMC Corporation.

PRIMECLUSTER is a trademark of Fujitsu LIMITED.

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

Notice

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

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

All Rights Reserved, Copyright(C) 2006 Fujitsu Siemens Computers GmbH.

Contents

Chapter 1	Function	1
1.1	GDS Features	2
1.2	Functions for High Availability	4
1.2.1	Disk Mirroring	4
1.2.1.1	System Disk Mirroring [Itanium]	5
1.2.1.2	Mirroring between Disk Array Unit	6
1.2.1.3	Shared Disk Mirroring	7
1.2.2	Hot Spare	8
1.2.3	Hot Swap	11
1.2.4	Just Resynchronization Mechanism (JRM)	12
1.3	Functions for High Manageability	13
1.3.1	Operation Management Interface	13
1.3.2	Centralized Disk Management	14
1.3.3	Name Management	15
1.3.4	Single System Image Environment	16
1.3.5	Access Control	17
1.3.6	Realizing Large Capacity and I/O Load Balancing	19
1.3.6.1	Logical Partitioning	20
1.3.6.2	Disk Concatenation	20
1.3.6.3	Disk Striping	21
1.3.6.4	Combining Disk Striping with Mirroring	22
1.3.7	Online Volume Expansion	23
1.3.8	Snapshots by Slice Detachment	25
1.4	GDS Snapshot Features	27
1.5	Proxy Volume	28
1.5.1	Snapshot by Synchronization	30
1.5.2	Snapshot Function with No Load to the Server/SAN	34
1.5.3	Instant Snapshot by OPC	35
1.5.4	Instant Restore	36
1.5.5	Online Disk Migration	38
1.5.6	Creating an Alternative Boot Environment [Itanium]	39
1.6	Shadow Volume	40
1.6.1	Accessing Volumes from an External Server	40
1.6.2	Using the Copy Function of a Disk Unit	41
Chapter 2	Objects	43
2.1	SDX Object	44
2.2	Disk Class	45
2.3	SDX Disk	48
2.4	Disk Group	50
2.5	Logical Volume	52
2.6	Logical Slice	57
2.7	GDS Snapshot Objects	60
2.7.1	Proxy Object	60
2.7.2	Shadow Object	63
2.7.2.1	Shadow Class	63
2.7.2.2	Shadow Disk	66
2.7.2.3	Shadow Group	69
2.7.2.4	Shadow Volume	70
2.7.2.5	Shadow Slice	75
Chapter 3	Starting and Exiting GDS Management View	77
3.1	Preparation for Starting GDS Management View	78
3.1.1	Deciding the User Group	78
3.1.1.1	User Group Types	78
3.1.1.2	Creating User Groups	78
3.1.1.3	Registering to a User Group	79
3.1.2	Setting up the Client Environment	79
3.1.3	Setting up the Web Environment	79
3.2	Starting the GDS Management View	80
3.2.1	Starting Web-Based Admin View Operation Menu	80
3.2.2	Web-Based Admin View Operation Menu Functions	81
3.2.3	Starting GDS Management View	82
3.3	Exiting GDS Management View	83
3.4	Changing the Web-Based Admin View Settings	84

Chapter 4	Management View Screen Elements	85
4.1	Screen Configuration	86
4.2	Menu Configuration and Functions	89
4.2.1	General	89
4.2.2	Settings	90
4.2.3	Operation	96
4.2.4	View	98
4.2.5	Help	99
4.3	Icon Types and Object Status	100
4.4	Object Information	104
Chapter 5	Operation	105
5.1	Operation Outline	106
5.1.1	System Disk Settings [Itanium]	106
5.1.2	Configuration Settings	106
5.1.2.1	Single Volume Configuration Settings	107
5.1.2.2	Other Volume Configuration Settings	108
5.1.3	Backup	109
5.1.3.1	Backup (by Slice Detachment)	109
5.1.3.2	Backup (by Synchronization)	110
5.1.3.3	Backup (by OPC)	111
5.1.4	Restore	112
5.1.5	Disk Swap	113
5.1.6	Disk Migration	114
5.1.7	Configuration Change	115
5.1.8	Unmirroring the System Disk [Itanium]	116
5.1.9	Operations from GDS Management View	117
5.2	Settings	121
5.2.1	System Disk Settings [Itanium]	121
5.2.2	Operating from the Settings Menu	128
5.2.2.1	Class Configuration	128
5.2.2.2	Cluster System Class Configuration	132
5.2.2.3	Group Configuration	134
5.2.2.4	Volume Configuration	137
5.2.3	File System Configuration	140
5.2.4	Proxy Configuration	143
5.2.4.1	Join	143
5.2.4.2	Relate	148
5.3	Operation in Use	151
5.3.1	Viewing Configurations and Monitoring Statuses	151
5.3.1.1	Confirming SDX Object Configuration	151
5.3.1.2	Viewing Proxy Object Configurations	157
5.3.1.3	Monitoring Object Status	162
5.3.2	Backup	164
5.3.2.1	Backup (by Slice Detachment)	164
5.3.2.2	Backup (by Synchronization)	171
5.3.2.3	Backup (by OPC)	176
5.3.3	Restore	179
5.3.4	Disk Swap	184
5.3.5	Disk Migration	187
5.3.6	Copying Operation	190
5.4	Changes	194
5.4.1	Class Configuration	194
5.4.2	Group Configuration	199
5.4.3	Volume Configuration	202
5.5	Removals	204
5.5.1	Removing a File System	204
5.5.2	Removing a Volume	205
5.5.3	Removing a Group	207
5.5.4	Removing a Class	208
5.5.5	Unmirroring the System Disk [Itanium]	209
5.5.6	Breaking a Proxy	211
Chapter 6	Backing Up and Restoring	213
6.1	Backing Up and Restoring a System Disk [Itanium]	214
6.1.1	Checking Physical Disk Information and Slice Numbers	214
6.1.2	Backing Up	217
6.1.3	Restoring (When the System Can Be Booted)	220
6.1.4	Restoring (When the System Cannot Be Booted)	223

6.2	Backing Up and Restoring a System Disk through an Alternative Boot Environment [Itanium].....	225
6.2.1	System Configuration	225
6.2.2	Summary of Backup	227
6.2.3	Summary of Restore.....	228
6.2.4	Summary of Procedure.....	229
6.2.5	Configuring an Environment	230
6.2.6	Backing Up	232
6.2.7	Restoring	236
6.3	Backing Up and Restoring Local Disks and Shared Disks.....	240
6.3.1	Offline Backup	240
6.3.2	Online Backup (by Slice Detachment)	242
6.3.3	Restoring	244
6.4	Online Backup and Instant Restore through Proxy Volume	245
6.4.1	Online Backup (by Synchronization).....	247
6.4.2	Online Backup (Snapshot by OPC)	252
6.4.3	Instant Restore	256
6.5	Backing Up and Restoring through Disk Unit's Copy Functions [Linux2.6]	260
6.5.1	Configuring an Environment	260
6.5.2	Backing Up	262
6.5.3	Restoring from Backup Disks.....	265
6.6	Backing Up and Restoring through an External Server.....	267
6.6.1	Backing Up and Restoring a Logical Volume with No Replication	269
6.6.1.1	System Configuration	269
6.6.1.2	Summary of Backup	270
6.6.1.3	Summary of Restore.....	271
6.6.1.4	Summary of Procedure.....	272
6.6.1.5	Configuring an Environment	273
6.6.1.6	Backing Up	274
6.6.1.7	Restoring	278
6.6.2	Backing Up and Restoring through Snapshot by Slice Detachment	282
6.6.2.1	System Configuration	282
6.6.2.2	Summary of Backup	283
6.6.2.3	Summary of Restore.....	285
6.6.2.4	Summary of Procedure.....	287
6.6.2.5	Configuring an Environment	288
6.6.2.6	Backing Up	289
6.6.2.7	Restoring	294
6.6.3	Backing Up and Restoring Using Snapshots from a Proxy Volume	300
6.6.3.1	System Configuration	300
6.6.3.2	Summary of Backup	301
6.6.3.3	Summary of Restore from a Proxy Volume.....	303
6.6.3.4	Summary of Restore from Tape.....	304
6.6.3.5	Summary of Procedure.....	306
6.6.3.6	Configuring an Environment	308
6.6.3.7	Backing Up	309
6.6.3.8	Restoring from a Proxy Volume	313
6.6.3.9	Restoring from Tape	314
6.6.4	Backing Up and Restoring by the Disk Unit's Copy Function.....	318
6.6.4.1	System Configuration	319
6.6.4.2	Summary of Backup	320
6.6.4.3	Summary of Restore from a BCV	322
6.6.4.4	Summary of Restore from Tape	324
6.6.4.5	Summary of Procedure.....	326
6.6.4.6	Configuring an Environment	328
6.6.4.7	Backing Up	329
6.6.4.8	Restoring form a BCV.....	334
6.6.4.9	Restoring from Tape	336
6.7	Backing Up and Restoring Object Configurations [Linux2.6]	342
6.7.1	Backing Up	342
6.7.2	Restoring	343

Appendix A	General Notes	347
A.1	Rules.....	347
A.1.1	Object Name.....	347
A.1.2	Number of Classes	348
A.1.3	Number of Disks	349
A.1.4	Number of Groups	349
A.1.5	Number of Volumes	349
A.1.6	Number of Keep Disks [Itanium]	350
A.1.7	Creating Group Hierarchy.....	350
A.1.8	Exception to Proxy Configuration.....	351
A.1.9	Number of Proxy Volumes	352
A.1.10	Proxy Volume Size	352
A.1.11	Proxy Group Size.....	352
A.2	Important Points.....	353
A.2.1	Managing System Disks	353
A.2.2	Restraining Access to Physical Special File.....	353
A.2.3	Bootng from a CD-ROM Device.....	354
A.2.4	Initializing Disk	354
A.2.5	Disk Size.....	355
A.2.6	Volume Size.....	356
A.2.7	Hot Spare.....	356
A.2.8	System Disk Mirroring [Itanium]	358
A.2.9	Keep Disk [Itanium].....	358
A.2.10	Creating a Snapshot by Slice Detachment.....	358
A.2.11	The Difference between a Mirror Slice and a Proxy Volume.....	359
A.2.12	Just Resynchronization Mechanism (JRM)	359
A.2.13	Online Volume Expansion	361
A.2.14	Swapping Physical Disks	362
A.2.15	Object Operation When Using Proxy	363
A.2.16	Using the Advanced Copy Function in a Proxy Configuration	364
A.2.17	Instant Snapshot by OPC.....	365
A.2.18	To Use EMC Symmetrix	366
A.2.19	Using EMC TimeFinder or EMC SRDF in a Proxy Configuration [Linux2.6].....	368
A.2.20	Ensuring Consistency of Snapshot Data.....	370
A.2.21	Data Consistency at the time of Simultaneous Access.....	370
A.2.22	Volume Access Mode.....	370
A.2.23	Operation in Cluster System	371
A.2.24	Changing Over from Single Nodes to a Cluster System [Linux2.6]	371
A.2.25	Disk Switch.....	372
A.2.26	Shadow Volume	373
A.2.27	Configuration Tables [Linux2.6].....	375
A.2.28	GDS Management View.....	376
A.2.29	File System Auto Mount	376
A.2.30	Raw Device Binding	377
A.2.31	Upgrading Instructions	377
A.3	General Points	380
A.3.1	Guidelines for Mirroring	380
A.3.2	Guidelines for Striping.....	380
A.3.3	Guidelines for Concatenation.....	381
A.3.4	Guidelines for Combining Striping with Mirroring	381
Appendix B	Log Viewing with Web-Based Admin View	383
Appendix C	Web-Based Admin View Operating Environment Setting	385
Appendix D	Command Reference	387
D.1	sdxcass - Class operations	388
D.2	sdxdisk - Disk operations	389
D.3	sdxgroup - Group operations	399
D.4	sdxvolume - Volume operations.....	403
D.5	sdxslice - Slice operations	408
D.6	sdxinfo - Display object configuration and status information	411
D.7	sdxattr - Set objects attributes.....	422
D.8	sdxswap - Swap disk	428
D.9	sdxfix - Restore a failed object.....	430
D.10	sdxcopy - Synchronization copying operation	432
D.11	sdxroot - Root file system mirroring definition and cancellation [Itanium]	434
D.12	sdxparam - Configuration parameter operations	437
D.13	sdxconfig - Object configuration operations [Linux2.6]	439
D.14	sdxproxy - Proxy object operations	443

D.15	sdxshadowdisk - Shadow disk operations	453
D.16	sdxshadowgroup - Shadow group operations	458
D.17	sdxshadowvolume - Shadow volume operations	461
D.18	Volume Creation Using Command	464
D.19	Snapshot Creation Using Command	468
D.20	Volume Expansion Using Commands [Itanium]	470
Appendix E	GDS Messages	477
E.1	Web-Based Admin View Messages (0001-0099)	477
E.2	Driver Messages	477
E.2.1	Warning Messages (22000-22099)	478
E.2.2	Information Messages (24000-24099)	481
E.2.3	Internal Error Messages (26000)	484
E.3	Daemon Messages	485
E.3.1	Halt Messages (40000-40099)	487
E.3.2	Error Messages (42000-42099)	488
E.3.3	Warning Messages (44000-44099)	498
E.3.4	Information Messages (46000-46199)	505
E.3.5	Internal Error Messages (48000)	520
E.4	Command Messages	521
E.4.1	Error Messages (60000-60399)	522
E.4.2	Warning Messages (62000-62099)	578
E.4.3	Information Messages (64000-64099)	582
E.4.4	Fix Messages (66000)	592
E.4.5	Internal Error Messages (68000)	593
E.5	Operation Management View Messages	594
E.5.1	Error Messages (5000-5099)	594
E.5.2	Warning Messages (5000,5100-5199)	597
E.5.3	Information Messages (5200-5299)	605
Appendix F	Troubleshooting	613
F.1	Resolving Problems	613
F.1.1	Slice Status Abnormality	613
(1)	Mirror slice configuring the mirror volume is in INVALID status.	613
(2)	The copy destination slice was made INVALID due to an I/O error generated on the copy source slice during synchronization copying.	617
(3)	Slice configuring the volume is in TEMP status.	620
(4)	Slice configuring volume is in TEMP-STOP status.	620
(5)	Slice configuring the volume is in COPY status.	620
(6)	Slice configuring the volume is in NOUSE status.	620
F.1.2	Disk Status Abnormality	621
(1)	Disk is in DISABLE status.	621
(2)	Disk is in SWAP status.	621
(3)	Disk is in I/O error status.	622
F.1.3	Volume Status Abnormality	622
(1)	Mirror volume is in INVALID status.	622
(2)	Single volume is in INVALID status.	624
(3)	Stripe volume or volume in concatenation group is in INVALID status.	625
(4)	Master volume is in INVALID status.	626
(5)	Proxy volume is in INVALID status.	628
(6)	Volume is in STOP status.	630
(7)	I/O error occurs although mirror volume is in ACTIVE status.	631
(8)	An I/O error occurs on a single volume.	633
(9)	An I/O error occurs on a stripe volume or a volume in a concatenation group.	635
F.1.4	Class Status Abnormality	638
(1)	Class becomes closed status during operation.	638
F.1.5	System Disk Abnormality [Itanium]	640
(1)	Create backup of root file system.	640
(2)	System can be booted, but the system disk data is invalid.	641
(3)	System cannot be booted. (Failure of original boot disk)	641
(4)	System cannot be booted. (Boot disk data damage)	642
F.1.6	GDS Management View Abnormality	642
(1)	Web-Based Admin View abnormality	642
(2)	GDS Management View cannot be started.	643
(3)	Disk information is not displayed on GDS Management View. Or, the disk sizes for disk units are not displayed correctly.	643
F.1.7	Proxy Object Abnormality	644
(1)	The Advanced Copy function cannot be used in master-proxy copying.	644

F.1.8	EMC Symmetrix Abnormality [Linux2.6].....	646
	(1) SRDF is not used in master-to- proxy synchronization copying.....	646
	(2) The joined proxy volume is in INVALID status.....	650
F.1.9	Cluster System Related Error.....	650
	(1) The error message "ERROR: class: cannot operate in cluster environment, ..." is output, and the operation cannot be conducted on the class class. [Linux2.6].....	651
	(2) The PRIMECLUSTER CF clinitreset(8) command ends abnormally outputting an error message # 6675.	653
	(3) Cluster applications become "Inconsistent".	654
	(4) The GFS Shared File System is not mounted on node startup.	654
F.1.10	Shadow Object Errors.....	655
	(1) A shadow volume became inaccessible.	655
F.1.11	Disk Unit Error.....	655
F.2	Collecting Investigation Material.....	656
F.2.1	Collecting with pclsnap Command.....	656
F.2.2	Collecting Initial Investigation Material with sdxsnap.sh Script.....	656
Appendix G	Frequently Asked Questions (FAQ).....	657
G.1	Operation Design.....	657
	Q: Is operation possible without using GUI (GDS Management View)?.....	657
	Q: Is it possible to perform GDS Snapshot operations through the GUI (GDS Management View)?.....	657
	Q: Can disks with different performance specifications, such as size and revolution speed, be mirrored?.....	657
	Q: Do disks have to have equivalent attributes, such as size and revolution speed when I relate them as master and proxy?.....	658
	Q: Can I combine mirroring and striping?.....	658
	Q: Do I need to specify certain settings in order to perform a hot swap?.....	658
	Q: What is the purpose of single volume?.....	658
	Q: Can the messages output by GDS driver, daemon and command be displayed in a language other than English?.....	658
	Q: Are there any restrictions on the LUN (Logical Unit) configuration within the disk array when managing the disk array using GDS?.....	658
	Q: Are the mirrored disks differentiated as "primary" and "secondary" disks?.....	658
	Q: When reading the mirror volume, which one of the slices in the volume will the I/O request be issued to?.....	659
	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?.....	659
	Q: Which of the GDS messages should be monitored?.....	659
	Q: I want to know GDS resident daemons to monitor resident processes.....	660
	Q: When configuring a 24-hour operation system, would files of incremental size such as message log file have any effect?.....	660
G.2	Environment Configuration.....	661
	Q: How do I make a backup of a system disk that is managed with GDS? [Itanium].....	661
	Q: How do I make a backup of a disk (other than a system disk) that is managed with GDS?.....	661
	Q: I cannot execute the command parted(8) on a physical disk managed by GDS (e.g. sda). How can I solve this problem?.....	661
G.3	Operation.....	662
	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?.....	662
	Q: How do I change the volume size?.....	663
	Q: I cannot boot the system. How do I restore it? [Itanium].....	663
	Q: When parting a proxy volume in a shared class, I want to lock the proxy volume not to start up on certain nodes.....	663
	Q: Is there a possibility that joined proxy volumes are parted automatically?.....	663
Appendix H	Shared Disk Unit Resource Registration.....	665
H.1	What Is Shared Disk Resource Registration?.....	665
H.2	The Flow of Shared Disk Resource Registration.....	665
H.3	Preconditions and Important Points.....	666
H.4	Shared Disk Unit Resource Registration Procedure.....	667
H.5	Command Reference.....	670
	H.5.1 clautoconfig(8)-Execute resource registration.....	670
	H.5.2 cldelrsc(8)-Delete resource.....	671
Glossary	673

Chapter 1 Function

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

1.1 GDS Features

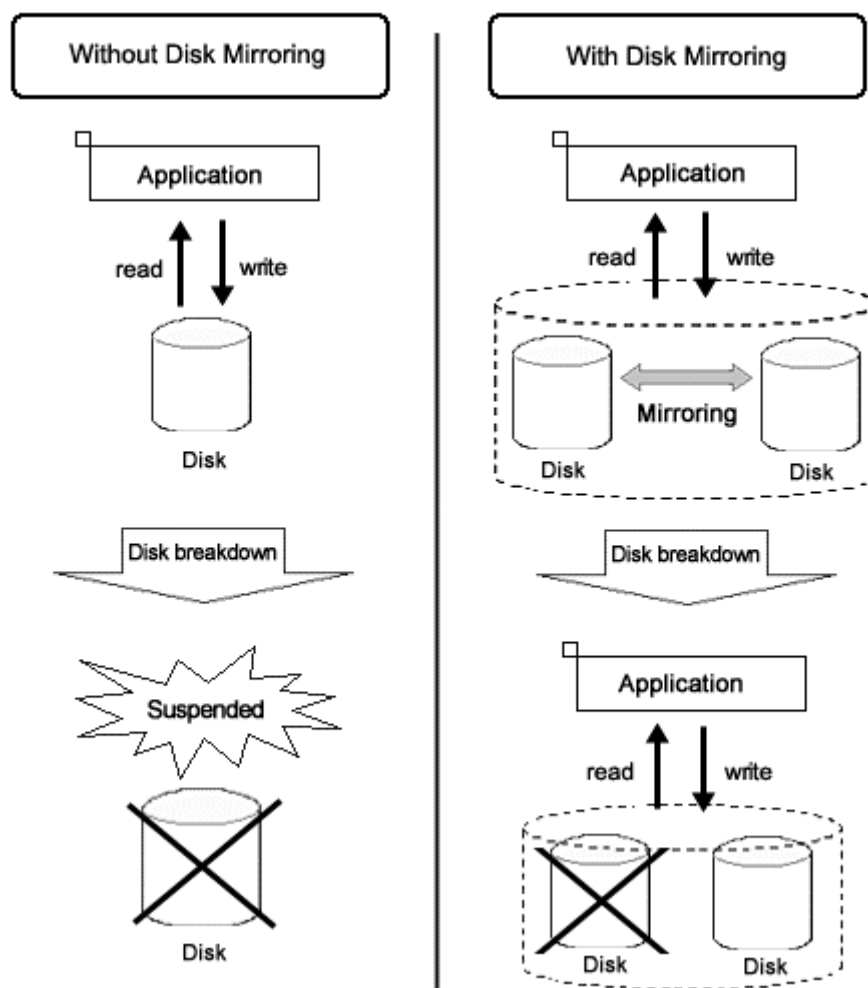
GDS is volume management software that improves the availability and manageability of disk-stored data. GDS protects disk data from hardware failures and operational mistakes and supports the management of disk units.

GDS has the following functions, which are closely related:

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

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

Figure: 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 integrate management of disk units that are shared by multiple servers in a SAN (Storage Area Network) environment as well as local disk units that are connected to specific servers.



Note Locally Connected Disk Unit Management

If the Linux kernel version is 2.4, locally connected disk units cannot be managed.

Figure: SAN (Storage Area Network)




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

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

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

Figure: Access Control Mechanism in a SAN Environment



1.2 Functions for High Availability

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


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

1.2.1 Disk Mirroring

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

GDS realizes disk mirroring by creating one logical disk from multiple physical disks. Disk mirroring allows applications to access a different disk even if one mirror disk (a disk that is mirrored) fails. Therefore, data will not be damaged and the applications will continue normal operation.

Figure: Disk Mirroring



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

Some examples of special mirroring implementation are explained below.

1.2.1.1 System Disk Mirroring [Itanium]

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

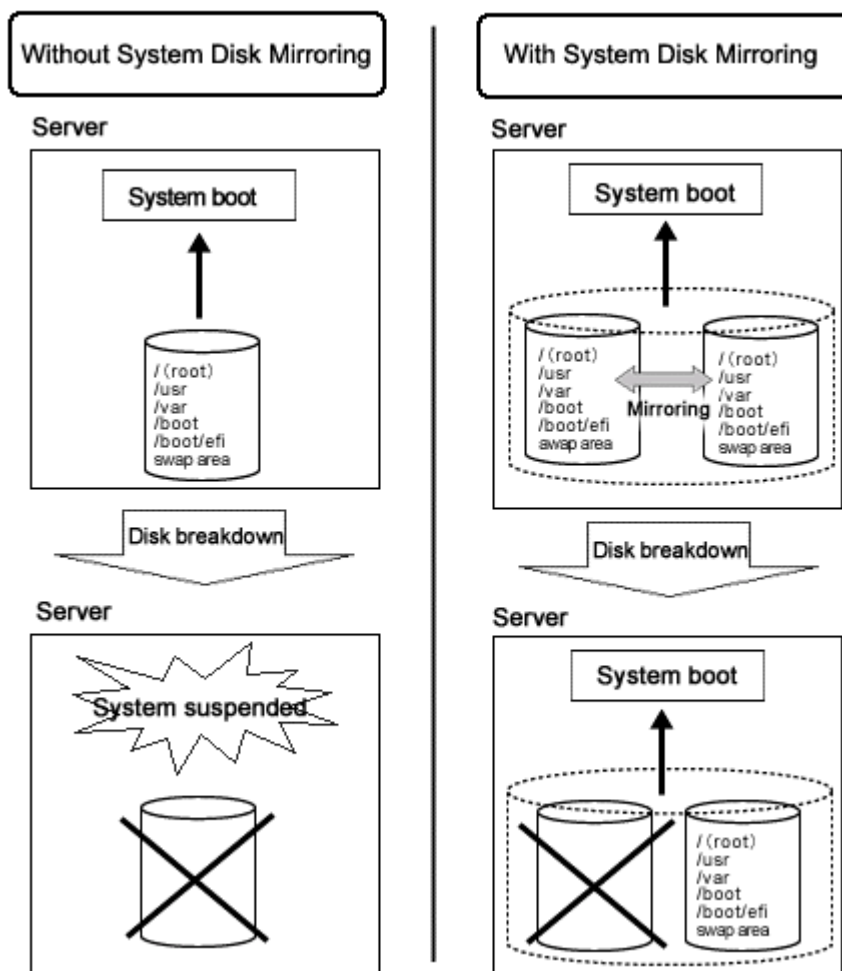
If a system disk failure occurs, the entire system probably stops and even booting will become impossible, and operating system reinstallation will be required. As a result, service stop may remain for a long time. System disk mirroring ensures continuous system operation even when a failure occurs in part of the disks.



Note SAN Boot

In a SAN boot environment, system disk mirroring is not supported.

Figure: System Disk Mirroring




1.2.1.2 Mirroring between Disk Array Unit

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

Mirroring two disk arrays connected with Fibre Channel provides data protection from unexpected accidents and blackouts. Also, by combining this feature with special software, disk units with redundant access paths can also be mirrored.

Figure: Mirroring between Disk Array Unit



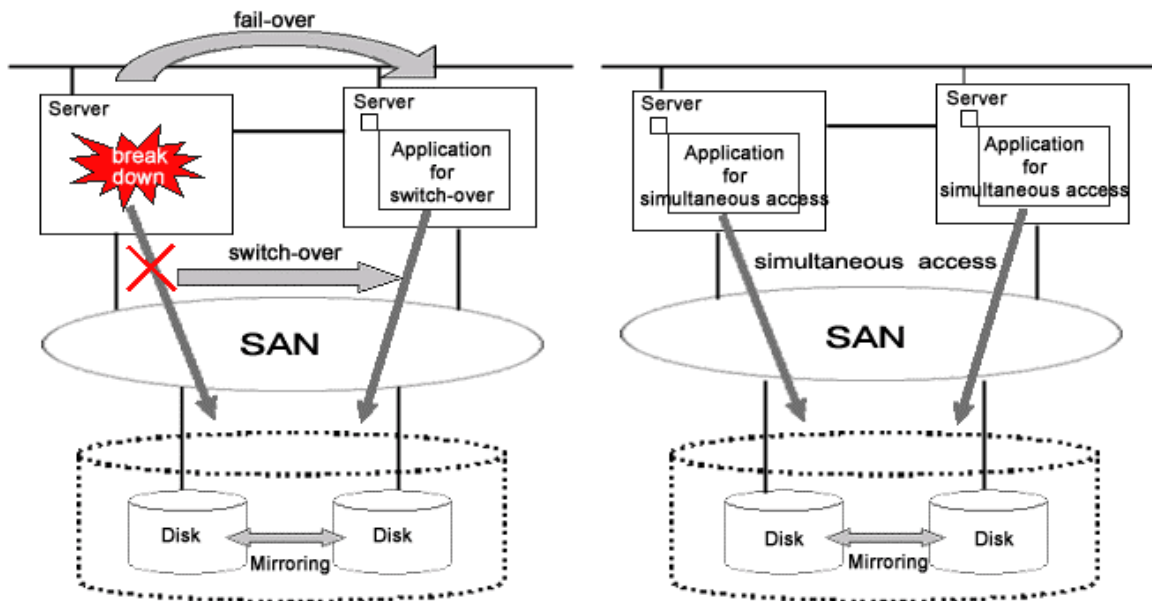
1.2.1.3 Shared Disk Mirroring

GDS provides a function that mirrors shared disk units connected a cluster system composed of multiple servers (also referred to as nodes).

This is called shared disk mirroring in distinction from local disk mirroring that mirrors disks connected to a single node.

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

Figure: Shared Disk Mirroring




1.2.2 Hot Spare

Overview

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

Figure: Hot Spare



Spare Disk Automatic Connection

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

Spare Disk Automatic Disconnection

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

Hot Spare Mode (Spare Disk Selection Mode)

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

- **External Mode (Default)**

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


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

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

Features:

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

Figure: 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: Hot Spare in Internal Mode



Note **Points of Concern for Hot Spare**
See "A.2.7 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.



Note Conditions for Hot Swap

GDS's function for hot swapping is available only when using a disk unit supporting hot swapping. Do not remove or insert a disk when using a disk unit not supporting hot swapping. It may cause breakdown, or damage data.

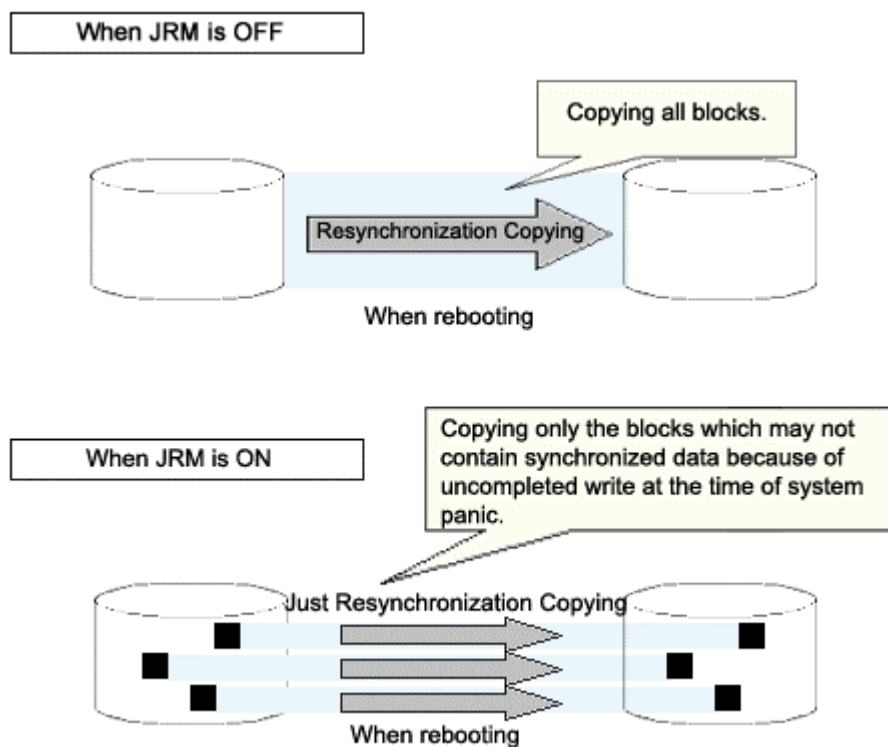
1.2.4 Just Resynchronization Mechanism (JRM)

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

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

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

Figure: Just Resynchronization Mechanism(JRM)



Information What is JRM?

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



Information Three Types of JRM

There are three types of Just Resynchronization Mechanism (JRM): for volumes, for slices and for proxy. For details, see "A.2.12 Just Resynchronization Mechanism (JRM)."

1.3 Functions for High Manageability

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

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

Functions improving disk manageability are explained below.

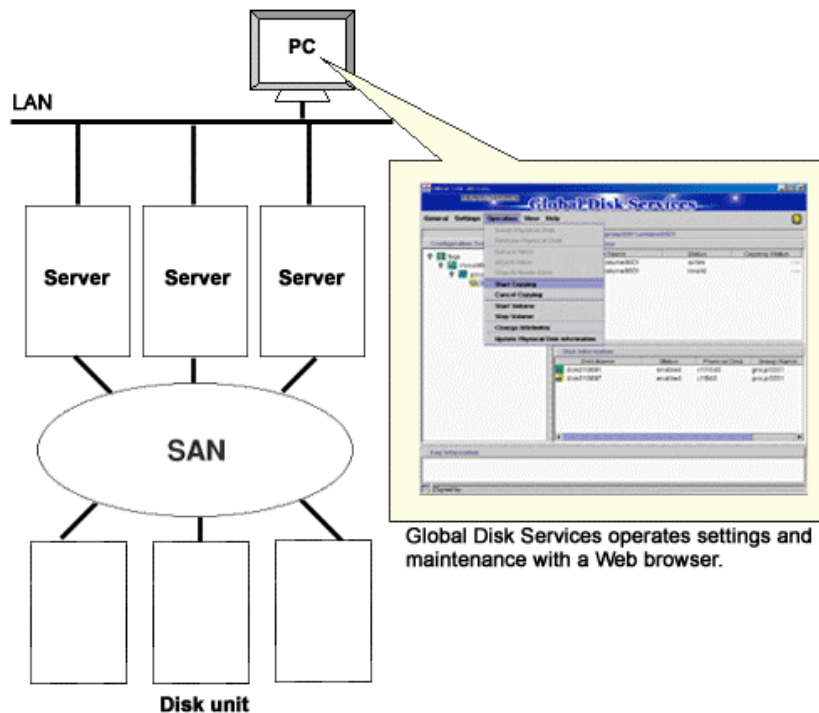
1.3.1 Operation Management Interface

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

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

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

Figure: Operation Management Interface



See

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

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

For the operations available on GUI and on CLI, see "5.1.9 Operations from GDS Management View."

1.3.2 Centralized Disk Management

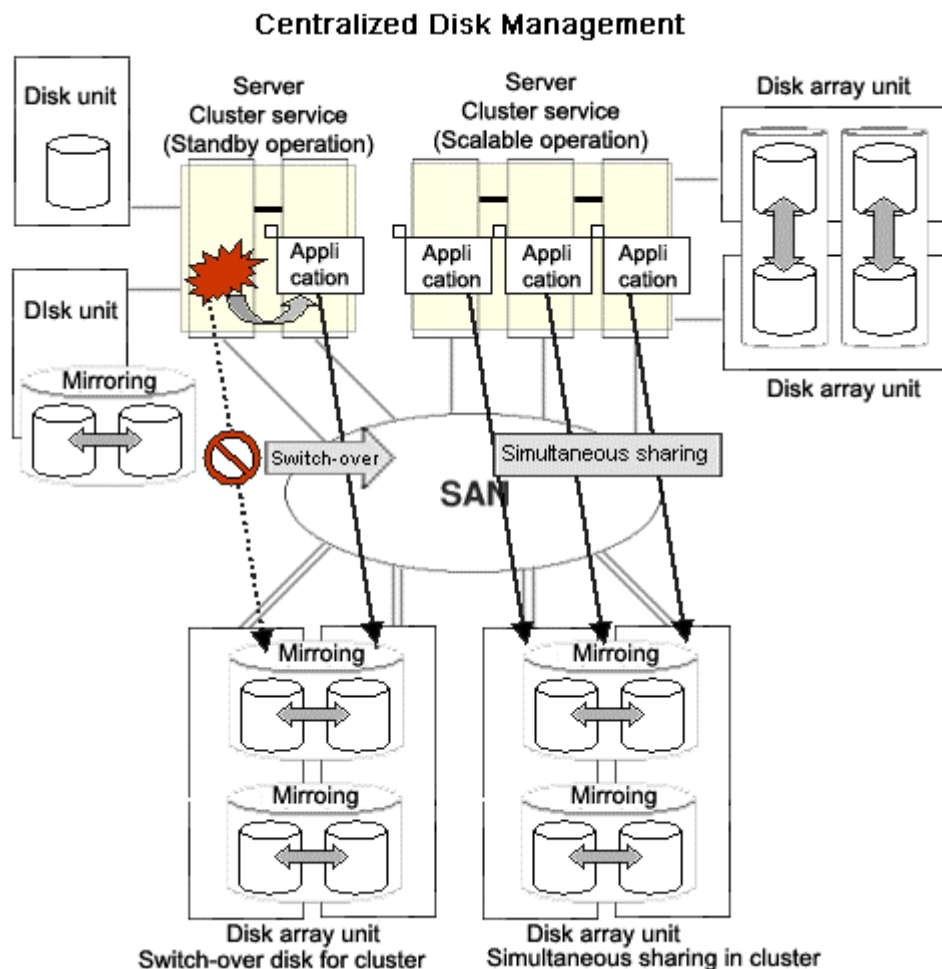
By registering all disk units connected to servers with GDS, it becomes possible to perform the centralized management of all operations on the disks from the GDS operation management interface.

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

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

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

Figure: Centralized Disk Management



Note System Disk and Local Disk Management

- With servers other than IPF servers, system disk management cannot be performed.
- If the Linux kernel version is 2.4, local disk management cannot be performed.

1.3.3 Name Management

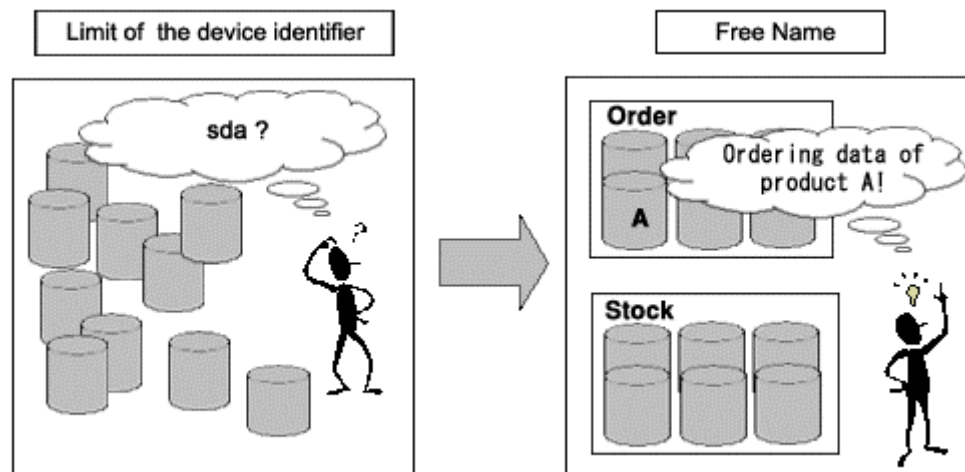
In Linux System, disks are named in "sdX" format, where "X" is an alphabetical character assigned by the OS, and the administrator differentiates the disks by the assigned consecutive characters.

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: 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: 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.

Class Scope


In order to manage a disk unit with GDS, it must be registered with a certain class.

A class is like a container that holds multiple disks.

By defining the class attribute called "scope", the user can specify which node set has access, or can perform operation on the disks registered to that class. Since operation on the disk is restricted to the nodes you specify, there is no risk of changing the configuration by mistake, or losing data consistency.

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

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




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: Starting and Stopping Volume



Access Mode

Logical volume has an attribute called "Access mode", which can be defined for each node sharing the volume. There are two access modes; "Read and write possible" mode, and "Read only possible" mode. For example, if a certain node will access the logical volume to create data backup, set the mode "Read only possible". That way, you can prevent writing to the volume by mistake.

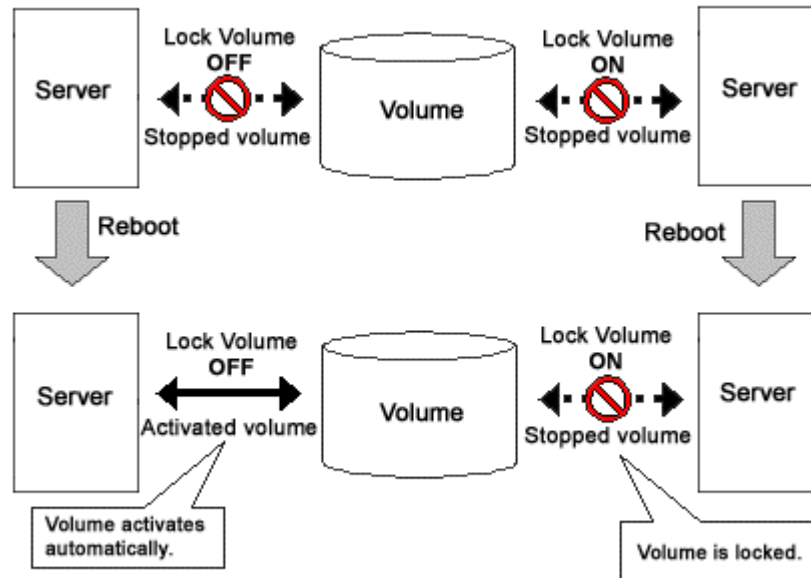
Figure: Access Mode



Lock Volume

When the node or cluster application is activated, logical volume starts automatically and becomes accessible. Likewise, when the cluster application is terminated, logical volume will also stop. This prevents a node with terminated application from accessing the logical volume. However, rebooting a node could result in starting the volume. In order to preclude the logical volume from starting in such a situation, the user can define the "Lock volume" attribute. When "Lock volume" is selected, volume will not be activated even when the node is rebooted, or cluster application activated.

Figure: 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 partition table.

With an external disk (physical disk) in a Linux system, in general, a maximum of 16 partitions can only be used.

GDS allows users to use physical disks and objects equivalent to disks dividing them into a maximum of 224 logical devices.

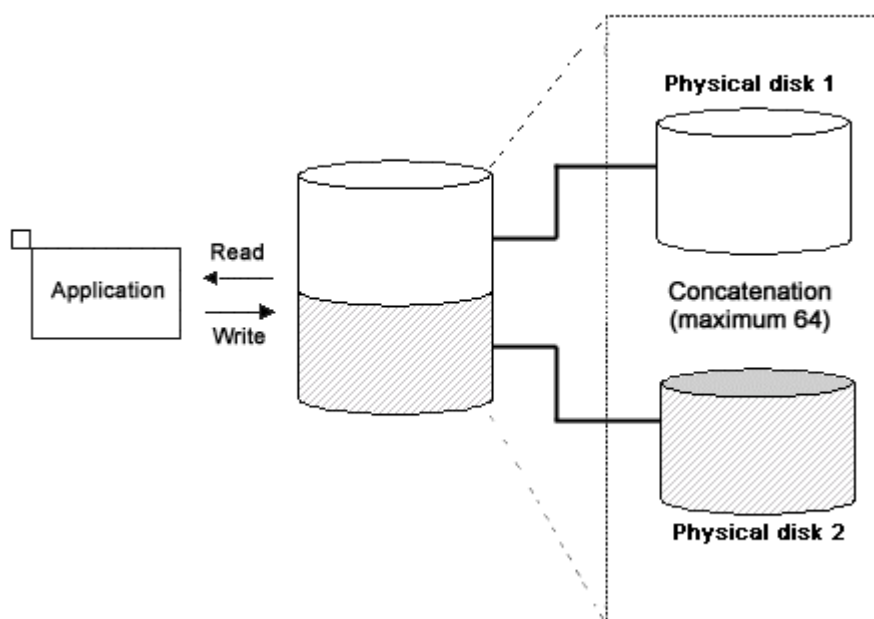
For details, see "2.5 Logical Volume" and "D.4 sdxvolume - Volume operations."

1.3.6.2 Disk Concatenation

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

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

Figure: Disk Concatenation



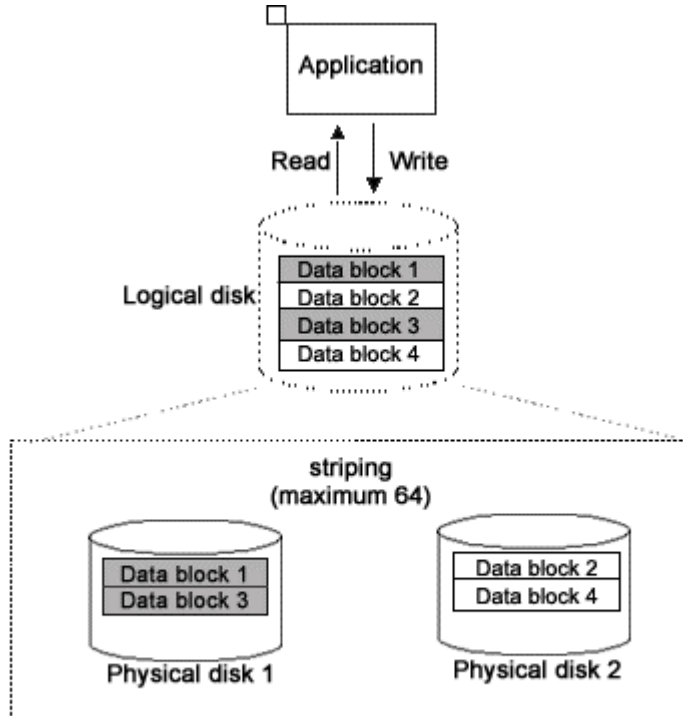
1.3.6.3 Disk Striping

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: 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: 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: Online Volume Expansion



Concatenation and Online Volume Expansion

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




Note Use Conditions on "Concatenation and Online Volume Expansion"

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

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

For details, see "Concatenation and Online Volume Expansion" in "A.2.13 Online Volume Expansion."

Figure: Concatenation and Online Volume Expansion



Online Disk Migration and Online Volume Expansion

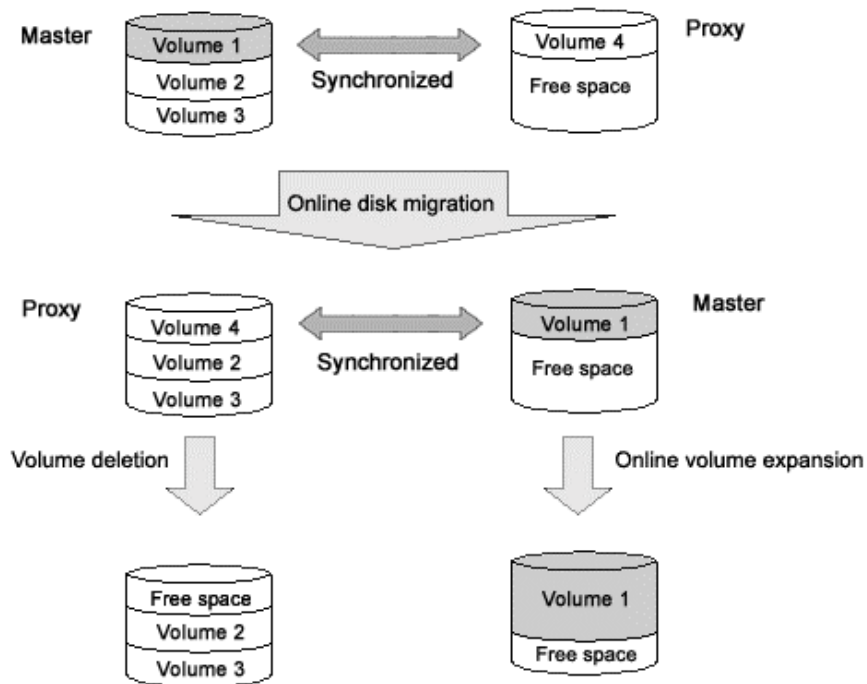
Even if there is no sufficient continuous free space after the last block of the volume, by migration of the volume to a disk with sufficient free space, online volume expansion will be available. For the volume migration, use the GDS Snapshot online disk migration function.



See

For online disk migration, see "1.5.5 Online Disk Migration."

Figure: Online Disk Migration and Online Volume Expansion



Online System Volume Expansion [Itanium]

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 an Alternative Boot Environment [Itanium]." For the procedures for expanding system volumes, see "D.20 Volume Expansion Using Commands [Itanium]."

Operating Instructions

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



Note Points of Concern for Online Volume Expansion

See "A.2.13 Online Volume Expansion."

1.3.8 Snapshots by Slice Detachment

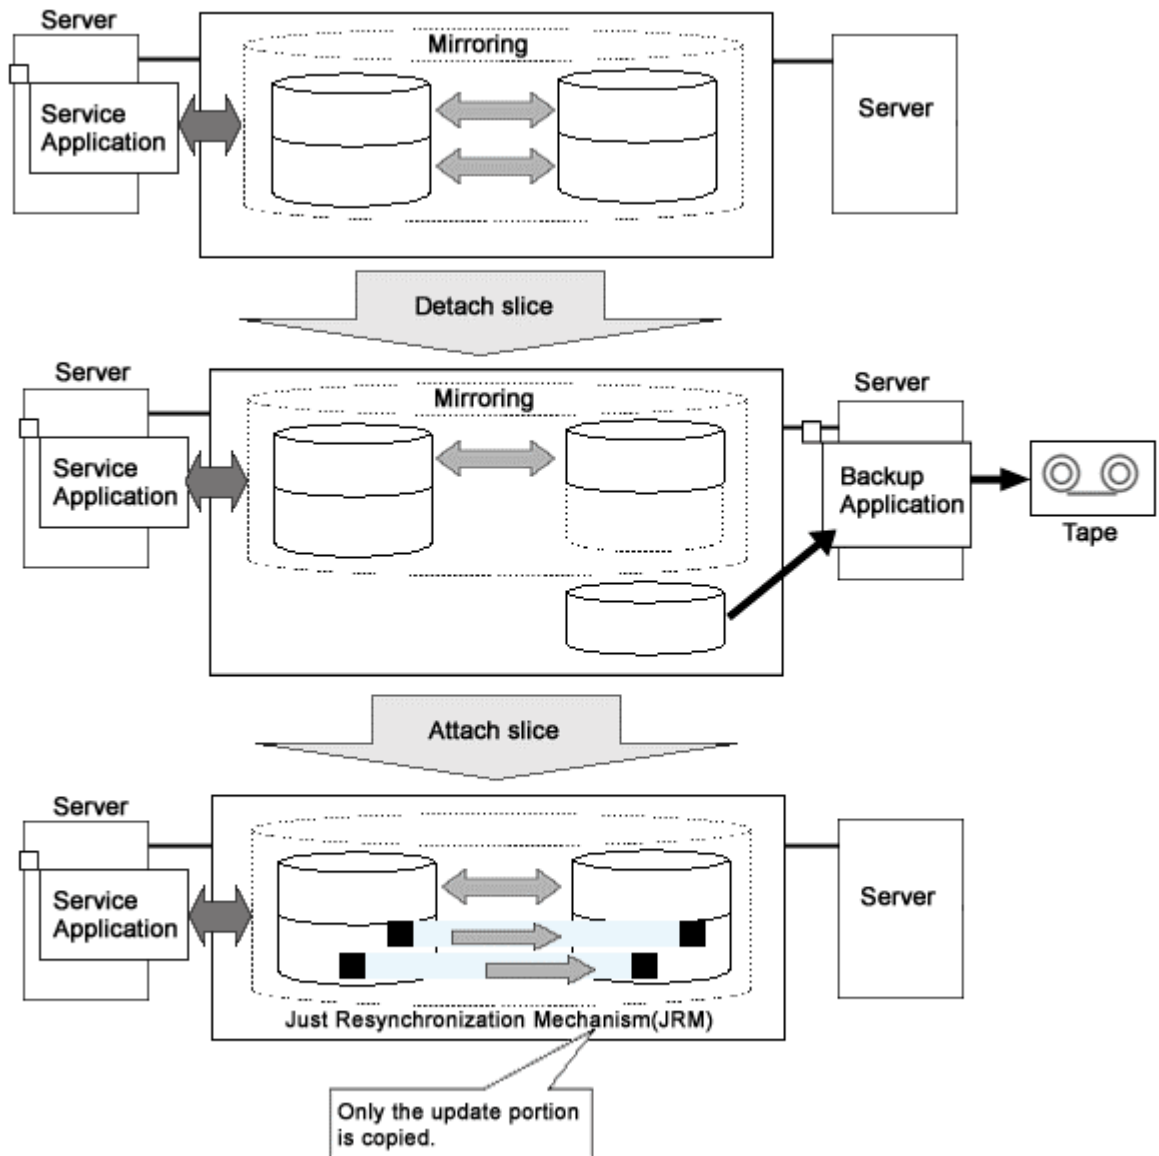
By temporarily detaching slices from mirror volumes, you can create snapshots (replications) of the volumes. In other words, by use of detached slices, volume data at the moment can be accessed.

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

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

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

Figure: Snapshot by Slice detachment





Note **Ensuring Consistency of Snapshot Data**

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

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

After creating the snapshot, start the application again.

For example, when using the volume as a file system such as GFS or ext3, unmount the volume with the `umount (8)` command before creating a snapshot, and mount it with the `mount(8)` command afterwards.

This way, you can ensure the consistency of the snapshot data.

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



Note **JRM for Slices**

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

JRM for slices becomes effective when a slice is detached while the `jrm` attribute of the slices is on.

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

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



Information **What is JRM?**

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



Information **Three Types of JRM**

There are three types of Just Resynchronization Mechanism (JRM): for volumes, for slices and for proxy. For details, see "A.2.12 Just Resynchronization Mechanism (JRM)."

1.4 GDS Snapshot Features

GDS Snapshot is software that provides additional functions to GDS.

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

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

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

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

1.5 Proxy Volume


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

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

Conventionally, data was backed up overnight when no one was on the system. However, as the amount of data grows, backup often cannot be completed by the following morning.

Also, the widespread usage of Internet and service diversification has made it impossible to stop many systems.


Figure: 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.

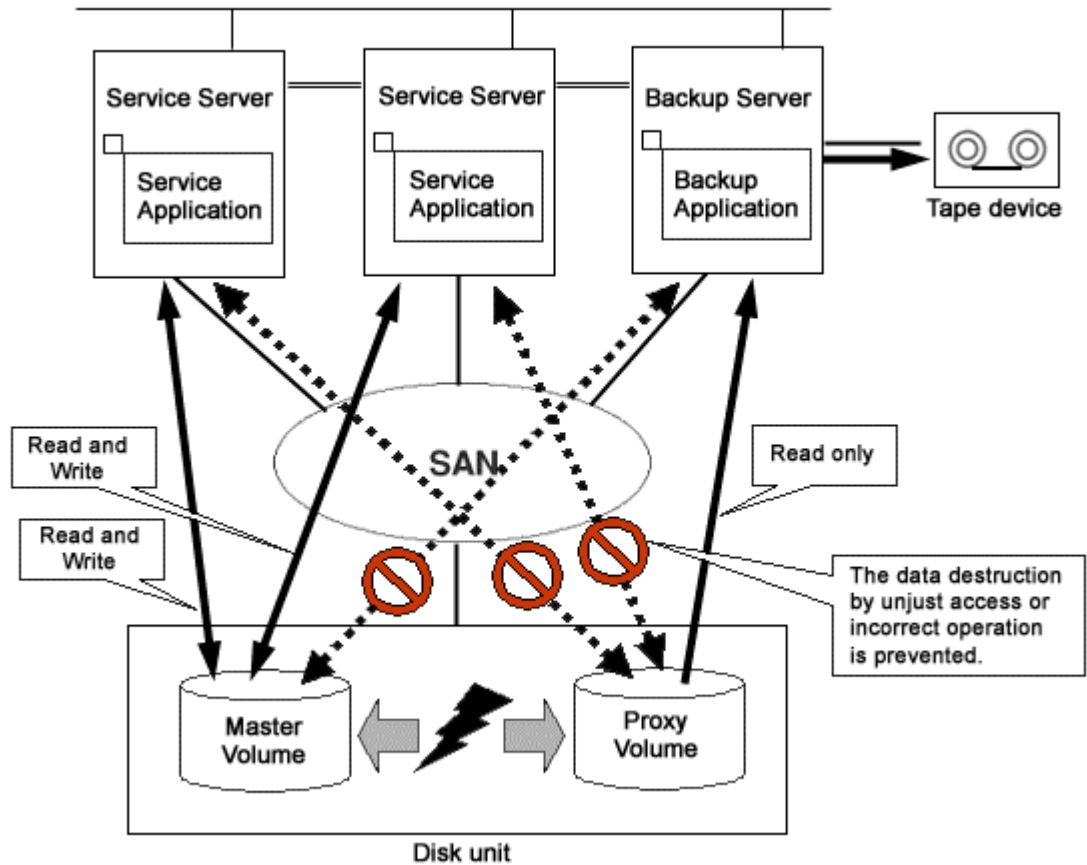
By using the proxy volume, users can conduct various tasks such as non-disruptive backup, data analysis, verification, and data protection from disaster, without affecting the main service application.

Figure: 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: Liaising with Volume Management Function
Cluster System (PRIMECLUSTER)



The snapshot functions using a proxy volume are explained here.

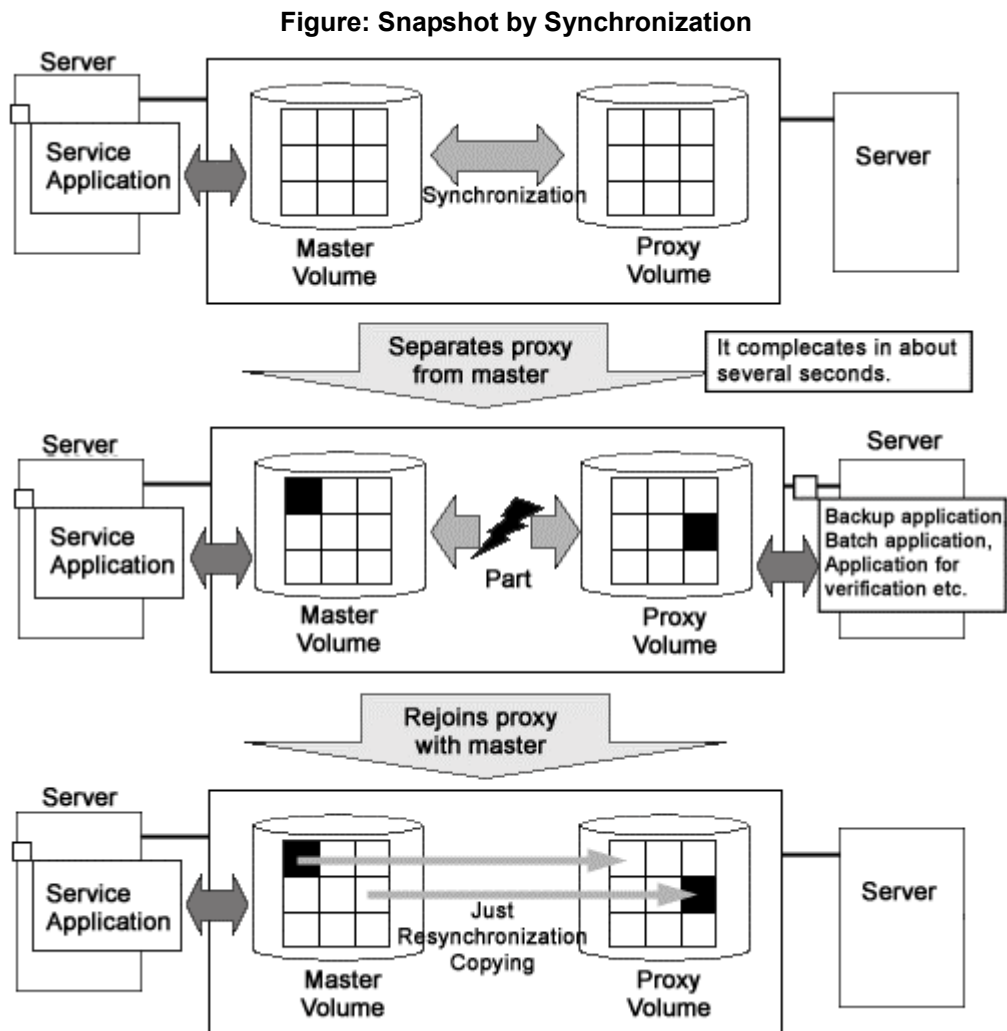
1.5.1 Snapshot by Synchronization

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.

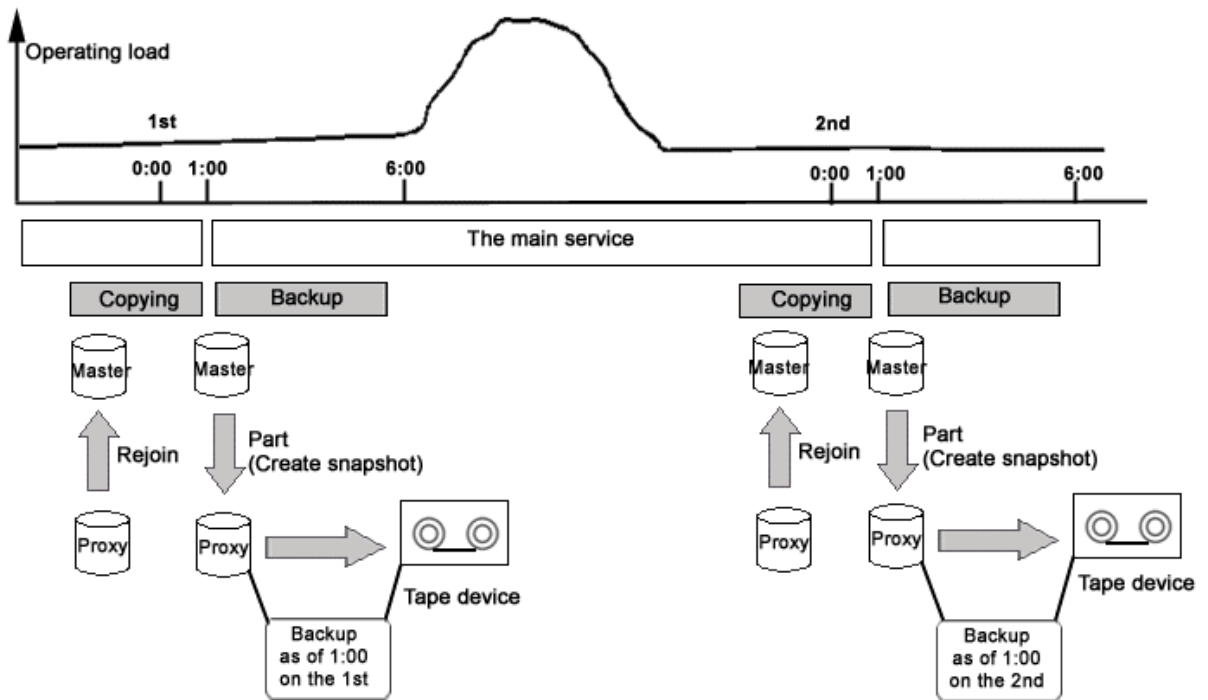
When the proxy volume is rejoined with the master volume for the next snapshot creation, high-speed resynchronization is performed by copying only updated portions on the master and proxy during the separation using JRM (Just Resynchronization Mechanism) for proxies.



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: Example of a Daily Backup Schedule



Note JRM for Proxies

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

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

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

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



Information Three Types of JRM

There are three types of Just Resynchronization Mechanism (JRM): for volumes, for slices and for proxy. For details, see "A.2.12 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 ext3, unmount the volume with the `umount (8)` command before creating a snapshot, and mount it with the `mount(8)` command afterwards. This way, you can ensure the consistency of the snapshot data.

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

For details, see "6.4 Online Backup and Instant Restore through Proxy Volume."



Note **The Difference between a Mirrored Slice and Synchronized Proxy Volume**

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

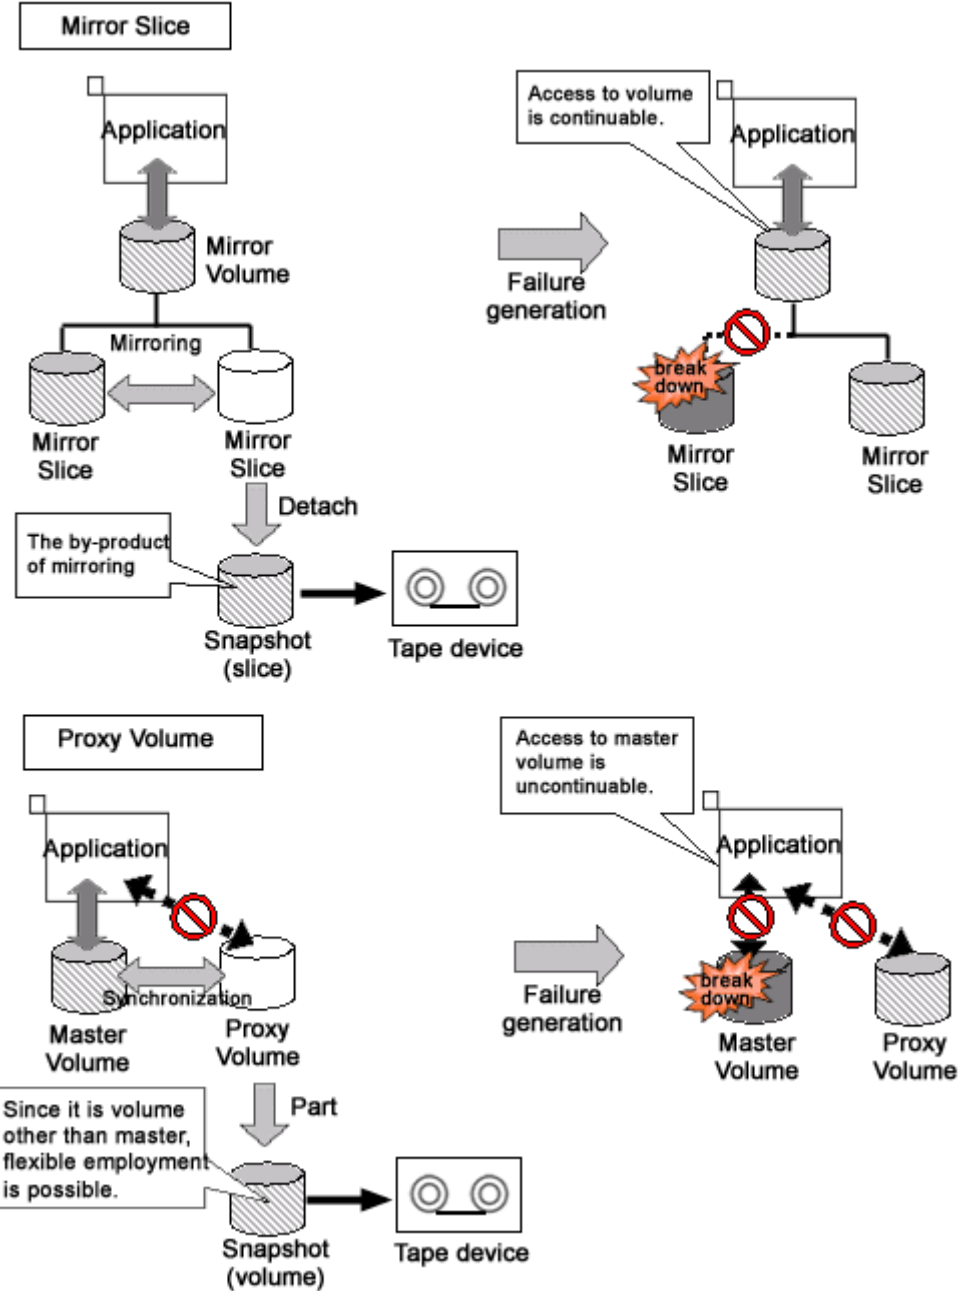
Mirrored slices are equals, and their purpose is to maintain data redundancy. So even if an abnormality is detected in one of the slices, the volume can be accessed as long as the other slice is normal.

However, even if the master volume and the proxy volume are synchronized, they are separate volumes and not equals. You may consider the master the primary volume, and the proxy the secondary volume. This means that even if the proxy volume is normal while all slices comprising the master volume are abnormal, you cannot continue accessing the master volume.

The purpose of proxy volumes is to separate the master volume in order to perform a different task in parallel, and not to improve the data redundancy used by the service application.

In short, the snapshot by slice detachment is a by-product of mirroring, which is the primary objective, but the primary objective of the snapshot through use of proxy volumes is creating snapshots. Therefore, GDS Snapshot provides more flexible disk configuration and supports various tasks.

Figure: Difference between a Mirrored Slice and Synchronized Proxy Volume




1.5.2 Snapshot Function with No Load to the Server/SAN

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

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

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



Note Working in Cooperation with EMC TimeFinder and SRDF
If the Linux kernel version is 2.4, EMC TimeFinder or SRDF cannot be used.

1.5.3 Instant Snapshot by OPC

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

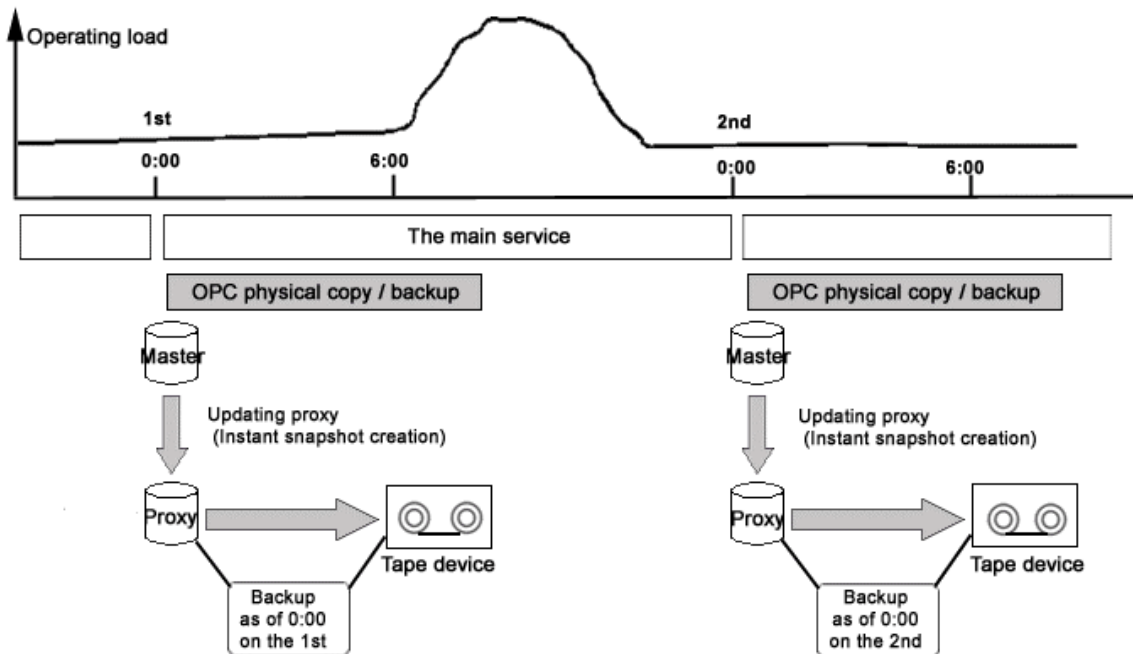
Unlike snapshot by synchronization, instant snapshot by OPC does not require synchronization with volumes in advance, so scheduling is unnecessary. However, as copying process will take place within the disk array after the snapshot has been created, I/O performance of the disk will be affected until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



Note Instant Snapshot by OPC

Instant snapshot by OPC is only available with the following storage systems supporting the function: ETERNUS8000 series, ETERNUS6000 series, ETERNUS4000 series, ETERNUS3000 series, GR720, GR730, GR740

Figure: 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.

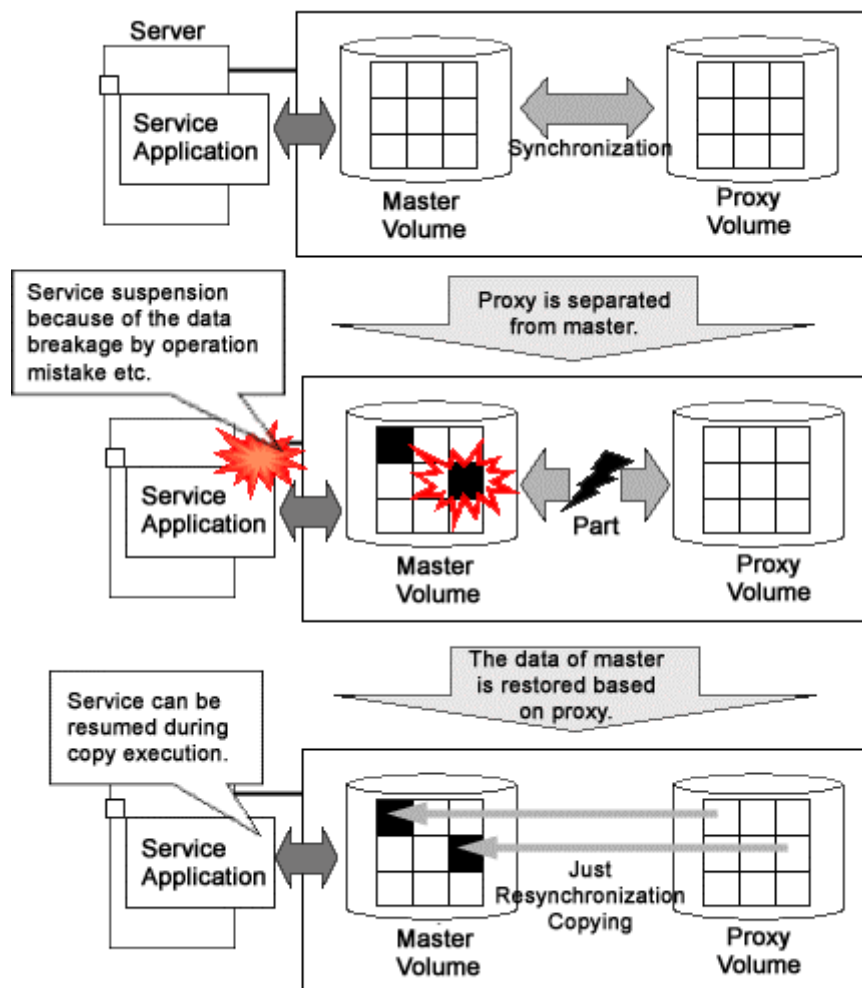
Users may use this function with any disk unit, as it is not dependent on a specific function of a disk unit.



Note System Volume Restoration [Itanium]


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

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

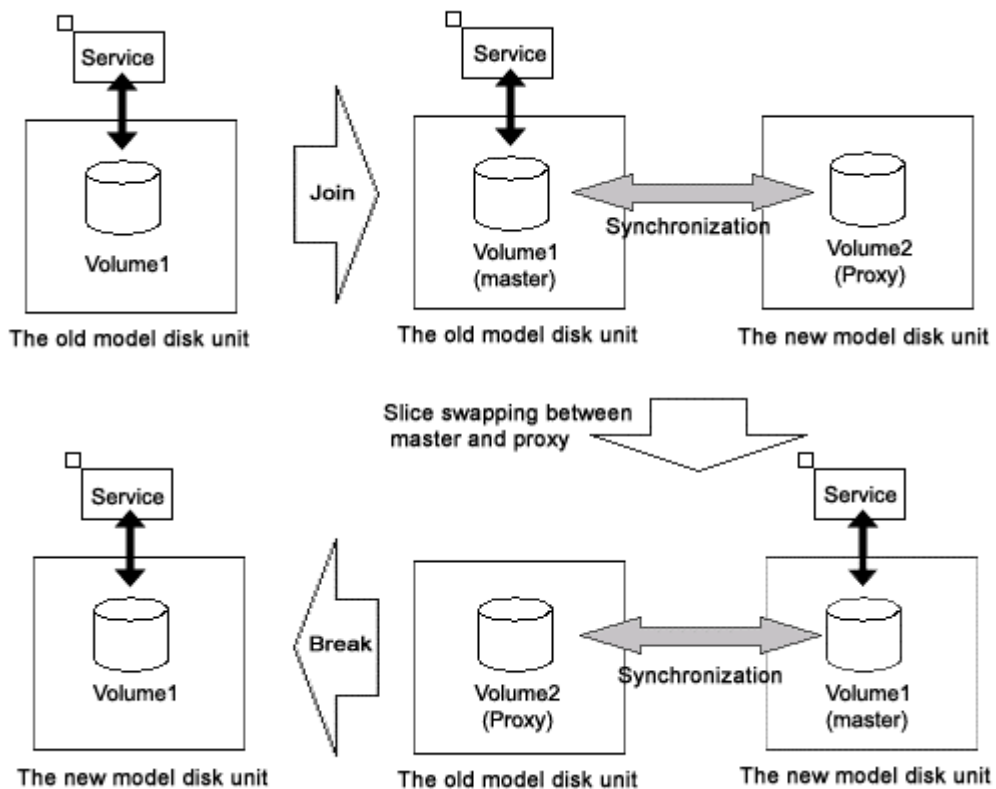


1.5.5 Online Disk Migration

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: Online Disk Migration



1.5.6 Creating an Alternative Boot Environment [Itanium]

Using the snapshot function by proxy volumes, another boot environment (alternative boot environment) can be created without influencing the current boot environment during system operation. System stop time at the time of backup and restore work for system volumes (`/`, `/usr`, `/var`, `/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**
Make booting available through snapshot volumes (proxy volumes) instead of system volumes in the current boot environment. This configuration changes the `fstab` files, the `elilo.conf` files and so on copied from the root volume (`/`) to the proxy volumes when snapshots were created. The current boot environment is not changed.

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

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

- **Backing up the system**
An alternative boot environment can be assumed as a backup of the system. Since an alternative boot environment can be created during system operation, stopping the system is no longer necessary at the time of backup work.
- **Restoring when the system fails**
When booting the system becomes unavailable due to a system disk failure or data crash, the system operation can be resumed by switching to an alternative boot environment. In addition, data on original system volumes can be restored from data in an alternative boot environment during system operation in the alternative boot environment, and the original boot environment can be switched back to by simply rebooting the system.
- **Reconfiguring the system**
The disk configuration of system volumes can be changed by creating an alternative boot environment and switching to the alternative boot environment. In addition, expanding the capacity and applying patches for volumes used in the alternative boot environment are available during system operation. These operations can be performed during system operation, and have no effect on the current boot environment. If booting an alternative boot environment fails or the performance of an alternative boot environment is improper, the current boot environment can be switched back to by simply rebooting the system.



See

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

1.6 Shadow Volume

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

This section describes the shadow volume feature.

1.6.1 Accessing Volumes from an External Server

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

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


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



See

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

Figure: Backup Server



1.6.2 Using the Copy Function of a Disk Unit

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


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



See

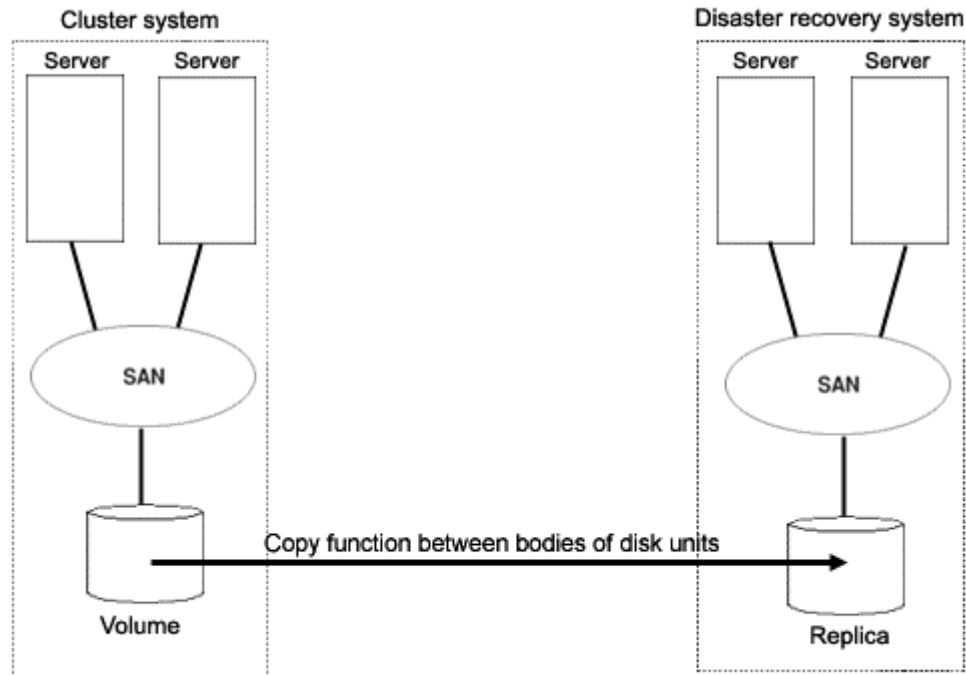
For the operations patterns and the procedure for backup and restore using the copy function of a disk unit and using shadow volumes, see "6.5 Backing Up and Restoring through Disk Unit's Copy Functions [Linux2.6]."

Figure: Backing Up a Replication Created with the Disk Unit's 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: Disaster Recovery System



Chapter 2 Objects

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

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

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

2.1 SDX Object

Virtual resources managed by GDS are called SDX objects, or objects.
There are five kinds of objects: classes, disks, groups, volumes and slices.

2.2 Disk Class

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

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

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

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



Note

Same Class Names [Linux2.6]

When multiple single nodes on which classes with the same name exist are changed over to a cluster system through installation of the cluster control facility, the duplicate class names come to exist in the cluster system. For details, see "A.2.24 Changing Over from Single Nodes to a Cluster System [Linux2.6]."

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 [Itanium]

Object managed in a class that is available only to its own node. "Root" type can manage system disk, while "local" type cannot.

Local [Linux2.6]

Object managed in a class that is available only to its own node.

Shared

Object managed in a class that is available to more than one node.

Scope

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

Hot Spare

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

on

Hot spare is enabled.

off

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

Hot Spare Mode

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

External

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

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

Internal

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

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

Operation

The following operations are available for classes.

Create

Registering a disk with the `sdxdisk -M` command and specifying a new class name will automatically create a new class. For details on GDS Management View, see "5.2.2.1 Class Configuration."

Delete

Deleting the last registered disk from a class using the `sdxdisk -R` command will automatically delete the class. You can also delete the classes with the `sdxclass -R` command or the `sdxconfig Remove` command [Linux2.6]. For details on GDS Management View, see "5.5.4 Removing a Class."

Status Display

Class status can be displayed with the `sdxinfo` command. For details on GDS Management View, see "5.3.1 Viewing Configurations and Monitoring Statuses."

Change Attributes

Class attribute values can be changed with the `sdxattr -C` command. For details on GDS Management View, see "5.4.1 Class Configuration."

Recover

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

Backup [Linux2.6]

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 [Linux2.6]

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 [Linux2.6]

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

Status

A class can have the following status.

Closed

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

Reference

There are the following points of concern for classes.

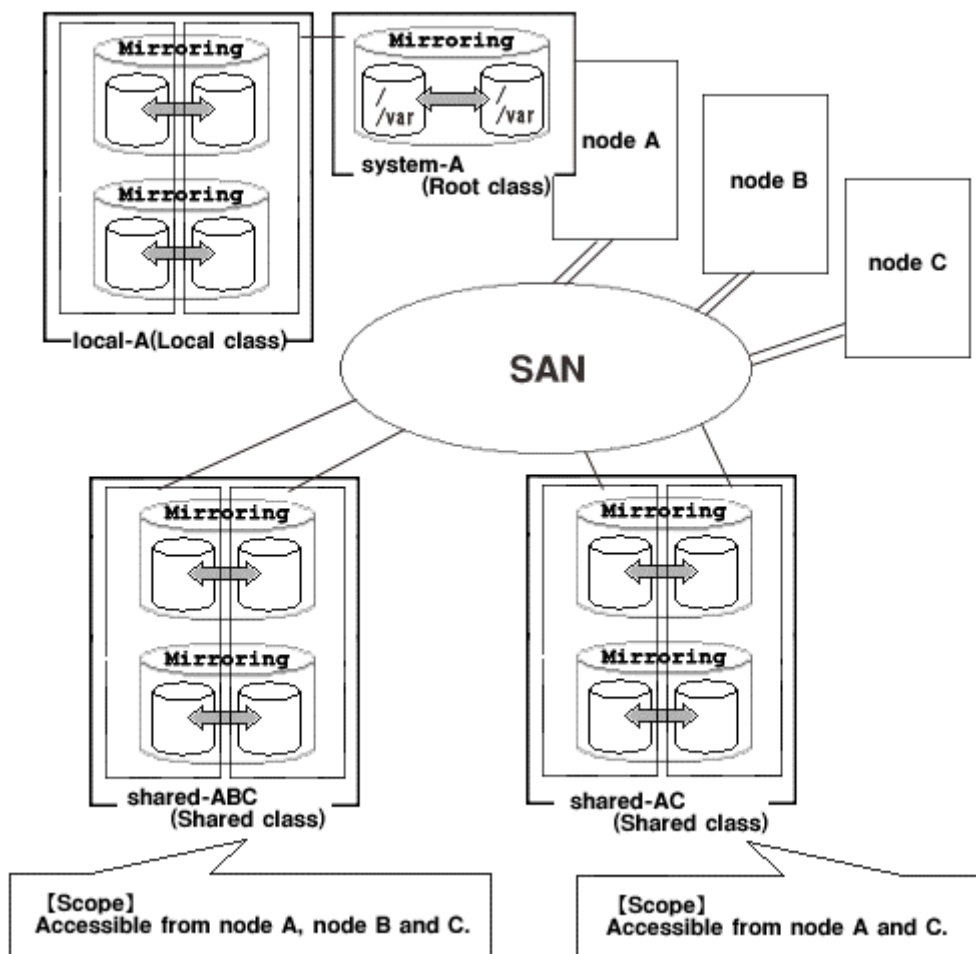
Rules

- A.1.1 Object Name
- A.1.2 Number of Classes

Important Points

- A.2.1 Managing System Disks
- A.2.3 Booting from a CD-ROM Device
- A.2.8 System Disk Mirroring [Itanium]
- A.2.23 Operation in Cluster System
- A.2.24 Changing Over from Single Nodes to a Cluster System [Linux2.6]
- A.2.27 Configuration Tables [Linux2.6]

Figure: Disk Class



2.3 SDX Disk

Physical disks managed by GDS are registered with a certain class and are called SDX disks, or simply "disks." Disks that no longer need to be managed with GDS can be deleted from the class and are returned to their original physical disk status.

Attributes

A disk has the following attributes.

Name

This attribute identifies the disk within a class.

Type

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

Mirror

The disk has been connected to a mirror group.

Stripe

The disk has been connected to a stripe group.

Concatenation

The disk has been connected to a concatenation group.

Switch

The disk is connected to a switch group.

Keep [Itanium]

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

Single

A volume can be created without connecting to a group.

Spare

The disk will be used as a spare disk.

Undefined

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

Operation

The following operations are available for disks.

Create

A disk can be created by registering a physical disk with a certain class, using the `sdxdisk -M` command. For details on GDS Management View, see "5.2.2.1 Class Configuration."

Delete

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

Connect

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

Disconnect

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

Swap

A disk can be prepared for swapping with the `sdxswap -O` command. For details on GDS Management View, see "5.3.4 Disk Swap."

Recover

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

Status Display

The disk status can be displayed with the `sdxinfo` command. For details on GDS Management View, see "5.3.1 Viewing Configurations and Monitoring Statuses."

Change Attributes

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

Clear Errors

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

Status

A disk can have the following status.

ENABLE

The disk is in the normal status.

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

DISABLE

Data on the disk is inaccessible.

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

For details, see "F.1.2 Disk Status Abnormality."

SWAP

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

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

Reference

There are the following points of concern for disks.

Rules

A.1.3 Number of Disks

A.1.6 Number of Keep Disks [Itanium]

Important Points

A.2.2 Restraining Access to Physical Special File

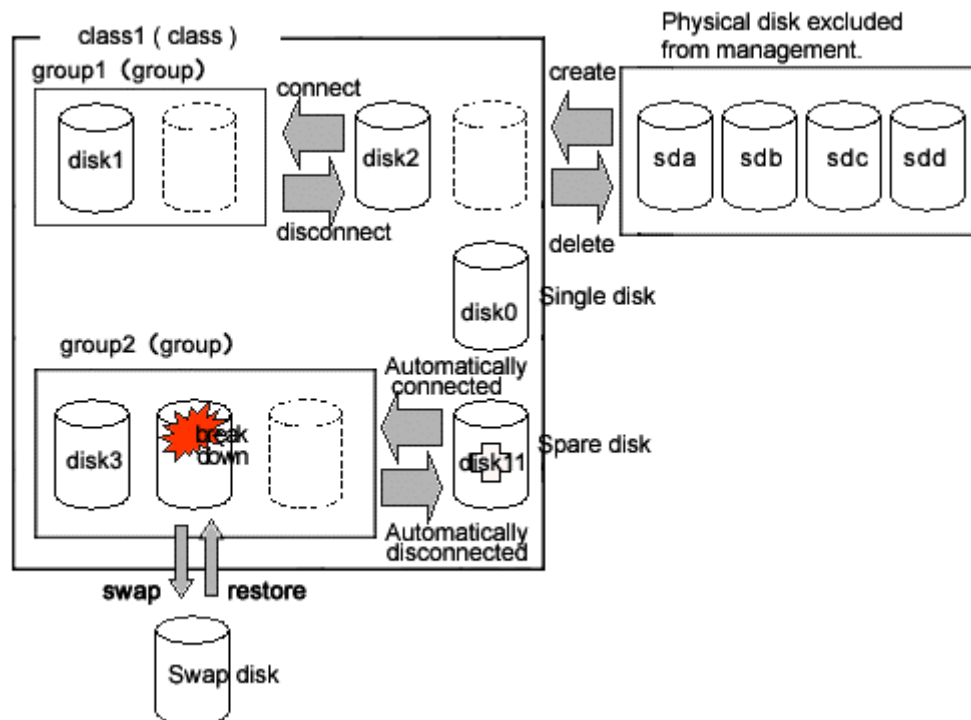
A.2.4 Initializing Disk

A.2.5 Disk Size

A.2.7 Hot Spare

A.2.9 Keep Disk [Itanium]

Figure: SDX Disk



2.4 Disk Group

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

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

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

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

Attributes

A group has the following attribute.

Name

This attribute identifies the group within a class.

Type

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

mirror

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

stripe

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

concat

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

switch

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

Stripe Width

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

Active Disk

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

Operation

The following operations are available for groups.

Create

Indicating a new (higher level) group when connecting a disk with the `sdxdisk -C` command, or connecting a group with `sdxgroup -C` command will automatically create a (higher level) group. For details on GDS Management View, see "5.2.2.3 Group Configuration."

Delete

Disconnecting the only remaining disk with the `sdxdisk -D` command, or disconnecting the only remaining lower level group with the `sdxgroup -D` command will automatically remove the (higher level) group. You can also delete a group using the `sdxgroup -R` command. For details on GDS Management View, see "5.5.3 Removing a Group."

Connect

You can add a group to another group with the `sdxgroup -C` command. A group that is connected to another group is called a lower level group, and a group to which the group is added is called a higher level group. For details on GDS Management View, see "5.4.2 Group Configuration."

Disconnect

The `sdxgroup -D` command disconnects the lower level group from the higher level group. For details on GDS Management View, see "5.4.2 Group Configuration."

Status Display

Group status can be displayed with the `sdxinfo` command. For details on GDS Management View, see "5.3.1 Viewing Configurations and Monitoring Statuses."

Change Attributes

Group attribute values can be changed with the `sdxattr -G` command. For details on GDS Management View, see "5.4.2 Group Configuration."

Reference

There are the following points of concern for groups.

Rules

- A.1.4 Number of Groups
- A.1.7 Creating Group Hierarchy

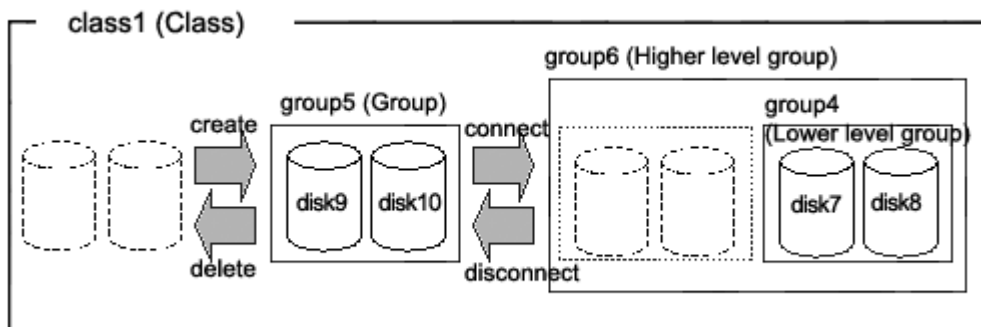
Important point

- A.2.25 Disk Switch

Guide

- A.3.1 Guidelines for Mirroring
- A.3.2 Guidelines for Striping
- A.3.3 Guidelines for Concatenation
- A.3.4 Guidelines for Combining Striping with Mirroring

Figure: Disk Group



2.5 Logical Volume

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

An application will access the logical volume instead of the physical disk. Logical volumes are also called "volumes."

There are five kinds of volumes as follows:

- **Single Volume Created in a Single Disk**
Data will not be redundant.
Single volumes are used when it does not necessary to be mirrored, but needs to be managed by GDS. By connecting the single disk to a mirror group, a single volume can be changed to a mirror volume while retaining its data.
- **Mirror Volume Created in a Mirror Group**
When multiple disks and lower groups are connected to the mirror group, data will be redundant through the mirroring function.
- **Stripe Volume Created in a Stripe Group**
The striping function enables I/O load sharing across multiple disks.
Data will not be redundant.
- **Volume Created in a Concatenation Group**
Concatenation group allows users to create a large capacity volume spanning multiple disks. Its data will not be redundant.
- **Switch Volume Created in a Switch Group**
If an inactive disk is connected to a switch group in addition to the active disk, the active disk can be changed to the inactive disk with the disk switch function.

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

With an external disk (physical disk) in a Linux system, in general, a maximum of 16 partitions are only available. 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 224 volumes and be used.

Attributes

A volume has the following attributes.

Name

This attribute identifies the volume within a class.

JRM

This attribute indicates the just resynchronization mechanism mode for volumes.

on

JRM for volumes is enabled.

off

JRM for volumes is disabled.

Lock Volume

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

on

The volume will be locked and prevented from activating.

off

The volume will not be locked.

Access Mode

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

rw

The default access mode is set to read and write.

ro

The default access mode is set to read only.

Physical Slice

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

on

When the volume is a single volume, a slice in the single disk will be registered with the partition table. 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 partition table. When the volume is a switch volume, slices on all of the disks connected to the switch group will be registered with partition table.

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 partition table.

PJRM

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

on

JRM for proxies is enabled.

off

JRM for proxies is disabled.

Operation

The following operations are available for volumes.

Create

A volume can be created in a single disk or the highest level group with the `sdxvolume -M` command. For details on GDS Management View, see "5.2.2.4 Volume Configuration."

Delete

Volume will be deleted with the `sdxvolume -R` command. For details on GDS Management View, see "5.5.2 Removing a Volume."

Start

Volume will be started with the `sdxvolume -N` command. For details on GDS Management View, see "4.2.3 Operation."

Stop

Volume will be stopped with the `sdxvolume -F` command. For details on GDS Management View, see "4.2.3 Operation."

Expand

The volume size can be expanded with the `sdxvolume -S` command. The operation from GDS Management View is unsupported.

Status Display

Volume status can be displayed with the `sdxinfo` command. For details on GDS Management View, see "5.3.1 Viewing Configurations and Monitoring Statuses."

Change Attributes

Volume attribute values can be changed with the `sdxattr -V` command. For details on GDS Management View, see "5.4.3 Volume Configuration."

Copy control

Synchronization copying of volume can be controlled with the `sdxcopy` command. For details on GDS Management View, see "5.3.6 Copying Operation."

Recover

Recovery attempt of an abnormal volume can be performed with the `sdxfix -V` command. The operation from GDS Management View is unsupported.

Status

A volume can have the following status.

ACTIVE

Valid data is accessible.

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

STOP

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

However, a proxy volume must first be parted from the master volume before activating. After the volume is stopped normally, it is given the STOP status.

Here, there are one or more slices with STOP or COPY (in process of background copying) status within the volume.

INVALID

Data is invalid and inaccessible.

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

For details, see "F.1.3 Volume Status Abnormality."

Reference

There are the following points of concern for volumes.

Rules

A.1.5 Number of Volumes

Important Point

A.2.6 Volume Size

A.2.12 Just Resynchronization Mechanism (JRM)

A.2.22 Volume Access Mode

Figure: Logical Volume

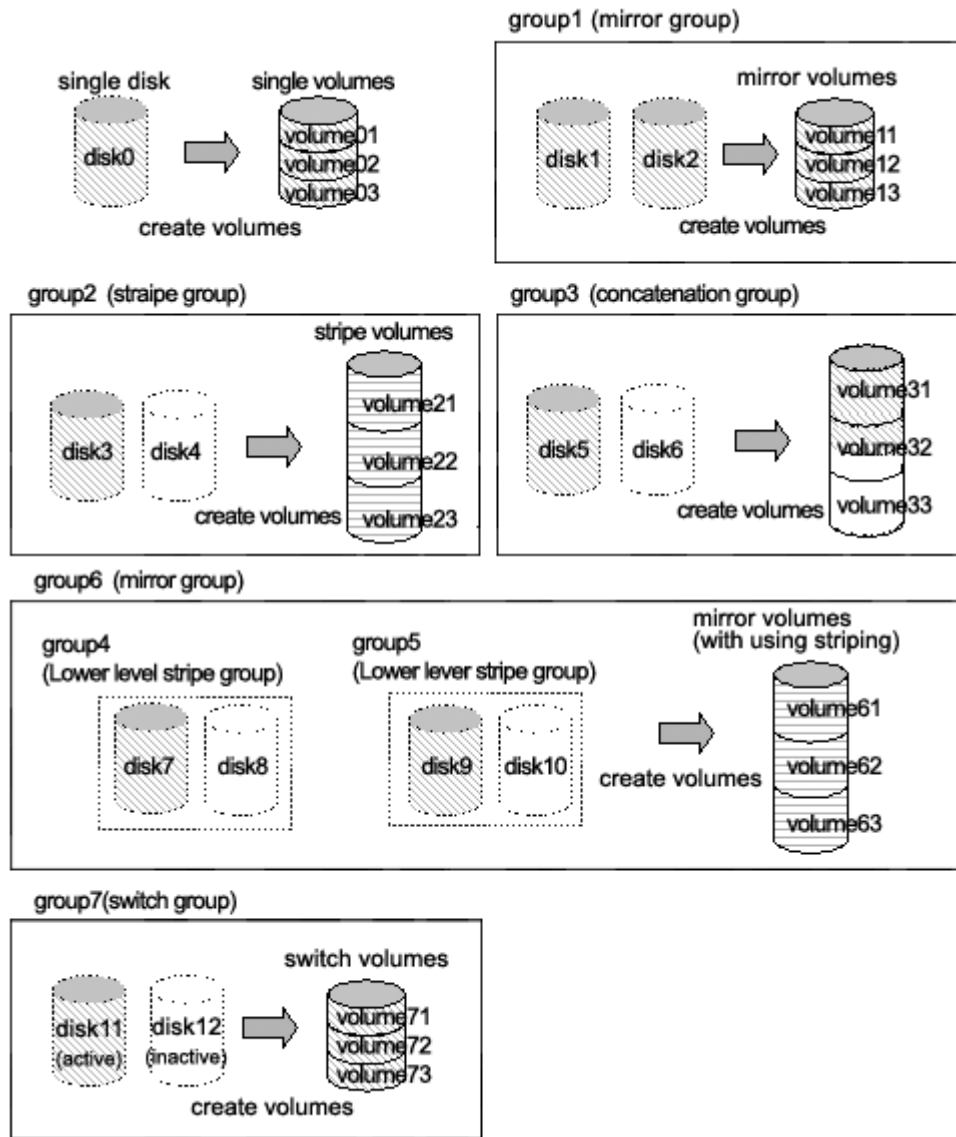



Figure: Logical Partitioning (Physical Slice Attribute)



2.6 Logical Slice

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

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

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

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

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

Attributes

A slice has the following attributes.

Name

This attribute identifies the slice within a class.

JRM

This attribute indicates the just resynchronization mechanism mode for slices.

on

JRM for slices is enabled.

off

JRM for slices is disabled.

Access Mode

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

rw

Sets access mode for read and write.

ro

Sets access mode for read only.

Operation

The following operations are available for slices.

Detach

A snapshot of a mirror volume can be created by detaching one of the slices consisting the mirror volume with the `sdxslice -M` command.

For details on GDS Management View, see "5.3.2.1 Backup (by Slice Detachment)."

Attach

The temporarily detached slice can be attached to the volume again with the `sdxslice -R` command.

For details on GDS Management View, see "5.3.2.1 Backup (by Slice Detachment)."

Activate

The detached slice can be activated with the `sdxslice -N` command.

For details on GDS Management View, see "5.3.2.1 Backup (by Slice Detachment)."

Stop

The detached slice can be stopped with the `sdxslice -F` command.

For details on GDS Management View, see "5.3.2.1 Backup (by Slice Detachment)."

Take over

Detached slice can be taken over from another node with the `sdxslice -T` command.

The operation from GDS Management View is unsupported.

Status Display

Slice status can be displayed with the `sdxinfor` command.

For details on GDS Management View, see "5.3.1 Viewing Configurations 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).

STOP

Data is normal, but inaccessible.

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

INVALID

Data is invalid and inaccessible.

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

Promptly identify the problem and restore the slice based on the disk driver message or the like.

After restoring the slice, execute the synchronization copying. When it ends successfully, the slice will be valid (the slice will be made the ACTIVE or STOP status). If it fails, it will be INVALID again.

COPY

Synchronization copying is in process to ensure the data integrity.

Synchronization copying is executed from a slice in the valid status (ACTIVE or STOP) to a slice in the COPY status.

When it ends successfully, the slice will be made valid (ACTIVE or STOP). If it fails, it will be INVALID.

TEMP

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

When the `sdxslice -M` command ends successfully, the slice is made the TEMP status.

To restore the status (ACTIVE or STOP), execute the `sdxslice -R` command.

TEMP-STOP

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

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

NOUSE

The slice is inaccessible for a special reason.

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

Reference

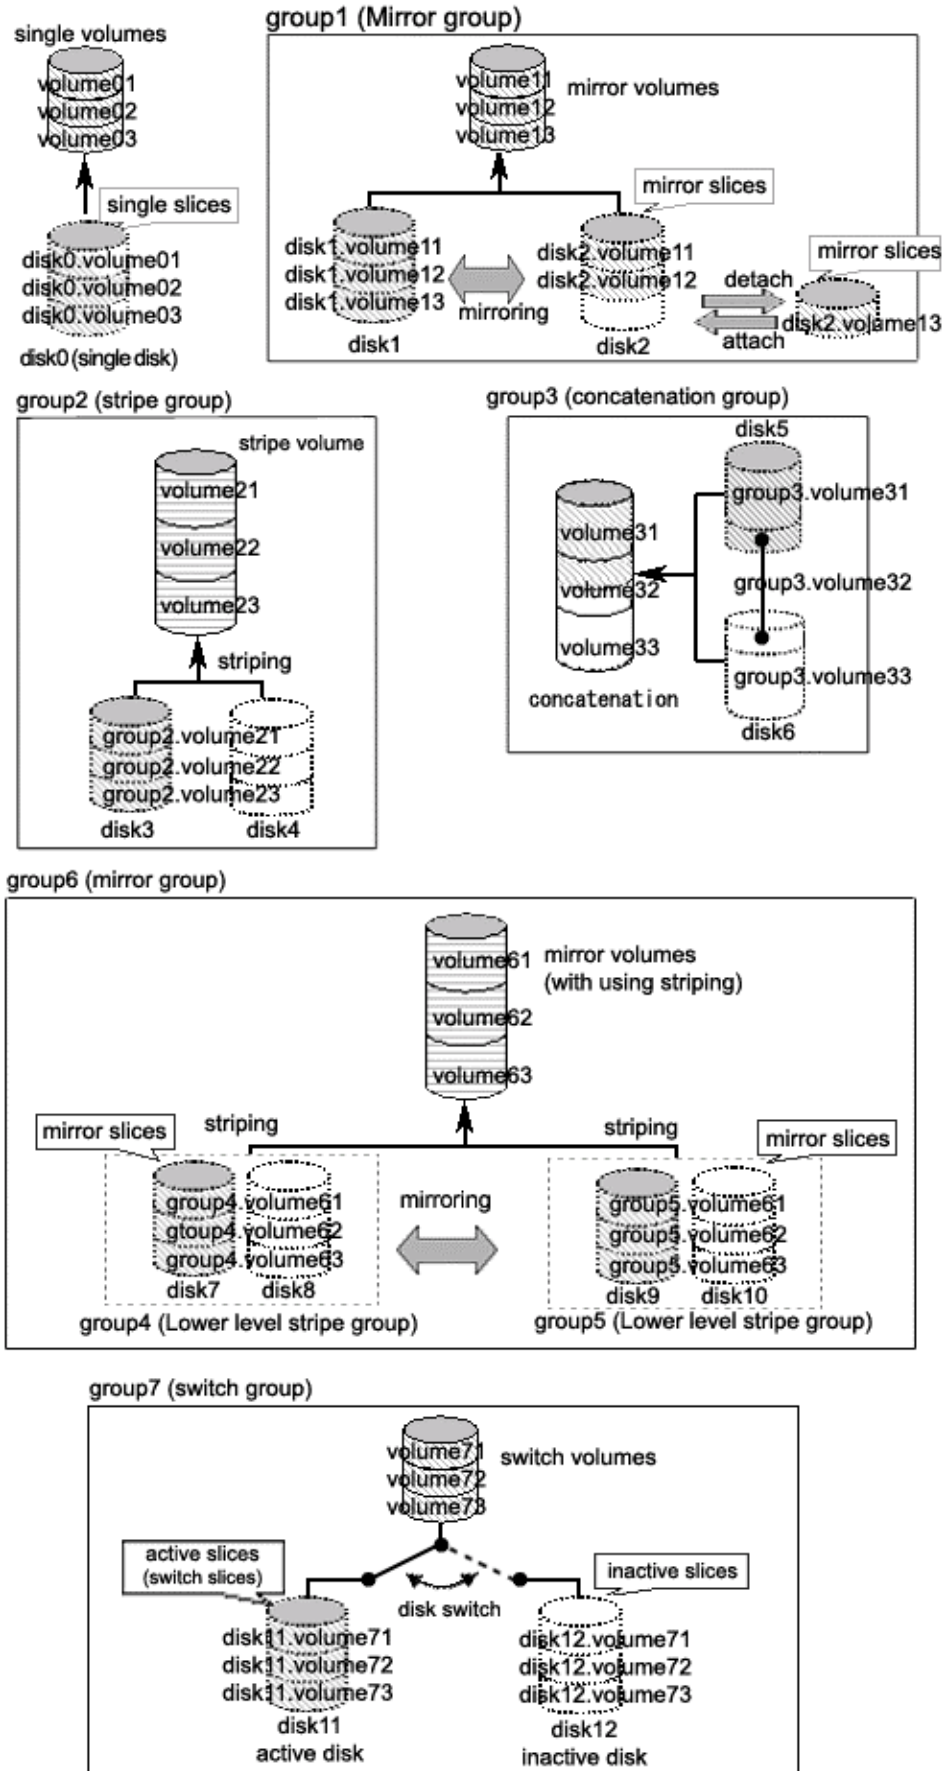
There are the following points of concern for slices.

Important Points

A.2.10 Creating a Snapshot by Slice Detachment

A.2.12 Just Resynchronization Mechanism (JRM)

Figure: Logical Slice



2.7 GDS Snapshot Objects

2.7.1 Proxy Object

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

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

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

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

Operation

The following operations are available for proxy objects.

Join

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

For details on GDS Management View, see "5.2.4.1 Join."

Part

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

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

Rejoin

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

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

Rejoin and Restore

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

For details on GDS Management View, see "5.3.3 Restore."

Swap

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

For details on GDS Management View, see "5.3.5 Disk Migration."

Relate

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

For details on GDS Management View, see "5.2.4.2 Relate."

Update

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

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

Restore

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

For details on GDS Management View, see "5.3.3 Restore."

Cancel Copy Session

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

The operation from GDS Management View is unsupported.

Configure an Alternative Boot Environment [Itanium]

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

The operation from GDS Management View is unsupported.

Break

The specified relationship between a pair of volumes or groups as the master and proxy will be cancelled and they will return to individual objects with the `sdxproxy Break` command. For details on GDS Management View, see "5.5.6 Breaking a Proxy."

Status

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

Joined

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

Parted

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

Reference

There are the following points of concern for proxy objects.

Rules

- A.1.8 Exception to Proxy Configuration
- A.1.9 Number of Proxy Volumes
- A.1.10 Proxy Volume Size
- A.1.11 Proxy Group Size

Important Points

- A.2.11 The Difference between a Mirror Slice and a Proxy Volume
- A.2.12 Just Resynchronization Mechanism (JRM)
- A.2.15 Object Operation When Using Proxy
- A.2.16 Using the Advanced Copy Function in a Proxy Configuration
- A.2.17 Instant Snapshot by OPC
- A.2.19 Using EMC TimeFinder or EMC SRDF in a Proxy Configuration [Linux2.6]
- A.2.20 Ensuring Consistency of Snapshot Data

Figure: Operating Proxy Objects

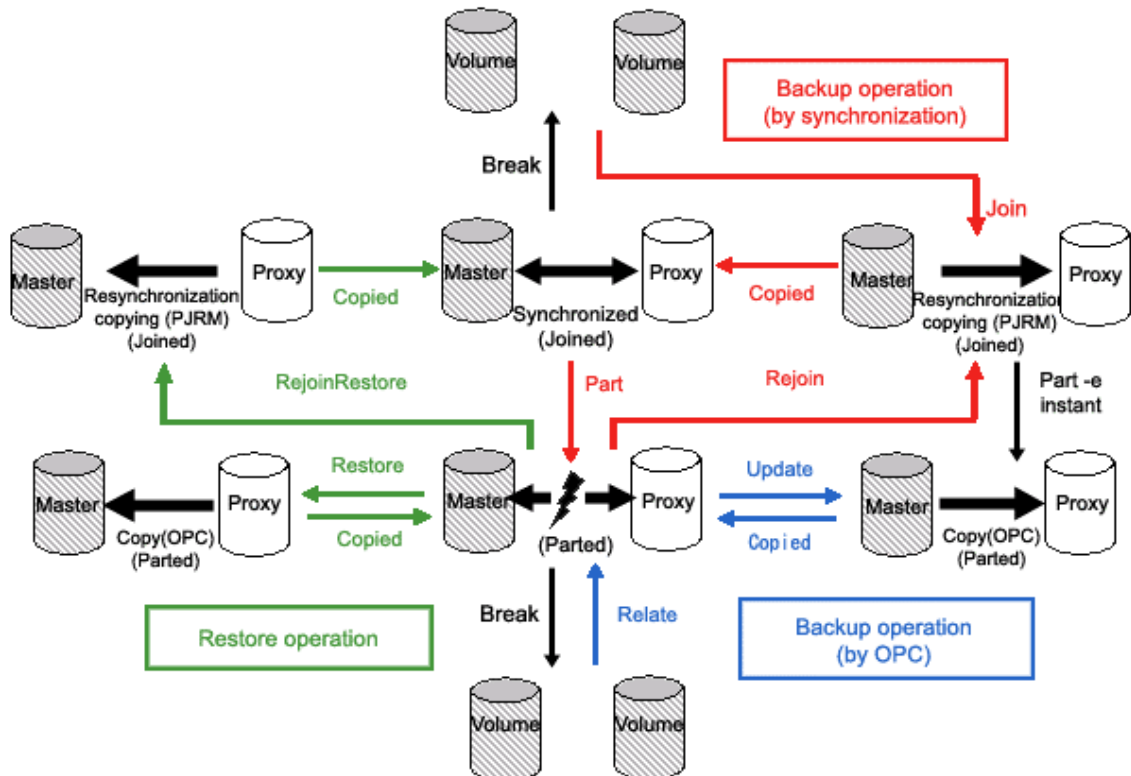


Figure: Swapping Master Volume and Proxy Volume

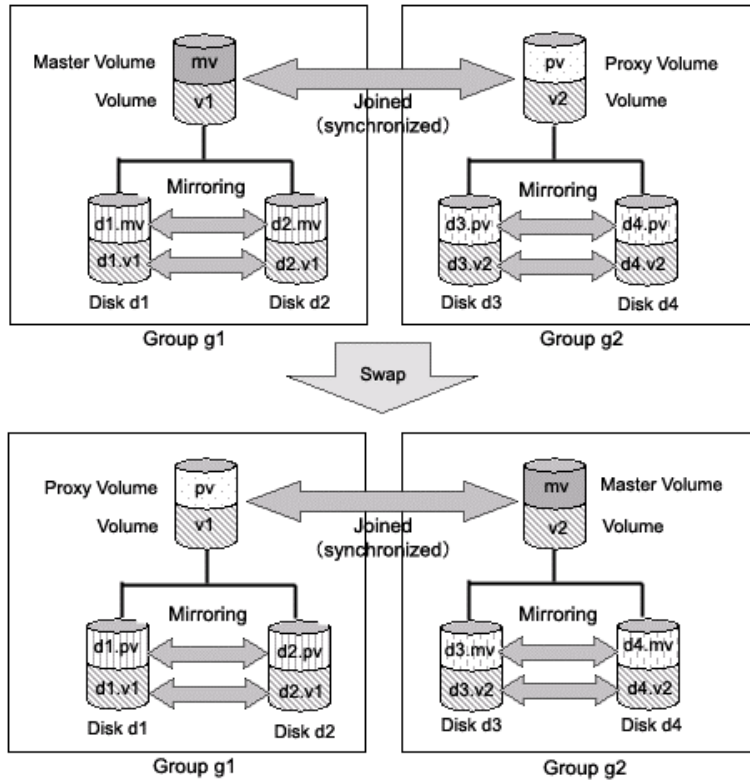
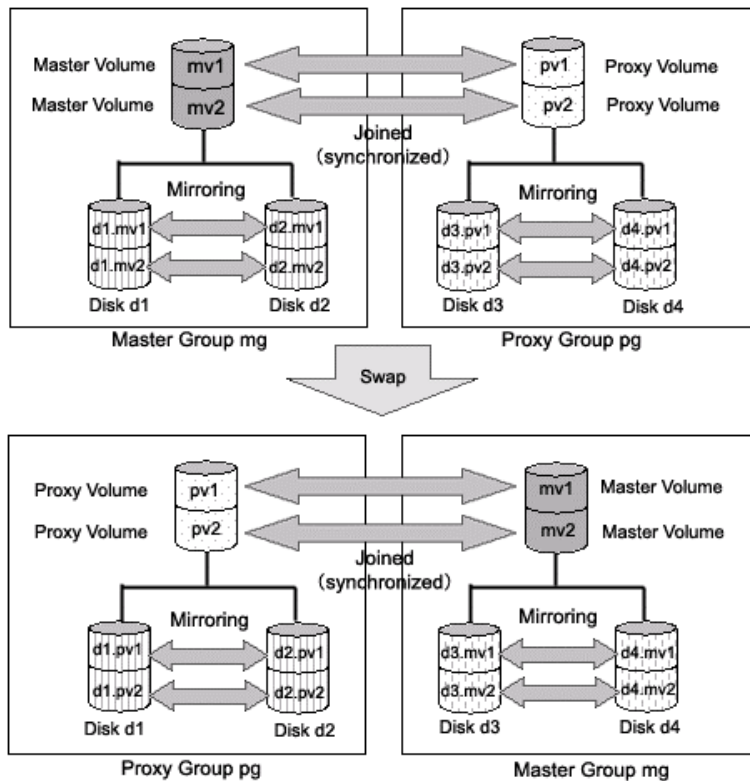


Figure: Swapping Master Group and Proxy Group



2.7.2 Shadow Object

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

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

On the server where GDS Snapshot is installed, physical disks that are not registered with disk classes in the same domain and that store SDX disk management data, can be managed as shadow objects and accessed. Physical disks conforming to 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.2.26 Shadow Volume."

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

2.7.2.1 Shadow Class

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

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

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

- Disks registered with a class (local class or shared class) of GDS 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 "5.3.1 Viewing Configurations and Monitoring Statuses."

Status

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

Reference

There are the following points of concern for shadow classes.

Rules

A.1.1 Object Name

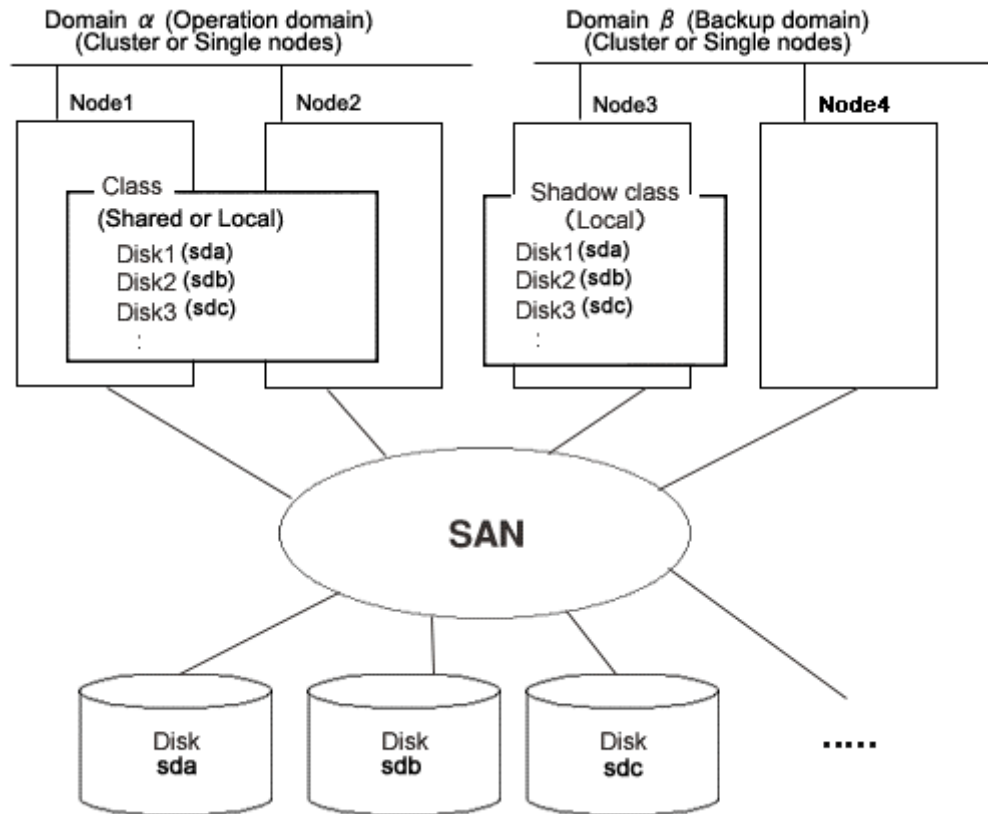
A.1.2 Number of Classes

Important Points

A.2.23 Operation in Cluster System

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


Figure: 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: 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.

2.7.2.2 Shadow Disk

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

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

Attributes

A shadow disk has the following attributes.

Name

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

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

Type

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

Mirror

The disk has been connected to a mirror group.

Stripe

The disk has been connected to a stripe group.

Concatenation

The disk has been connected to a concatenation group.

Single

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

Undefined

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

Operation

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

Create

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

Delete

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

Connect

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

Disconnect

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

Status Display

The shadow disk status can be displayed with the `sdxinfo` command. For details on GDS Management View, see "5.3.1 Viewing Configurations and Monitoring Statuses."

Status

A shadow disk can have the following status.

ENABLE

The shadow disk is in the normal status.

Reference

There are the following points of concern for shadow disks.

Rules

A.1.3 Number of Disks

Important Points

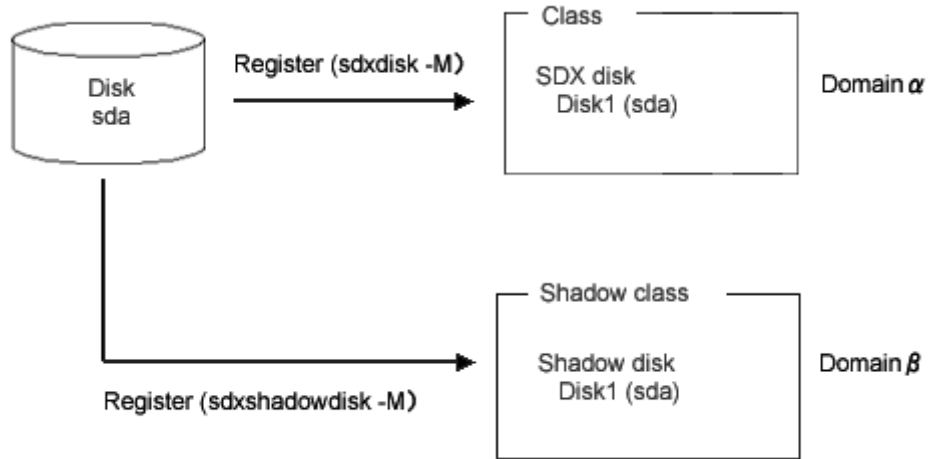
A.2.2 Restraining Access to Physical Special File

A.2.4 Initializing Disk

A.2.5 Disk Size

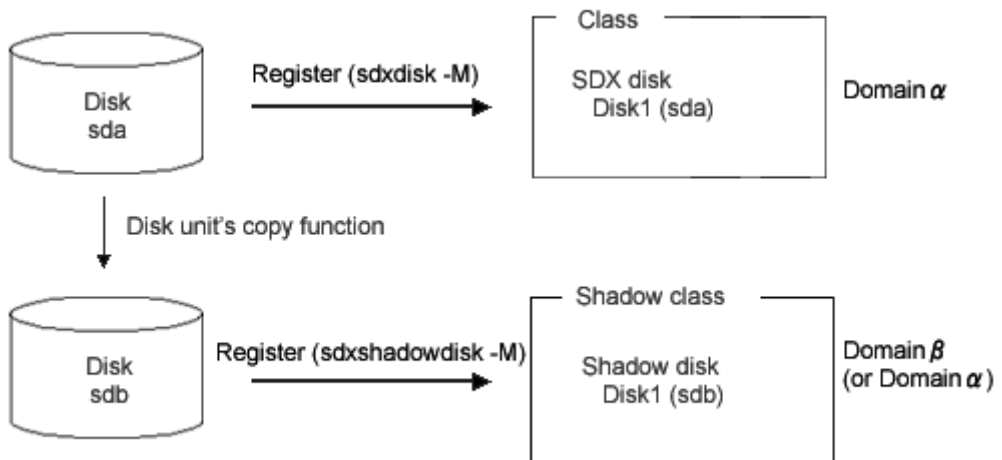
The physical disk that is registered with a class in a certain domain (domain alpha) with the `sdxdisk -M` command and accordingly is managed as an SDX disk can be registered with a shadow class in another domain (domain beta) with the `sdxshadowdisk -M` command and accordingly be managed as a shadow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

Figure: An SDX Disk and a Shadow Disk



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

Figure: 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.

2.7.2.3 Shadow Group

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

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

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

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

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

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

Attributes

A shadow group has the following attributes.

Name

This attribute identifies the shadow group within a shadow class.

Type

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

mirror

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

stripe

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

concat

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

Stripe Width

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

Operation

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

Create

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

Delete

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

Connect

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

Disconnect

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

Status Display

The shadow group status can be displayed with the `sdxinfo` command. For details on GDS Management View, see "5.3.1 Viewing Configurations and Monitoring Statuses."

Reference

There are the following points of concern for shadow groups.

Rules

- A.1.4 Number of Groups
- A.1.7 Creating Group Hierarchy

2.7.2.4 Shadow Volume

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

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

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

- Equal to the corresponding logical volumes in size. For volume sizes, check the `BLOCKS` field displayed with the `sdxinfo -V` command.
- Whose the 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.

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.2.26 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 partition table.

Operation

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

Create

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

Delete

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

Start

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

Stop

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

Status Display

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

For details on GDS Management View, see "5.3.1 Viewing Configurations and Monitoring Statuses."

Status

A shadow volume can have the following status.

ACTIVE

Data is accessible.

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

STOP

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

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

Reference

There are the following points of concern for shadow volumes.

Rules

A.1.5 Number of Volumes

Important Point

A.2.6 Volume Size


A.2.22 Volume Access Mode

A.2.26 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: A Logical Volume and a Shadow Volume

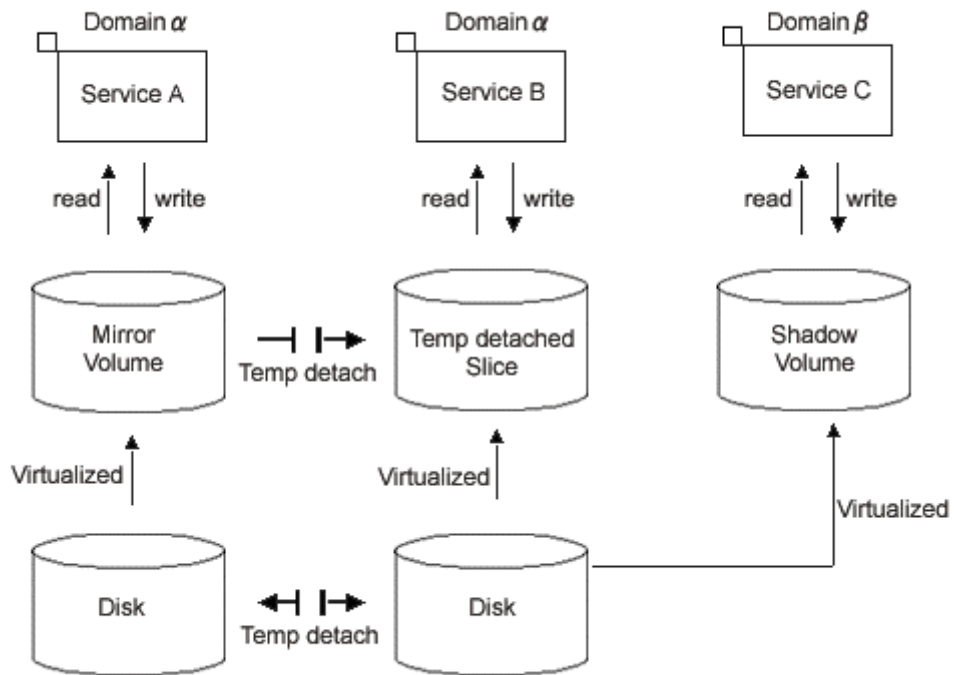
* Service A and service B cannot be run simultaneously



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: Mirror Slice Temporary Detachment and a Shadow Volume

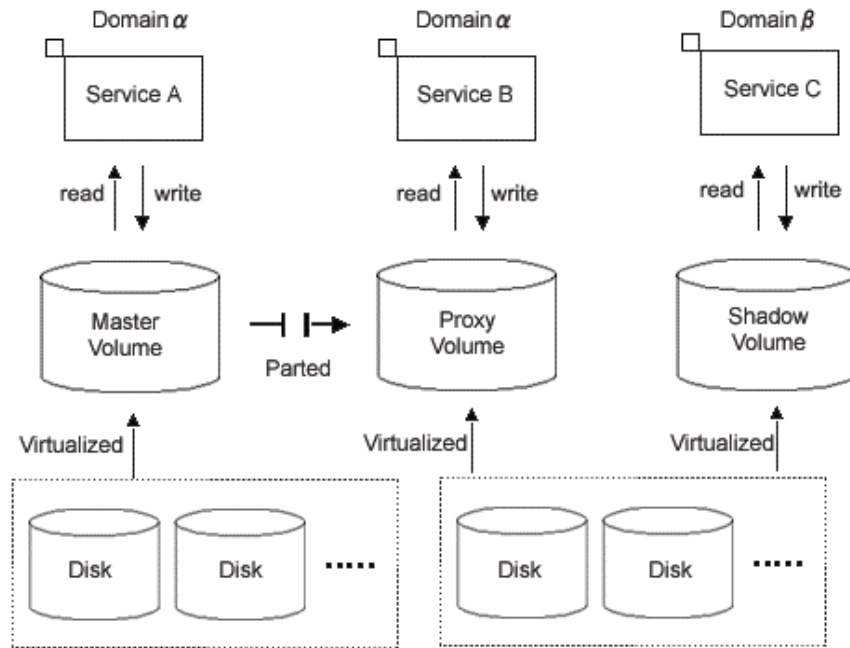
* Service B and service C cannot be run simultaneously



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.

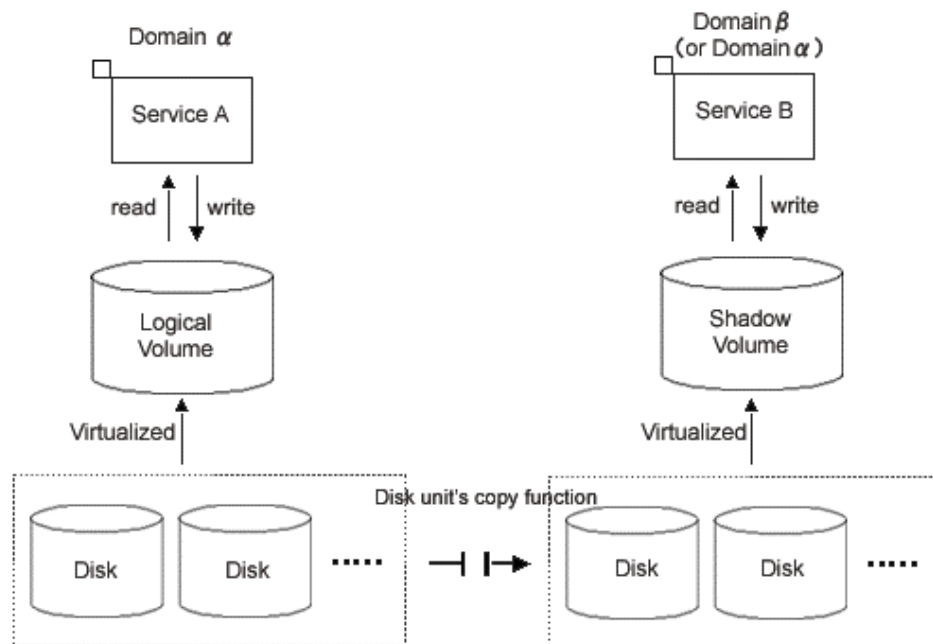
Figure: A Proxy Volume and a Shadow Volume

* Service B and service C cannot be run simultaneously



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

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



2.7.2.5 Shadow Slice

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

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

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

Attributes

A shadow slice has the following attributes.

Name

This attribute identifies the shadow slice within a shadow class.

Operation

Status Display

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

For details on GDS Management View, see "5.3.1 Viewing Configurations and Monitoring Statuses."

Status

ACTIVE

Data is accessible.

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

STOP

Data is inaccessible.

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

INVALID

Data is inaccessible due to an I/O error.

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

For details, see "F.1.1 Slice Status Abnormality."

Chapter 3 Starting and Exiting GDS Management View

GDS Management View manages and monitors objects by using the Web browser. This chapter explains how to start and exit the GDS Management View.

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

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

3.1 Preparation for Starting GDS Management View

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

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

3.1.1 Deciding the User Group

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

3.1.1.1 User Group Types

There are two types of user groups.

wvroot

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

sdxroot

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

3.1.1.2 Creating User Groups

After installing the software, only user group wvroot will be automatically created. The other user group sdxroot, explained in "3.1.1.1 User Group Types" must be created as necessary.

User group can be created using the following command.

```
# groupadd sdxroot
```

3.1.1.3 Registering to a User Group

You can register a user name to a user group with a command. Type an appropriate user group such as `sdxroot`, and execute the following command.

- **When registering an existing user to a user group**

- When registering a group to "Primary Group"

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

- When registering a group to "Secondary Group"

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

- **When registering a new user to a user group**

- When registering a group to "Primary Group"

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

- When registering a group to "Secondary Group"

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



Note

Registering a User to the `wvroot` Group

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

3.1.2 Setting up the Client Environment

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

3.1.3 Setting up the Web Environment

In order to use GDS Management View, you must define the environment for the client and the Web browser. For details on setting the Web browser, see "Web-Based Admin View Operation Guide."

3.2 Starting the GDS Management View

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

3.2.1 Starting Web-Based Admin View Operation Menu

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

Procedures

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

When using Java Plug-in.

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

host name

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

port number

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



See

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



Note When the Web-Based Admin View does not start

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

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

Figure: User Name Input Screen

The screenshot shows a Java applet window titled "Web-Based Admin View" with a standard Windows-style title bar (minimize, maximize, close buttons). The window contains two text input fields: "User name:" and "Password:". Below these fields is a single "OK" button. The status bar at the bottom of the window indicates it is an "Unsigned Java Applet Window".

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

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

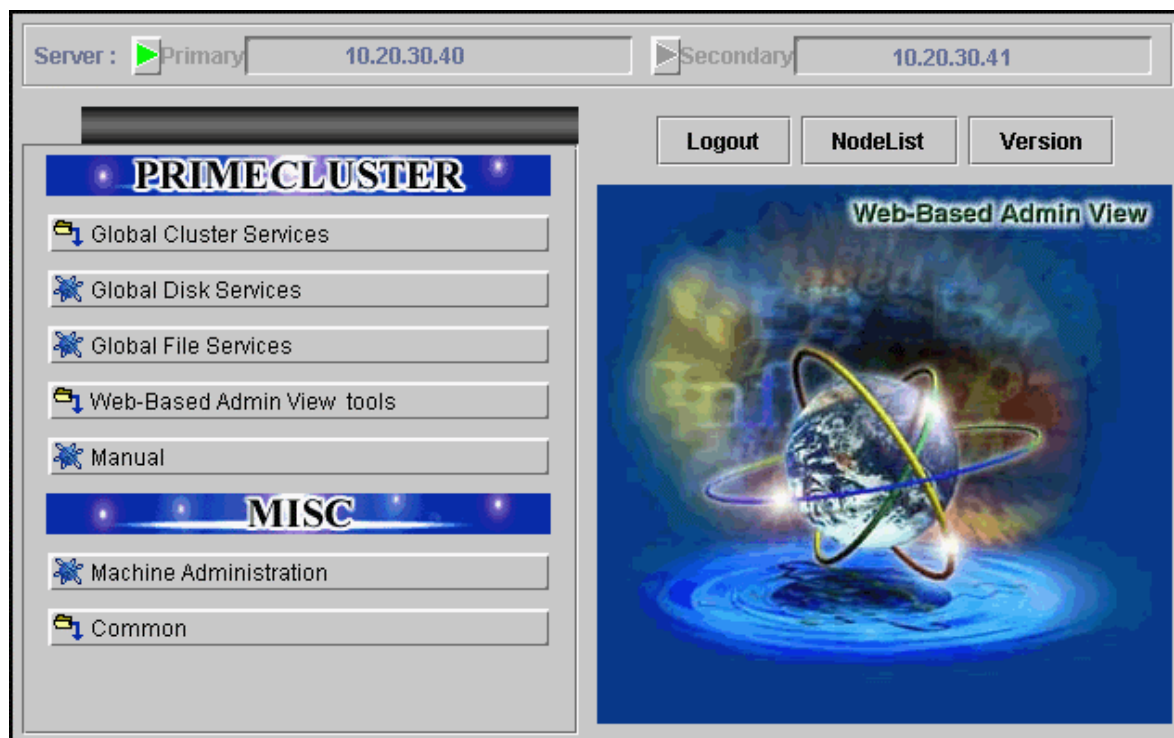
3.2.2 Web-Based Admin View Operation Menu Functions

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

Menu	Outline
Global Disk Services	Start the GDS Management View.

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

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



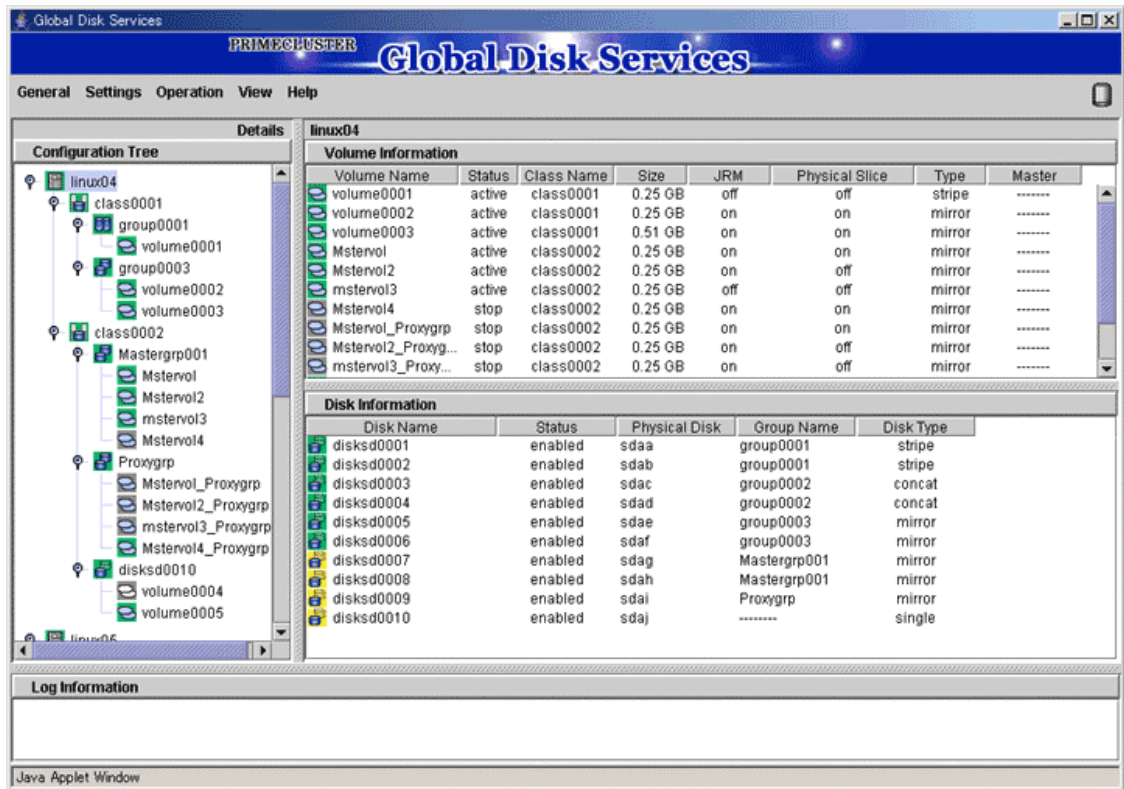
Note The state of Web-Based Admin View Operation Menu

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

3.2.3 Starting GDS Management View

Click the GDS management icon on the Web-Based Admin View Operation menu to start the GDS Management screen (hereinafter main screen).
 From the main screen, you can perform GDS object configuration, such as a class or a volume, status confirmation, and disk swap.
 For details, see "Chapter 5 Operation"


Figure: GDS Management Main Screen



3.3 Exiting GDS Management View

How to exit the GDS Management View is described below.
On the [General] menu, click [Exit]. The following message appear

Figure: Exit Message



Click <Yes>, and Web-Based Admin View (top menu) appears.
Click <No>, and you will return to the main screen.

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

- Click <Log Out> from one of the following menus; top menu, GDS Management Operation menu, or Common Operation menu.
- Log In screen appears. Exit the browser, or click Back on the browser and exit GUI.



Note **When the login screen remains displayed**

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

3.4 Changing the Web-Based Admin View Settings

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

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

Chapter 4 Management View Screen Elements

This chapter explains the screen elements of GDS Management View.
Screen elements refer to the following items:

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

4.1 Screen Configuration

Main Screen

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

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

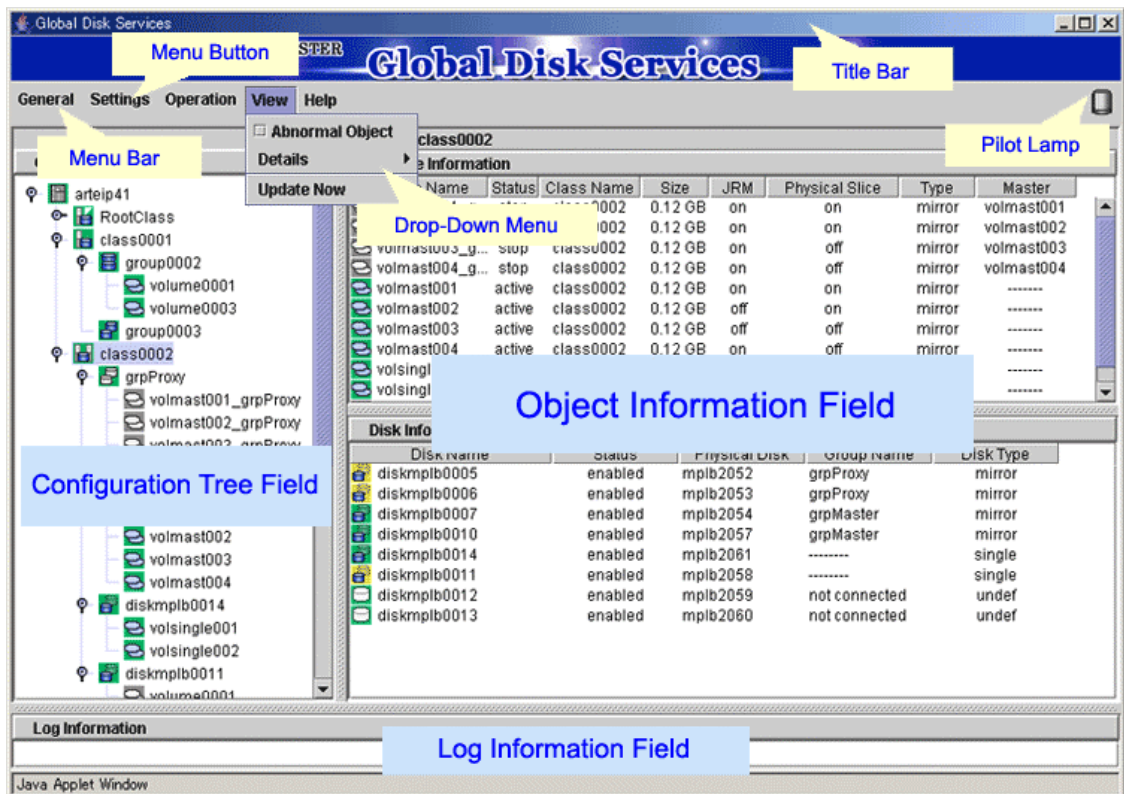
You can also configure and view the statuses of GDS Snapshot proxy objects and shadow objects and also operate those proxy objects.

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

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

The screen configuration of the main screen is shown below.

Figure: GDS Management Screen (Main Screen)



Configuration Tree Field

Objects managed with GDS are displayed in a tree-structured directory system. Each object has an icon depicting the type and the status of the object.

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

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

Object Information Field

Detailed information of objects is displayed in table format.

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

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

Log Information Field

Displays error message output by the GDS daemon program.

Title Bar

Displays screen title (Global Disk Services).

Menu Bar

Displays the menu buttons.

Menu Button

Allows you to control the objects selected on screen.

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

Drop-Down Menu






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

For details on drop-down menu, see "4.2 Menu Configuration and Functions."

Pilot Lamp

Shows the status of the monitored objects.

The lamp will indicate the following status.

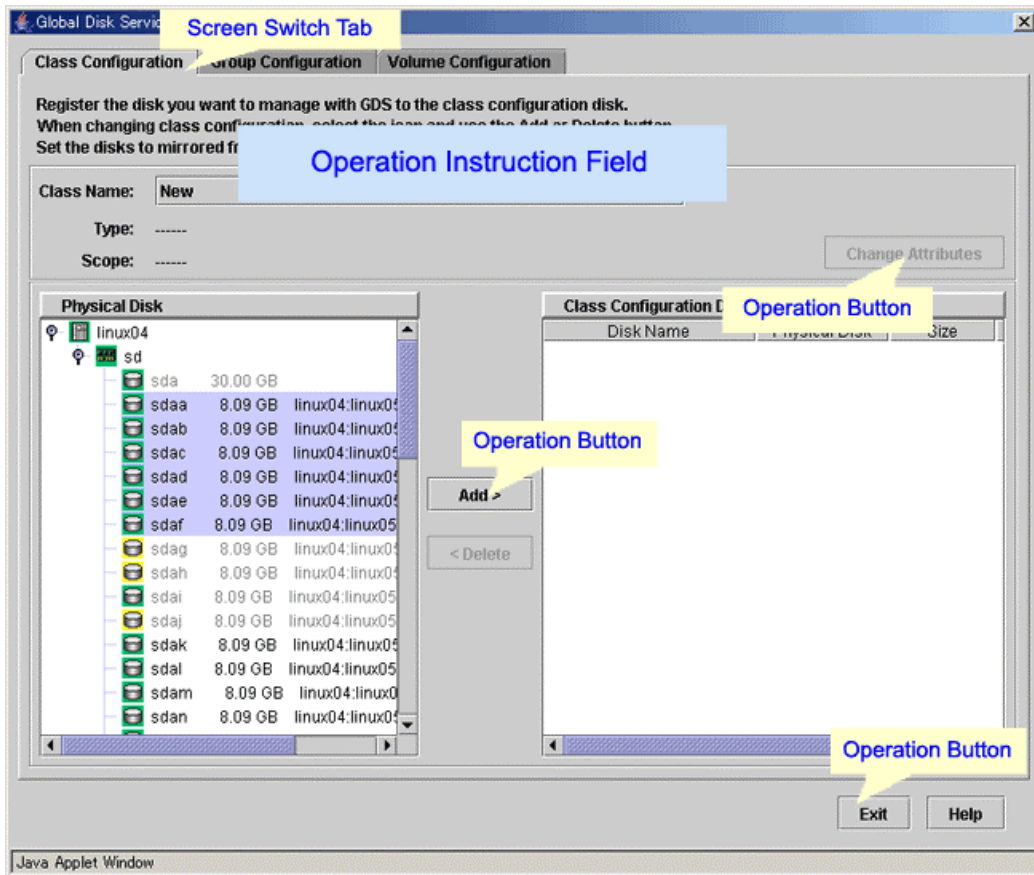
Pilot Lamp	Status	Meaning
 (Gray lit up)	Normal	Normal status.
 (Red blinking)	Abnormal	Critical abnormality (e.g. closed class, invalid volume)
 (Red lit up)	Abnormal	When red blinking warning lamp is single-clicked.
 (Yellow blinking)	Degradation	<ul style="list-style-type: none">Volume operating at 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: GDS Configuration Settings Screen



For information on how to set each configuration, see "5.2.2 Operating from the Settings Menu."

4.2 Menu Configuration and Functions

Each menu button has a drop-down menu allowing you to operate the selected object on screen. This section explains the menu configuration and functions. The operations for shadow objects available with GDS Snapshot are not supported.

4.2.1 General



Change Monitoring Intervals

Sets the monitoring interval (min) of objects.

Figure: [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: [General]: [Exit] Screen



4.2.2 Settings

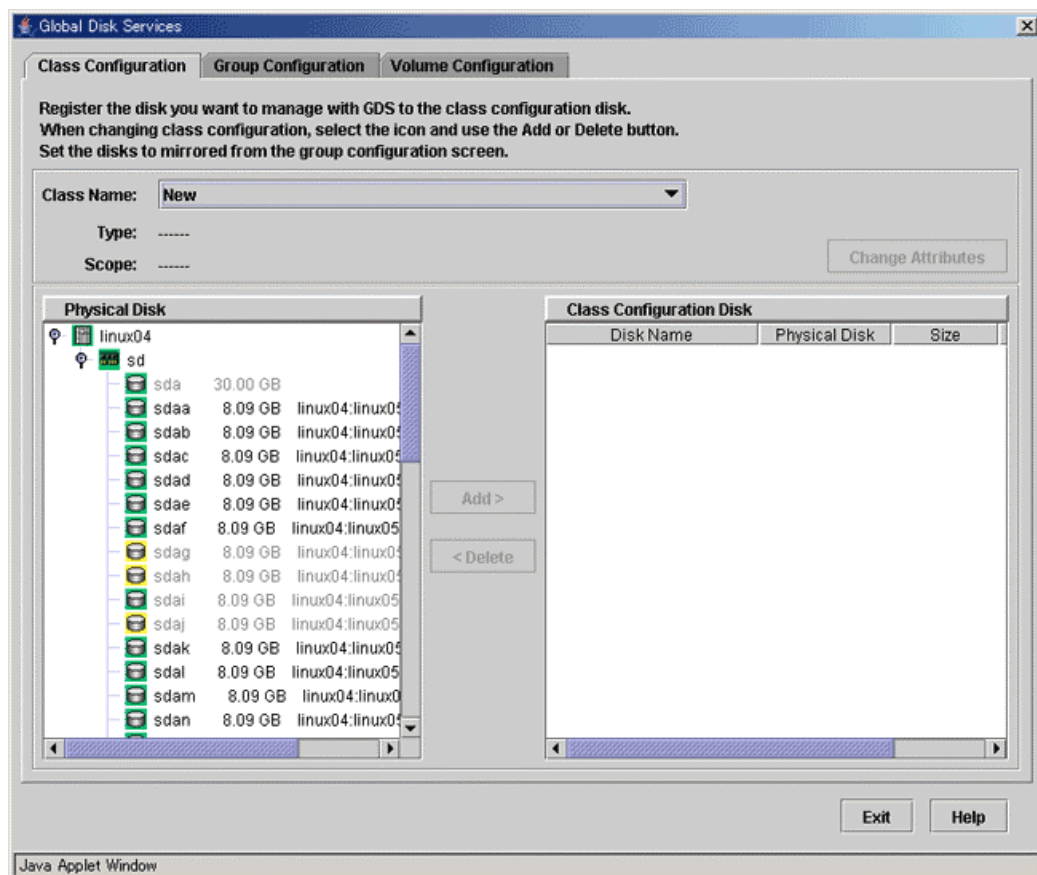


Class Configuration

Sets the class configuration.

For details, see "5.2.2 Operating from the Settings Menu."

Figure: [Settings]: [Class Configuration] Screen

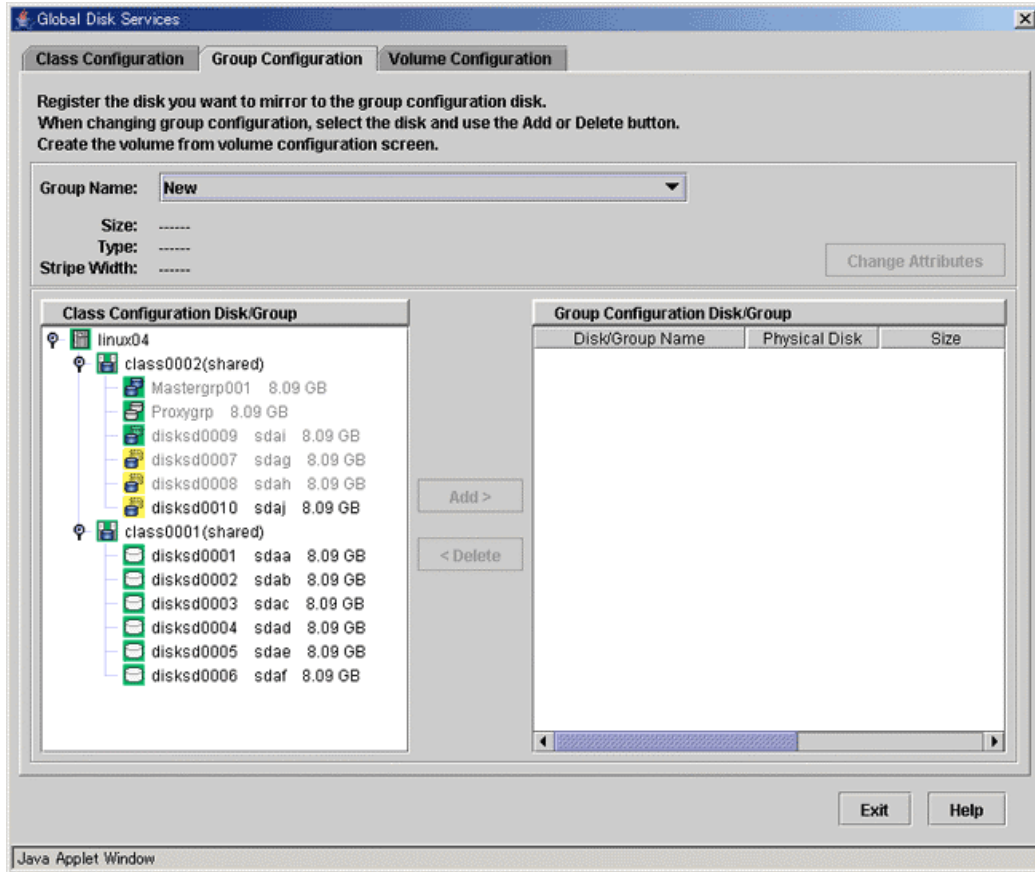


Group Configuration

Sets the group configuration.

For details, see "5.2.2 Operating from the Settings Menu."

Figure: [Settings]: [Group Configuration] Screen

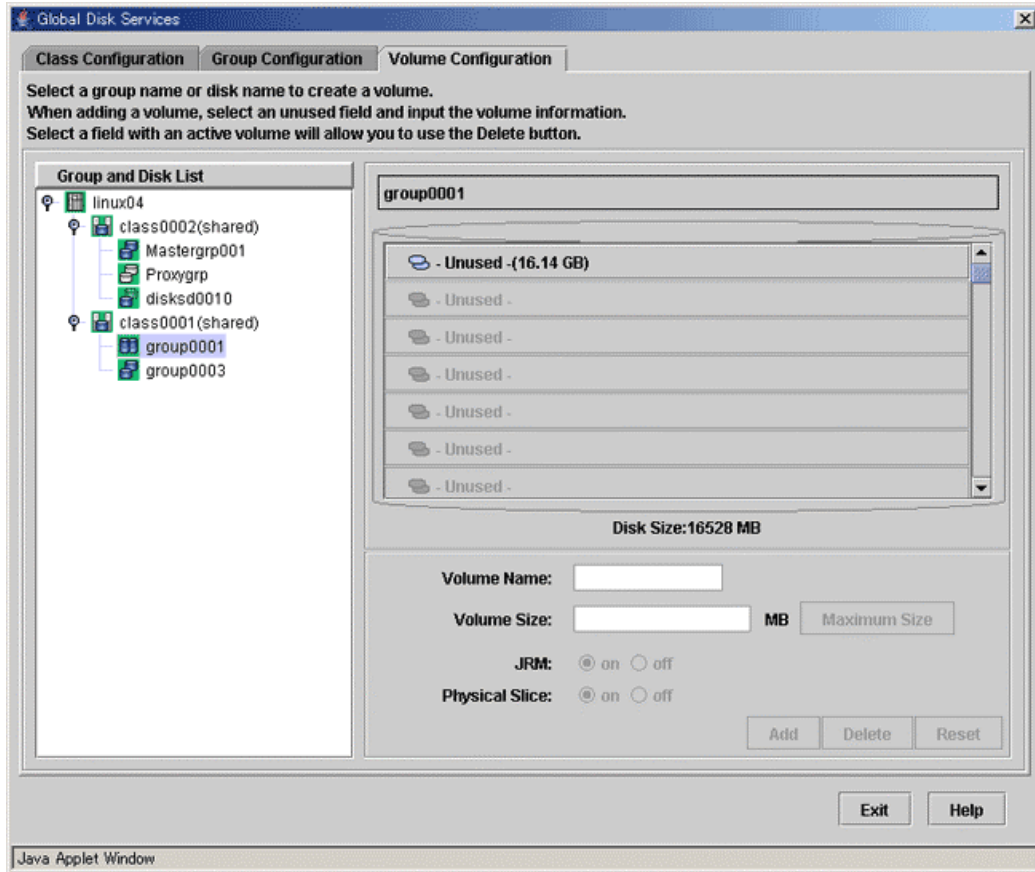


Volume Configuration

Sets the volume configuration.

For details, see "5.2.2 Operating from the Settings Menu."

Figure: [Settings]: [Volume Configuration] Screen

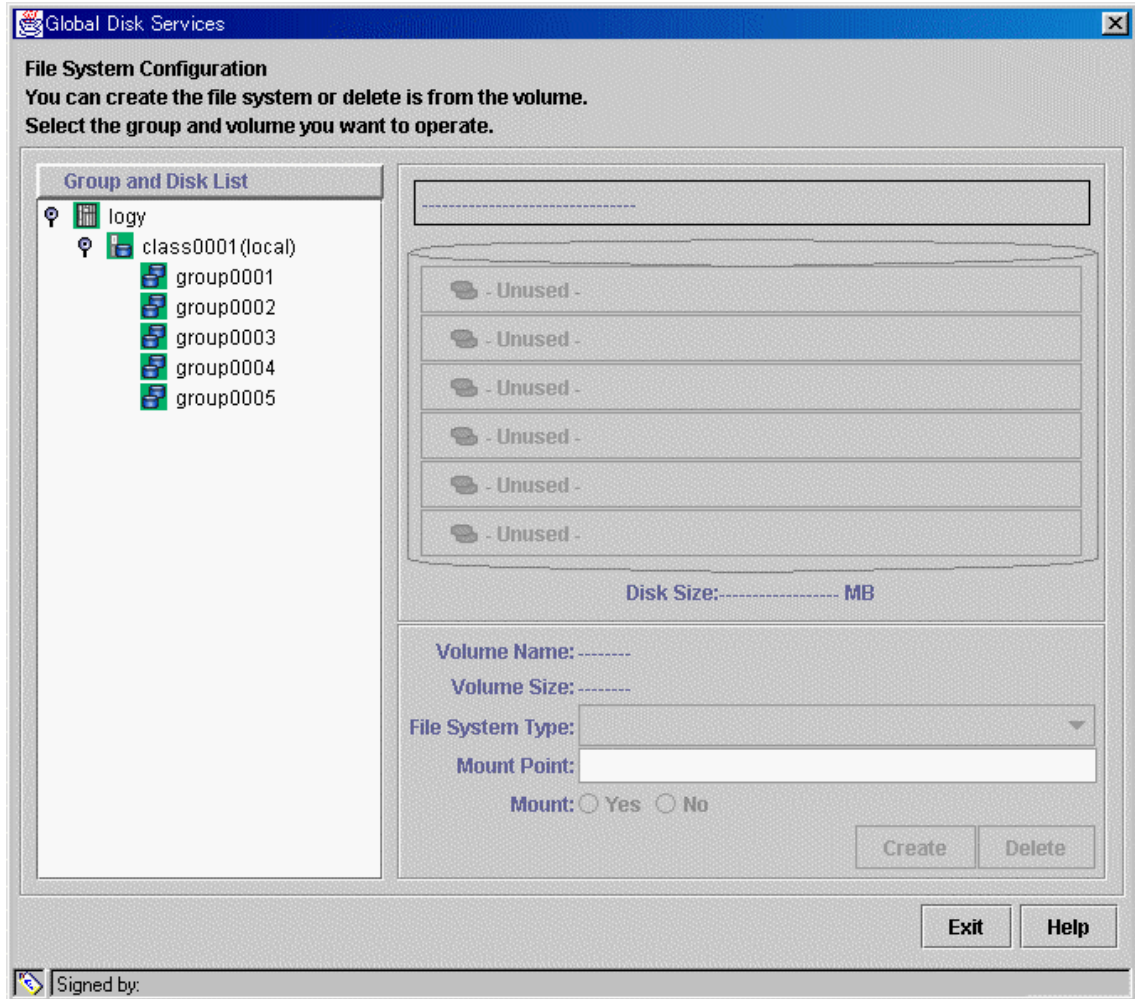


File System Configuration

Sets the file system configuration.

For details, see "5.2.3 File System Configuration."

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

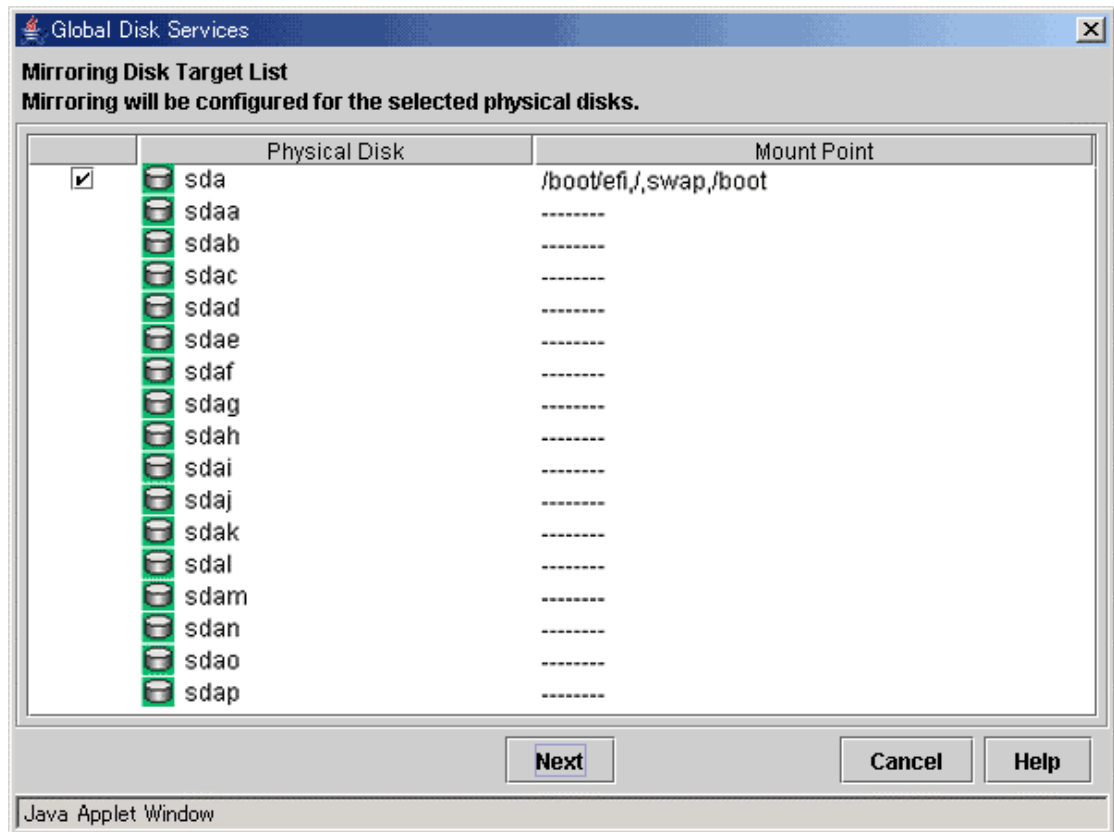


System Disk Settings [Itanium]

Mirrors the system disk.

For details, see "5.2.1 System Disk Settings [Itanium]."

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

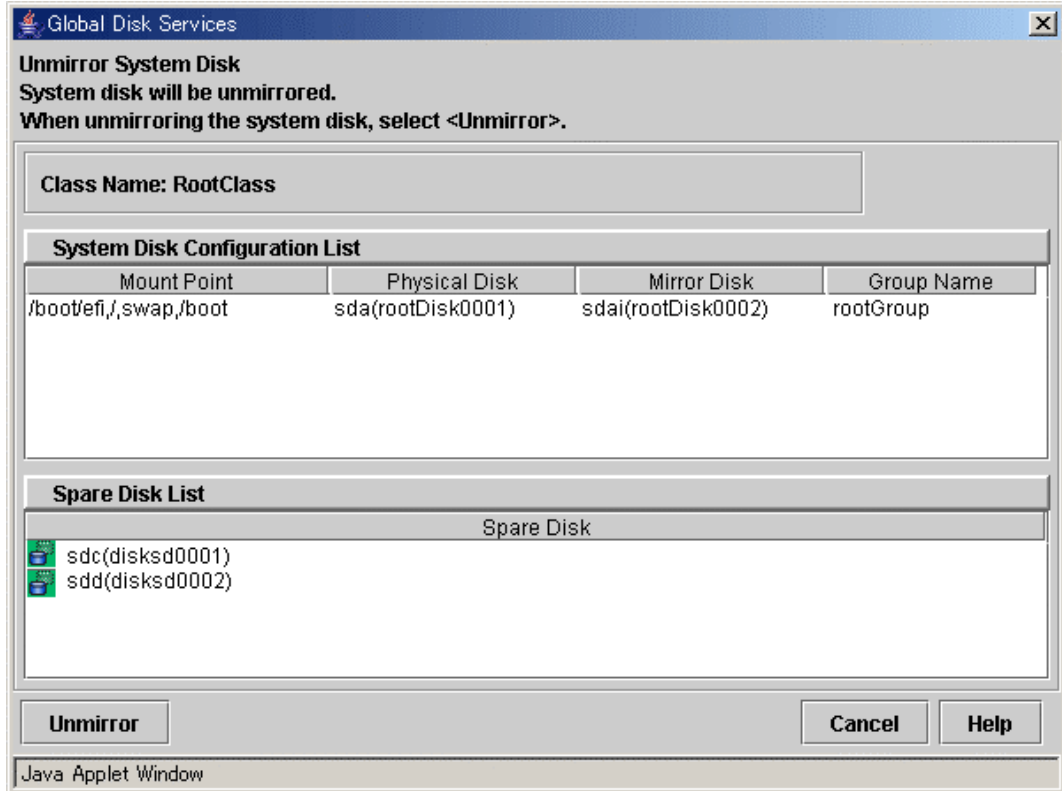


Unmirror System Disk [Itanium]

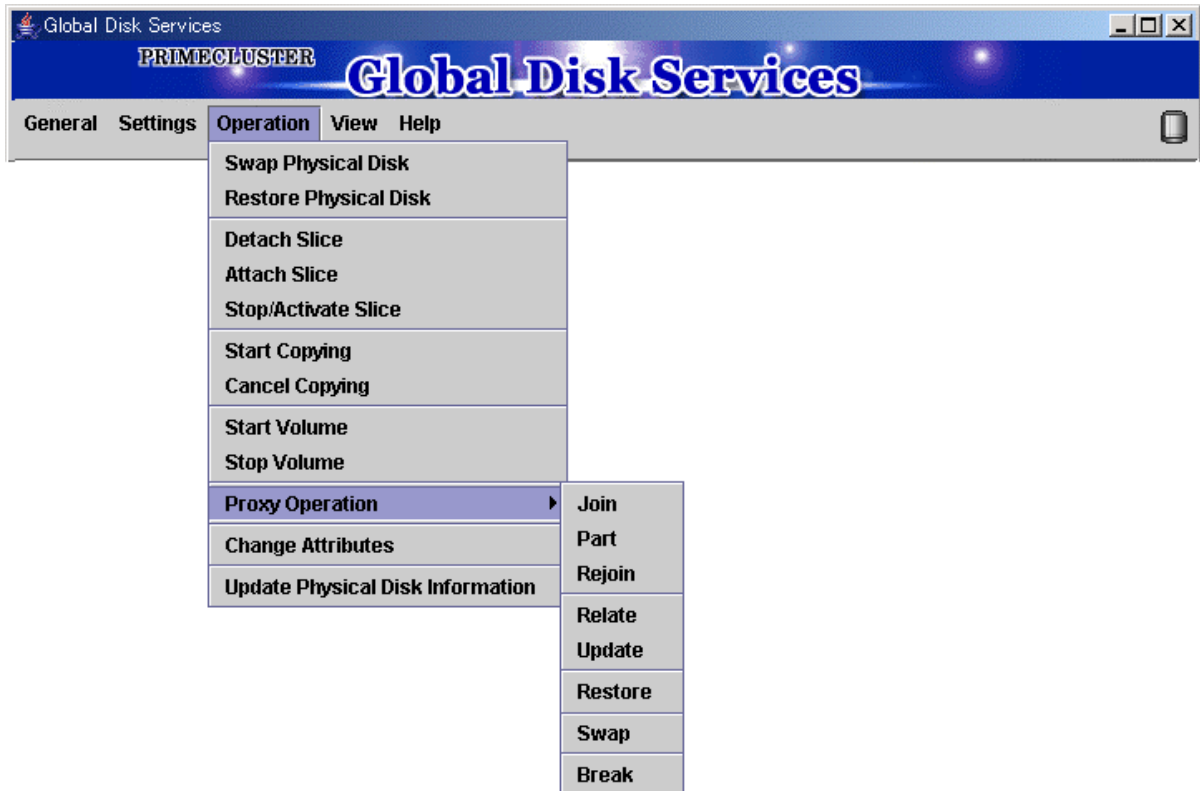
Unmirrors the system disk.

For details, see "5.5.5 Unmirroring the System Disk [Itanium]."

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



4.2.3 Operation



Swap Physical Disk

Places the physical disk off-line when swapping physical disks.
For details, see "5.3.4 Disk Swap."

Restore Physical Disk

Places the swapped physical disk on-line for restoration, after swapping the physical disks.
For details, see "5.3.4 Disk Swap."

Detach Slice

Detaches one of the slices from Mirror volume to prepare for backup.
The detached slice will become accessible as a separate logical device.
For details, see "5.3.2.1 Backup (by Slice Detachment)."

Attach Slice

Slice detached by [Detach Slice] is attached to Mirror volume again.
For details, see "5.3.2.1 Backup (by Slice Detachment)."

Stop/Activate Slice

Stop Slice

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

Activate Slice

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

For details, see "5.3.2.1 Backup (by Slice Detachment)."

Start Copying

A slice with "invalid" or "copy-stop" status as a result of [Cancel Copying] operation will be attached to a mirror volume, and synchronization copying will be performed.
For details, see "5.3.6 Copying Operation."

Cancel Copying

Execution of copying will be stopped to avoid effects caused by accessing the disk in the process of synchronization copying.
For details, see "5.3.6 Copying Operation."

Start Volume

Starts the stopped volume.

Stop Volume

Stops the volume.

Proxy Operation

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

Join

Relates proxy objects to master objects and synchronizes them as preparation for snapshot creation by synchronization, or online disk migration.
For details, see "5.2.4.1 Join."

Part

Parts joined proxies from masters temporarily to make them accessible as logical devices that are different from the masters for snapshot creation by synchronization. The parted proxies can be used as snapshots (replicas) of the masters at the moment.
For details, see "5.3.2.2 Backup (by Synchronization)."

Rejoin

Joins parted proxies to masters again and synchronizes them as preparation for snapshot re-creation by synchronization.
For details, see "5.3.2.2 Backup (by Synchronization)."

Relate

Relates proxy objects to master objects and parts them as preparation for snapshot creation by OPC.
For details, see "5.2.4.2 Relate."

Update

Copies (overwrites) data from masters to parted proxies for snapshot creation by OPC. The updated proxies can be used as snapshots (replicas) of the masters at the moment.
For details, see "5.3.2.3 Backup (by OPC)."

Restore

Copies (overwrites) data from parted proxies to masters for restoration of damaged master data. The masters are recovered by the proxy data at the moment.
For details, see "5.3.3 Restore."

Swap

Swaps slices of synchronized masters and proxies for online disk migration.
For details, see "5.3.5 Disk Migration."

Break

Breaks the relationships between masters and proxies and makes them unrelated again.
For details, see "5.5.6 Breaking a Proxy."

Change Attributes

Changes the attributes of the selected object.
For details, see "5.4 Changes."

Update Physical Disk Information

You can update the disk information without rebooting the system. This feature is useful in the situations given below.

- When physical disk size is not displayed properly.
- When you turn on the disk array or the disk unit after booting the system.
- When the disk has become inaccessible while operating the system.

4.2.4 View



Abnormal Object

Only displays objects with abnormalities; not all objects.

Details

Changes displayed contents in the Object Information Field.
By default, [SDX Object] is selected.

SDX Object

Displays information of volumes, disks and slices.

Proxy Object

Displays information of proxy volumes, proxy groups and slices.

For details on the Objects Information Field, see "4.4 Object Information."

Update Now

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

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

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

4.2.5 Help



Help

Displays Help information.

4.3 Icon Types and Object Status

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






Information


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

The status and the icons of objects are shown below.




1. Nodes

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




2. Adapter

Icon	Status	Meaning
 (Green)	-	-



3. Classes (local) [Linux2.6]

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) [Itanium]

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




6. Groups (mirror)

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




7. Groups (stripe)

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




8. Groups (concat)

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







9. Groups (switch)

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






10. Physical disks

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







11. Disks connected to a group

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






12. Single disks

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







13. Spare disks

Icon	Status	Meaning
 (Green)	enabled	Operation enabled
 (Yellow)	enabled	Some volumes or slices with status other than active
 (Red)	enabled	I/O error on the disk
 (Red)	disabled	Operation disabled
 (Red)	close	Closure of the class to which the spare disk belongs
 (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. You should pinpoint the 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 "F.1.12 Disk Unit Error."

4.4 Object Information

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

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

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

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

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

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

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

The Object Information Field displays the following information.

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

Chapter 5 Operation

This chapter describes the GDS operations from GDS Management View.

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

For details on operations supported in GDS Management View, see "5.1.9 Operations from GDS Management View."

5.1 Operation Outline

This section explains the GDS settings and operation management.



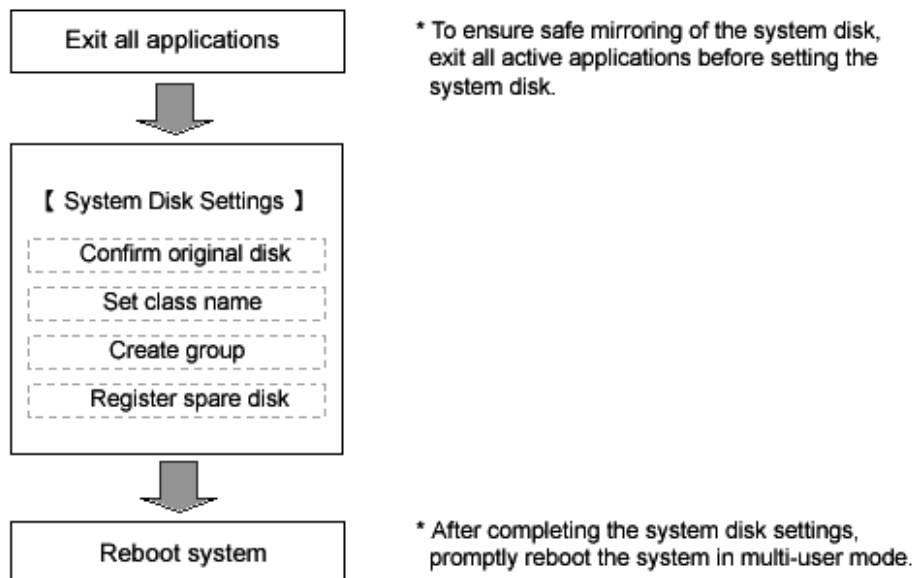
Note In PRIMECLUSTER systems

Before defining the configuration of GDS objects such as classes and volumes, conduct registration of shared disk devices in the PRIMECLUSTER resource database.
For details about resource registration, see "Appendix H Shared Disk Unit Resource Registration."

5.1.1 System Disk Settings [Itanium]

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

Figure: System Disk Settings Operation



See

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

5.1.2 Configuration Settings

For configuring objects other than system disks, use the relevant menus for configuration.

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

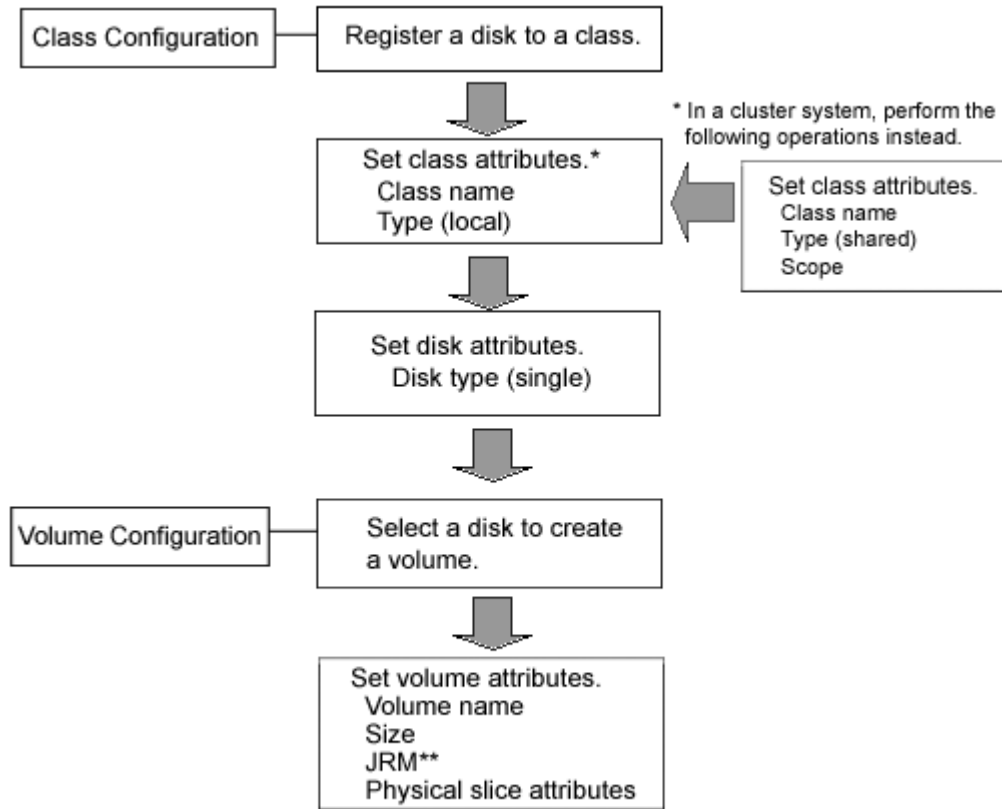
- Single volume configuration settings
- Other volume (mirror volume, stripe volume, volume in a concatenation group) configuration settings

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

5.1.2.1 Single Volume Configuration Settings

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

Figure: Single Volume Configuration Settings Operation



** Just Resynchronization Mechanism



See

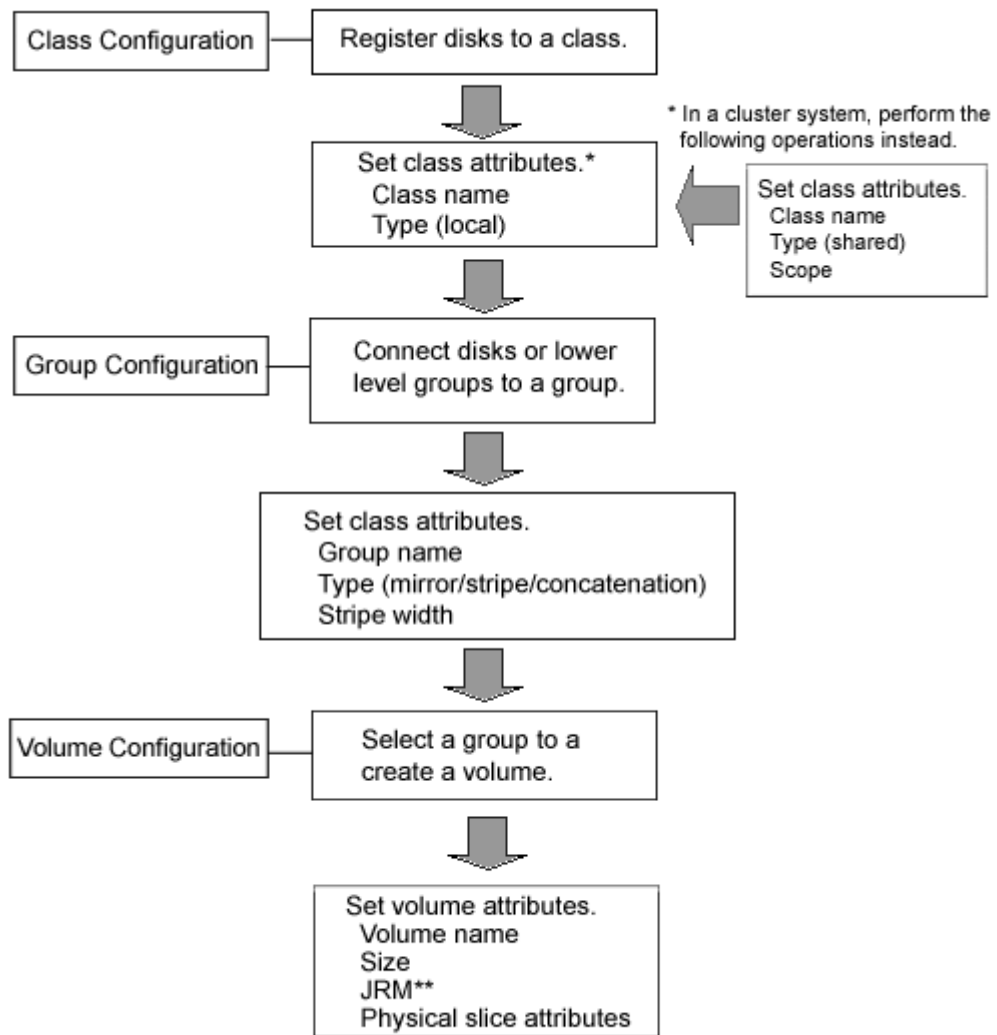
For details on the operation methods, see "5.2.2 Operating from the Settings Menu."

5.1.2.2 Other Volume Configuration Settings

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

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

Figure: Volume Configuration Settings Operation



** Just Resynchronization Mechanism



See

For details on the operation methods, see "5.2.2 Operating from the Settings Menu."

5.1.3 Backup

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

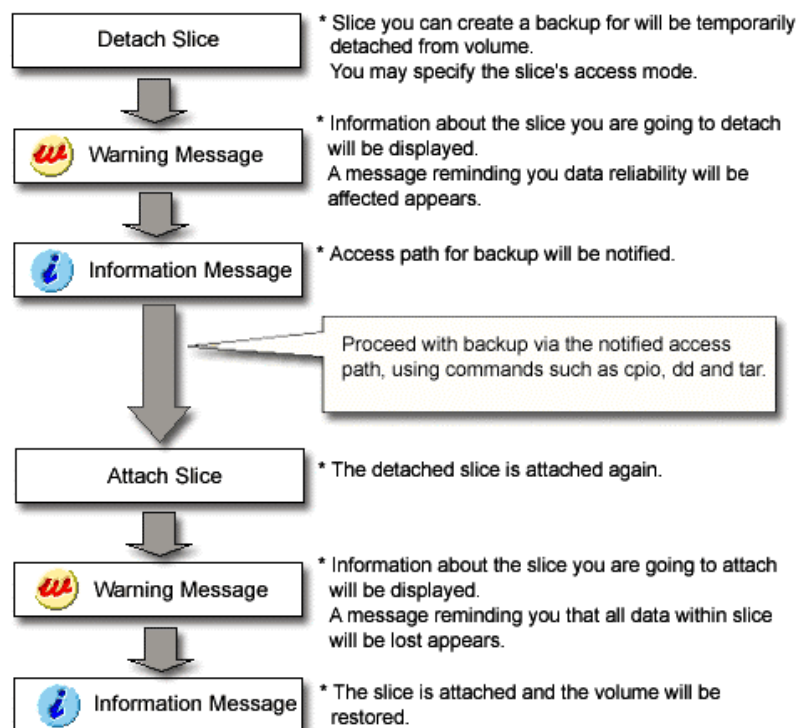
The following three backup methods are available.

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

5.1.3.1 Backup (by Slice Detachment)

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

Figure: Backup Operation (by Slice Detachment)



Note Use Conditions on Snapshot by Slice Detachment

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



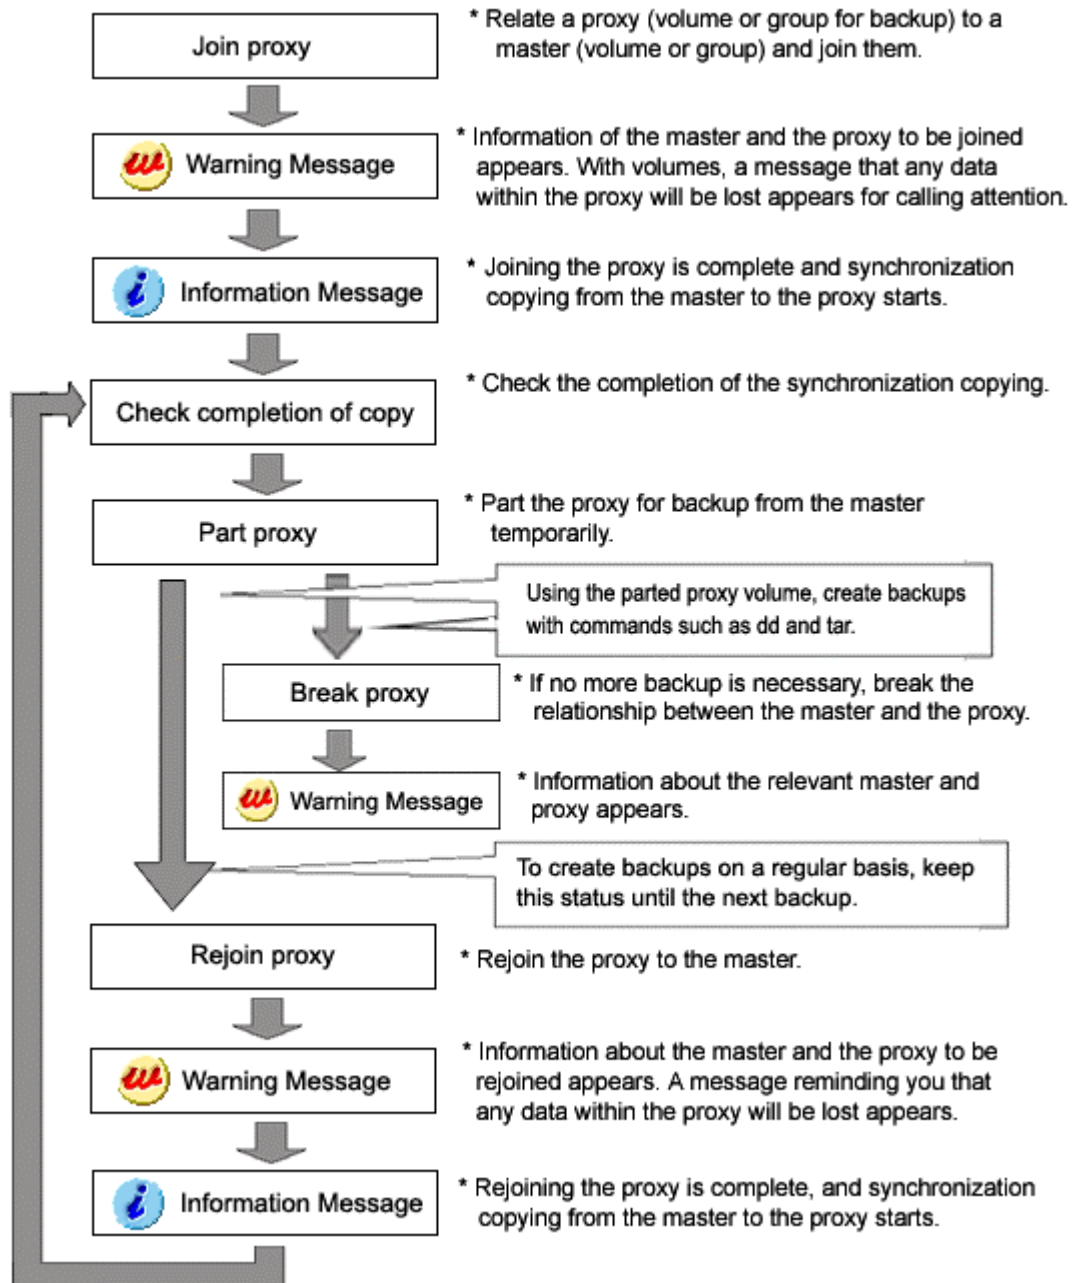
See

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

5.1.3.2 Backup (by Synchronization)

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

Figure: Backup Operation (by Synchronization)



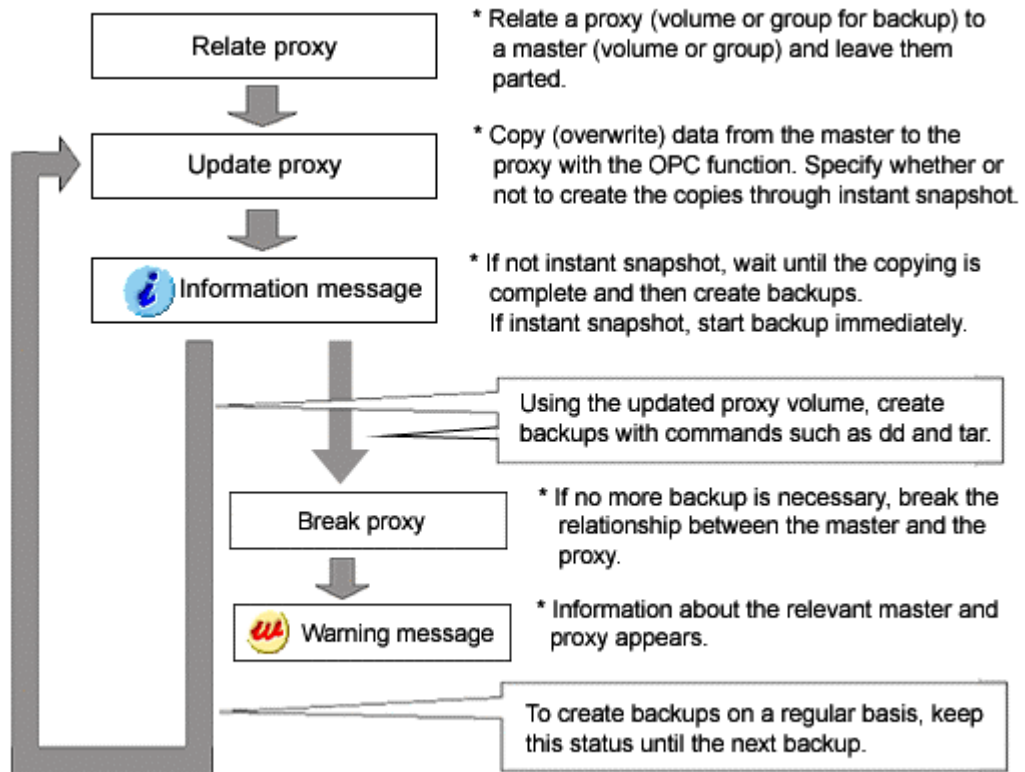
See

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

5.1.3.3 Backup (by OPC)

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

Figure: Backup Operation (by OPC)



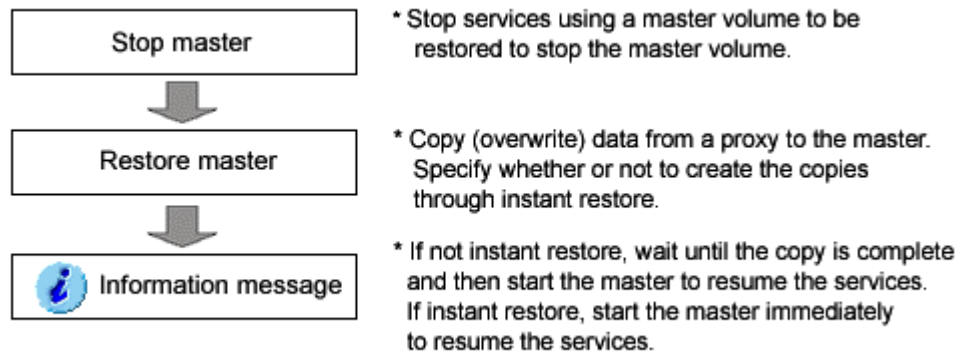
See

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

5.1.4 Restore

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

Figure: Restore Operation



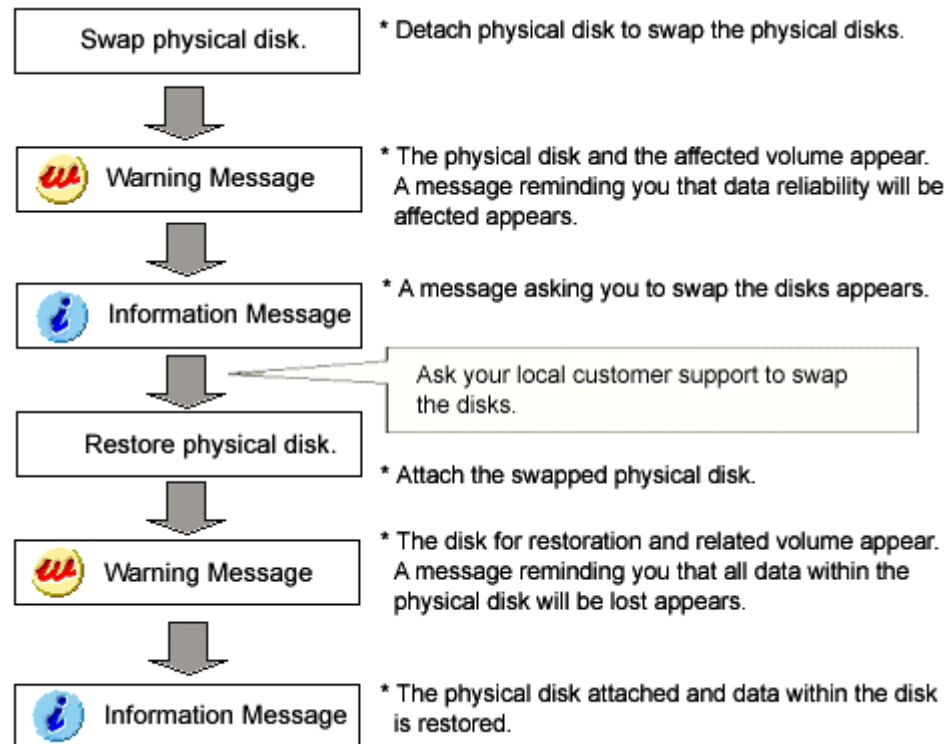
See

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

5.1.5 Disk Swap

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

Figure: Disk Swap and Restoration Operation




See

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

5.1.6 Disk Migration

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

Figure: Disk Migration Operation



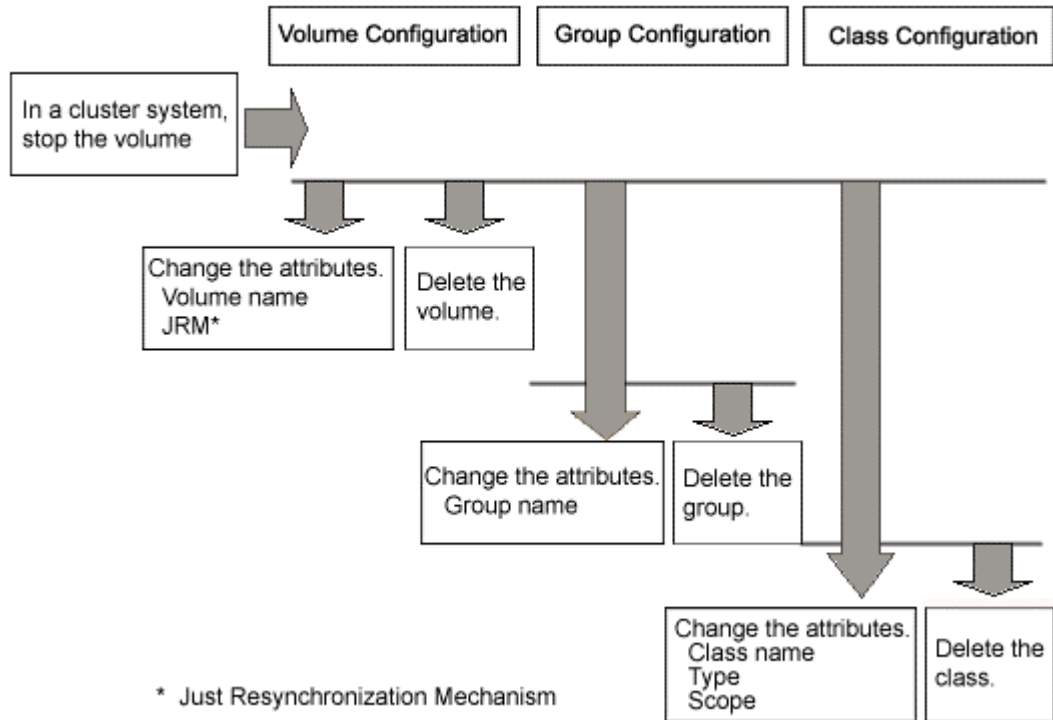
See

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

5.1.7 Configuration Change

Procedures for changing or removing the configuration settings is shown below.

Figure: Configuration Changing Operation




See

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

5.1.8 Unmirroring the System Disk [Itanium]

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

Figure: Unmirroring System Disk Operation



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



[See](#)

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

5.1.9 Operations from GDS Management View

GDS Management View supports the following operations.

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

Class operation

Operation	Command Operation	GDS Management View Operation
Create	# sdxdisk -M	For details, see "5.2.2.1 Class Configuration" in "5.2.2 Operating from the Settings Menu."
Remove	# sdxdisk -R or # sdxclass -R	For details, see "5.5.4 Removing a Class" in "5.5 Removals"
Change attribute (Class name)	# sdxattr -C -a name= <i>classname</i>	For details, see "5.2.2.1 Class Configuration" in "5.4 Changes."
Change attribute (Type)	# sdxattr -C -a type={local shared}	For details, see "5.4.1 Class Configuration" in "5.4 Changes."
Change attribute (Scope)	# sdxattr -C -a scope= <i>node: node:...</i>	For details, see "5.4.1 Class Configuration" in "5.4 Changes."
Change attribute (Hot spare)	# sdxattr -C -a hs={on off}	Unsupported
Change attribute (Hot spare mode)	# sdxattr -C -a hsmode={ebox bybox}	Unsupported
Restore	# sdxfix -V	Unsupported

Group operation

Operation	Command Operation	GDS Management View Operation
Connect	# sdxdisk -C or # sdxgroup -C	For details, see "5.2.2.3 Group Configuration" in "5.2.2 Operating from the Settings Menu."
Remove	# sdxdisk -D or # sdxgroup -D or # sdxgroup -R	For details, see "5.5.3 Removing a Group" in "5.5 Removals."
Change attribute (Group name)	# sdxattr -G -a name= <i>groupname</i>	For details, see "5.4.2 Group Configuration" in "5.4 Changes."
Change attribute (Active disk)	# sdxattr -G -a actdisk= <i>disk</i>	Unsupported

Volume operation

Operation	Command Operation	GDS Management View Operation
Create	# sdxvolume -M	For details, see "5.2.2.4 Volume Configuration" in "5.2.2 Operating from the Settings Menu."
Remove	# sdxvolume -R	For details, see "5.5.2 Removing a Volume" in "5.5 Removals."
Start	# sdxvolume -N	For details, see "Start Volume" in "4.2.3 Operation."
Stop	# sdxvolume -F	For details, see "Stop Volume" in "4.2.3 Operation."
Expand	# sdxvolume -S	Unsupported
Change attribute (Volume name)	# sdxattr -V -a name= <i>volumename</i>	For details, see "5.4.3 Volume Configuration" in "5.4 Changes."
Change attribute (JRM)	# sdxattr -V -a jrm={on off}	For details, see "5.4.3 Volume Configuration" in "5.4 Changes."
Change attribute (Physical slice)	# sdxattr -V -a pslice={on off}	For details, see "5.4.3 Volume Configuration" in "5.4 Changes."
Change attribute (Lock volume)	# sdxattr -V -a lock={on off}	Unsupported
Change attribute (Access mode)	# sdxattr -V ... -a mode={rw ro}	Unsupported
Change attribute (JRM for proxy)	# sdxattr -V -a pjrm= off	Unsupported
Start copying	# sdxcopy -B	For details, see "5.3.6 Copying Operation" in "5.3 Operation in Use."
Cancel copying	# sdxcopy -C	For details, see "5.3.6 Copying Operation" in "5.3 Operation in Use."
Restore	# sdxfix -V	Unsupported

Slice operation

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

Disk operation

Operation	Command Operation	GDS Management View Operation
Swap	# sdxswap -O	For details, see "5.3.4 Disk Swap" in "5.3 Operation in Use."
Restore	# sdxswap -I	For details, see "5.3.4 Disk Swap" in "5.3 Operation in Use."
Change attribute	# sdxattr -D	For details, see "5.4.1 Class Configuration" in "5.4 Changes."
Clear errors	# sdxfix -D	Unsupported

Proxy operation

Operation	Command Operation	GDS Management View Operation
Join	# sdxproxy Join	For details, see "5.2.4.1 Join" in "5.2 Settings."
Part	# sdxproxy Part	For details, see "5.3.2.2 Backup (by Synchronization)" in "5.3 Operation in Use."
Rejoin	# sdxproxy Rejoin	For details, see "5.3.2.2 Backup (by Synchronization)" in "5.3 Operation in Use."
Rejoin and restore	# sdxproxy RejoinRestore	For details, see "5.3.3 Restore" in "5.3 Operation in Use."
Swap	# sdxproxy Swap	For details, see "5.3.5 Disk Migration" in "5.3 Operation in Use."
Relate	# sdxpsoxy Relate	For details, see "5.2.4.2 Relate" in "5.2 Settings."
Update	# sdxpsoxy Update	For details, see "5.3.2.3 Backup (by OPC)" in "5.3 Operation in Use."
Restore	# sdxpsoxy Restore	For details, see "5.3.3 Restore" in "5.3 Operation in Use."
Cancel hard copy session	# sdxpsoxy Cancel	Unsupported
Create alternative boot environment [Itanium]	# sdxproxy Root	Unsupported
Break	# sdxproxy Break	For details, see "5.5.6 Breaking a Proxy" in "5.5 Removals."

Shadow operation

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

Operation to display the configuration information

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

Other operations

Operation	Command Operation	GDS Management View Operation
System disk settings [Itanium]	# sdxdisk -M # sdxdisk -C # sdxroot -M <reboot>	For details, see "5.2.1 System Disk Settings [Itanium]" in "5.2 Settings."
Unmirror system disk [Itanium]	# sdxdisk -D # sdxroot -R <reboot> # sdxvolume -F # sdxvolume -R # sdxdisk -D # sdxdisk -R	For details, see "5.5.5 Unmirroring the System Disk [Itanium]" in "5.5 Removals."
Configuration parameter operations	# sdxparam	Unsupported
Object configuration operations [Linux2.6]	# sdxconfig	Unsupported

5.2 Settings

This section explains the GDS setting procedure.

5.2.1 System Disk Settings [Itanium]

System Disk Settings allows you to mirror system disks.

System disks cannot be mirrored under the following conditions.

Change the disk configuration before performing the system disk settings.

- The partitioning method is not the GUID partition table (GPT).
- The disk does not have an unused (unassigned to a slice) slice number.
- The system disk does not have sufficient free space or a swap area.

For details on necessary free space (or a swap area), refer to "A.2.5 Disk Size."



Note **Disk that will be mirrored by system disk settings**

For system disk settings from the GDS Management View, disks with /, /var, /usr, /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].



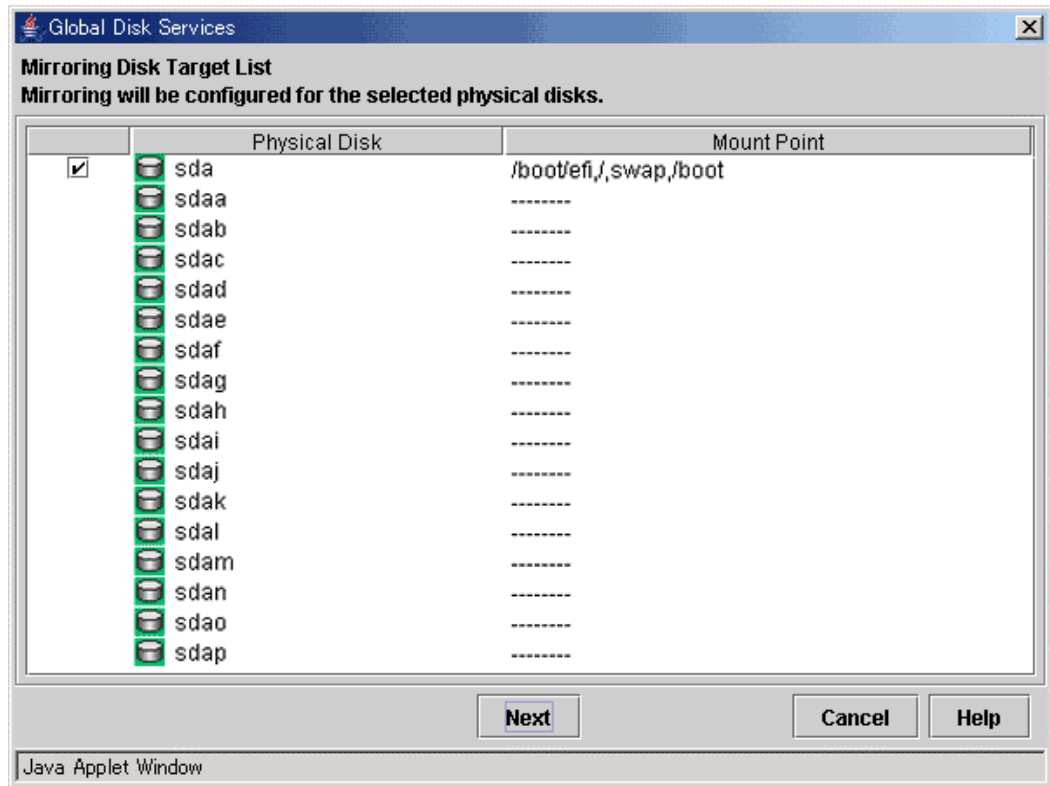
Note **For Safe System Disk Mirroring**

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.

How to mirror a system disk
 In the [Settings] menu, select [System Disk Settings].

1. Confirming original disks

Figure: 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.

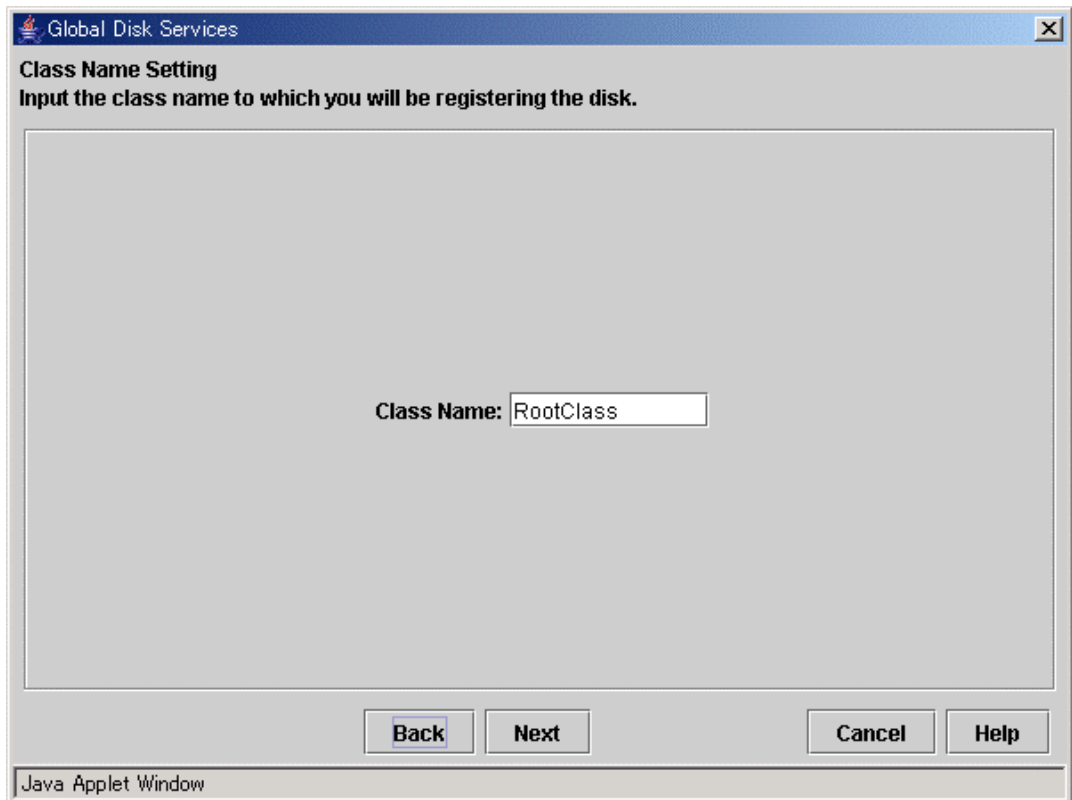
- It is possible to uncheck system disks for which only swap among /, /var, /usr, /boot, /boot/efi and swap is displayed in the [Mount Point] field.
- It is impossible to uncheck system disks for which /, /var, /usr, /boot, and /boot/efi are displayed in the [Mount Point] field.

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

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

2. Creating root class

Figure: Class Name Setting



Type the class name of the root class.

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



Note Inputting the Class Name

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

`/dev/sfdsk/class_name/rdsk/volume_name` [Linux2.4]

`/dev/sfdsk/class_name/dsk/volume_name`

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



Note When Setting the System Disk in Cluster System

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



See

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

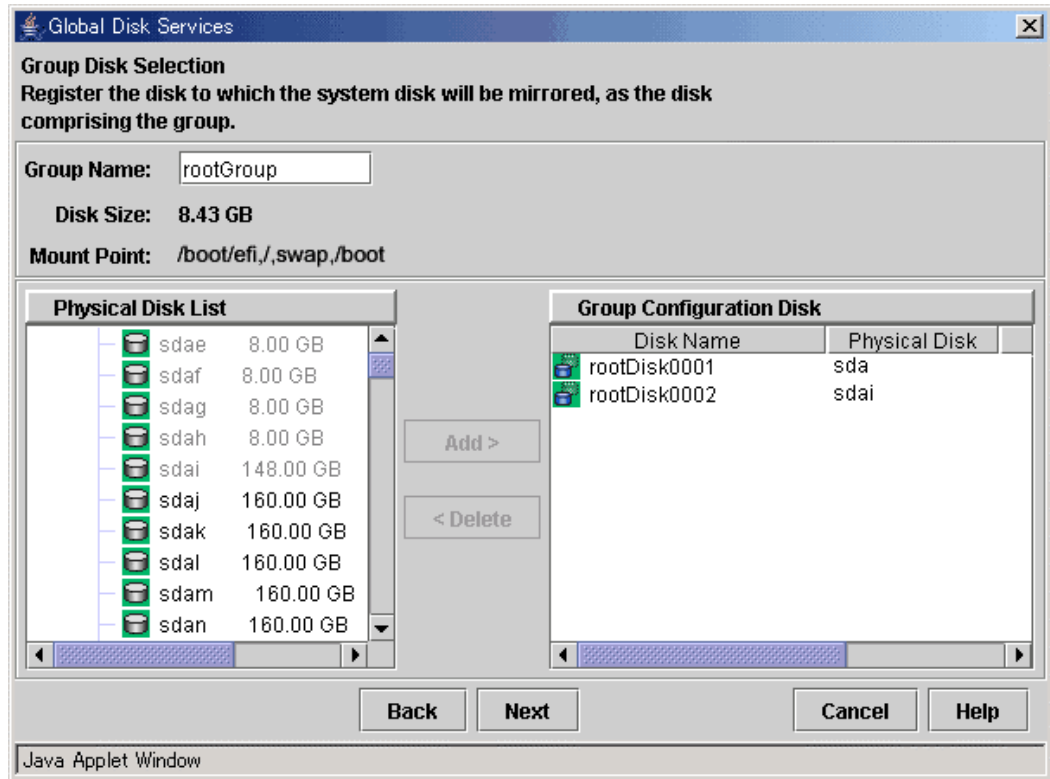
Click <Next> to continue.

3. Creating group

Create the group by selecting the mirror disks.

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

Figure: Group Disk Selection: rootGroup



Group name, size and mount point will be displayed.

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

Change the [Group Name] if necessary.

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

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

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

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

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

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

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

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



Note

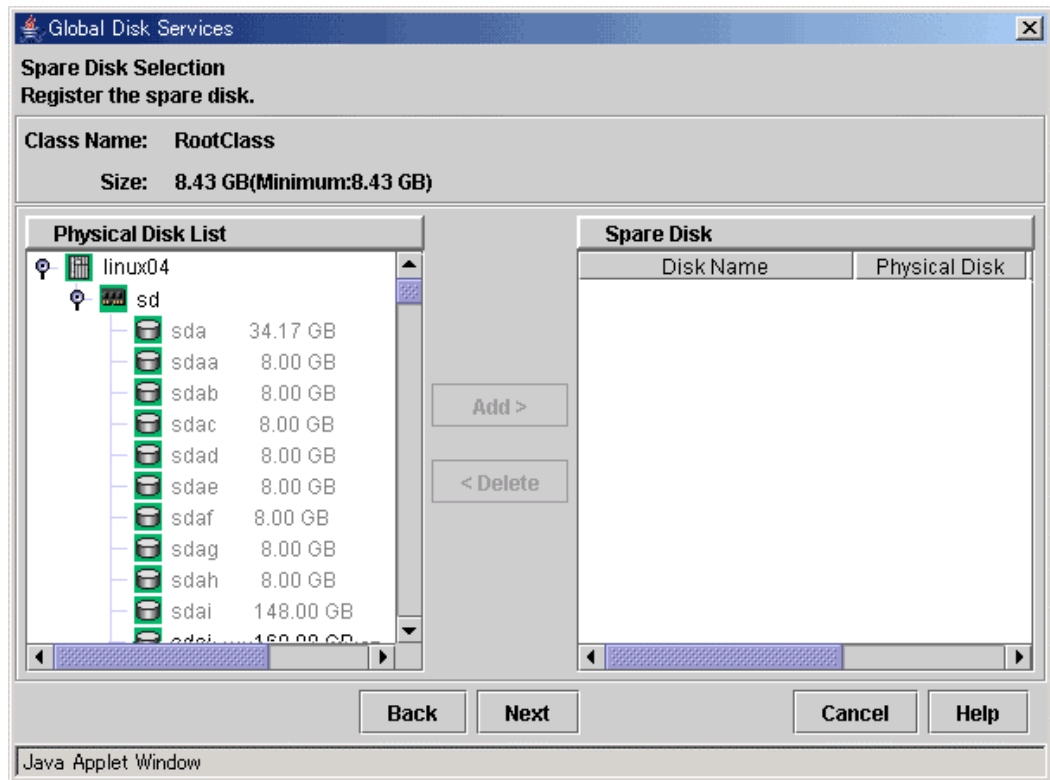
Physical disks that can be registered as the Group Configuration Disk

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

Users should always add a physical disk that is larger than the original disk.

4. Registering spare disk

Figure: 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>.

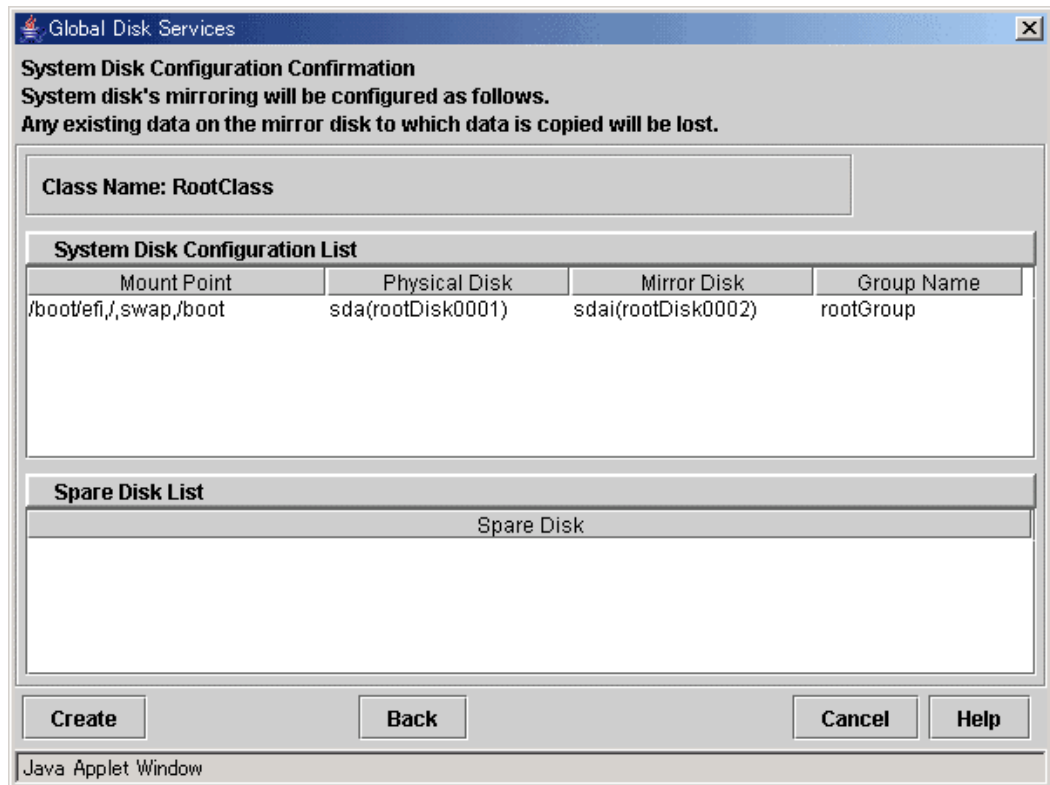


Point Spare Disk Size

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: 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.




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`)

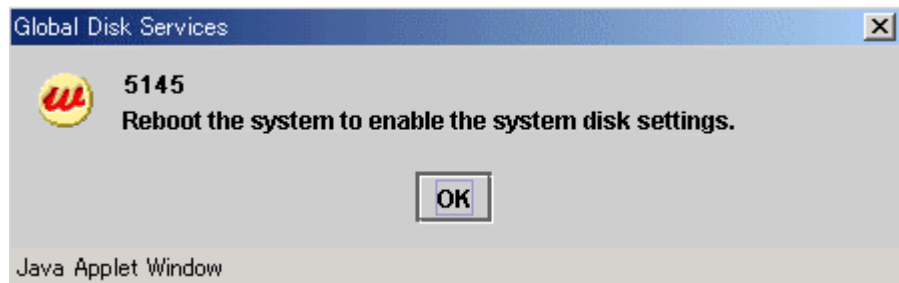
6. Completing system disk configuration

Figure: Setting System Disk Mirroring Complete



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

Figure: System Reboot Message



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



Note JRM and System Disk Settings

JRM (Just Resynchronization Mechanism) for a volume created by system disk settings will be set as "Yes." To disable the JRM feature, select the volume in the Main screen and change the attribute by clicking [Change Attributes] in the [Operation] menu.



Note Checking Physical Disk Information and Slice Numbers

After rebooting the system, check physical disk information and the slice numbers of system disks and make a note of them. For details, refer to "6.1.1 Checking Physical Disk Information and Slice Numbers."

5.2.2 Operating from the Settings Menu

From the Main screen, select [XXXXX Configuration] in the [Settings] menu. Settings screen for different features appears. Use the screen switch tab to switch between "Class Configuration," "Group Configuration" and "Volume Configuration."



Note **Configuring a System of High Reliability**

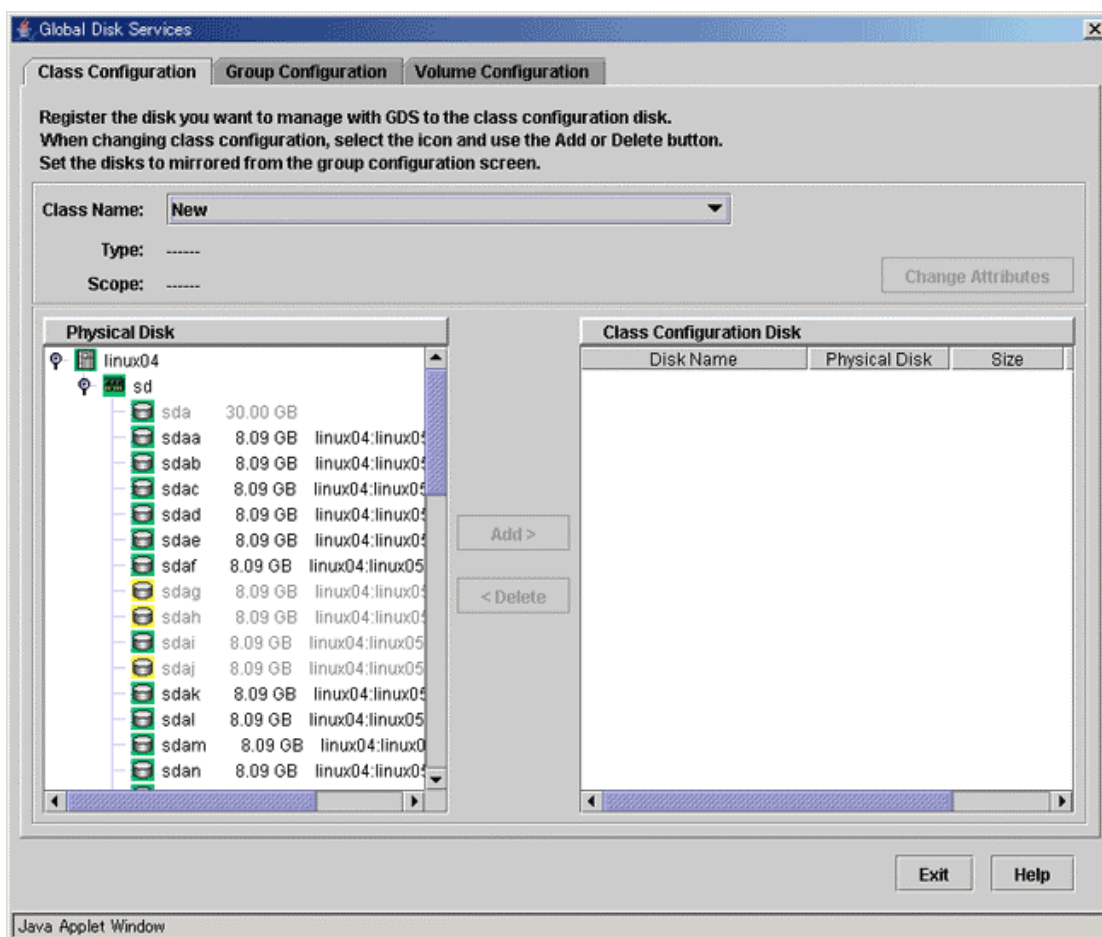
In order to ensure a system of high reliability, mirroring disks connected to separate I/O adapter, controller and cable is recommended.

5.2.2.1 Class Configuration

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

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

Figure: Class Configuration

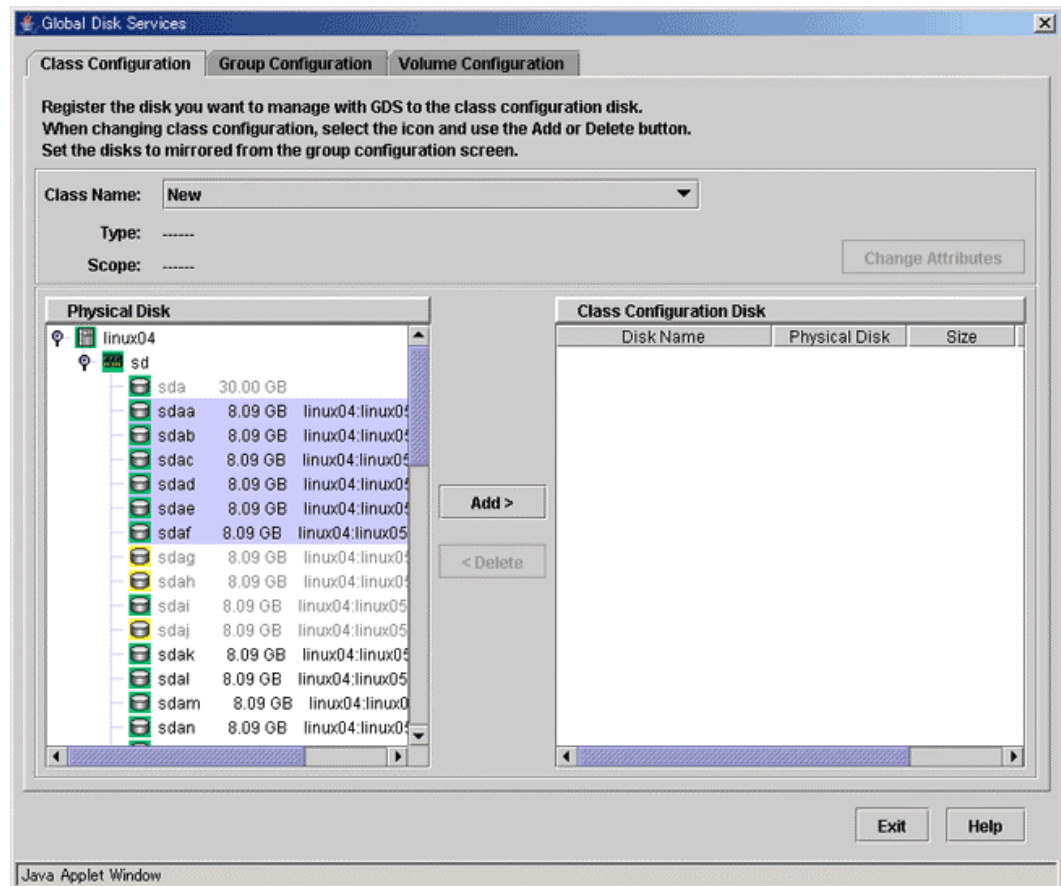


1. **Selecting class name**
In the [Class Name] list, select "New."

2. Selecting physical disk

In the [Physical Disk] field, select the physical disk you want to include in the disk class. You can select more than one physical disk at a time. Selecting a physical disk will make the <Add> button available.

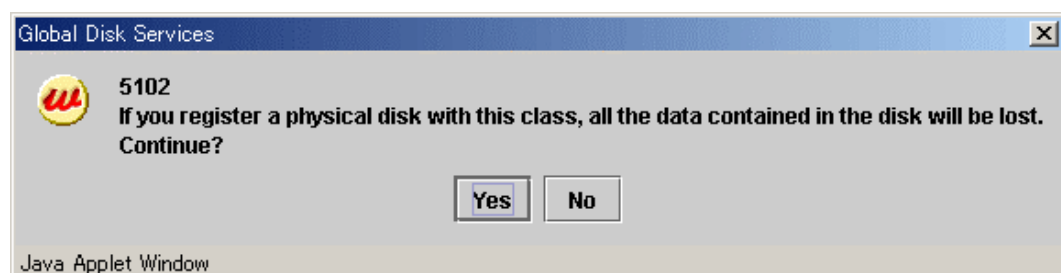
Figure: Selecting Physical Disk to Configure Class



3. Creating a class

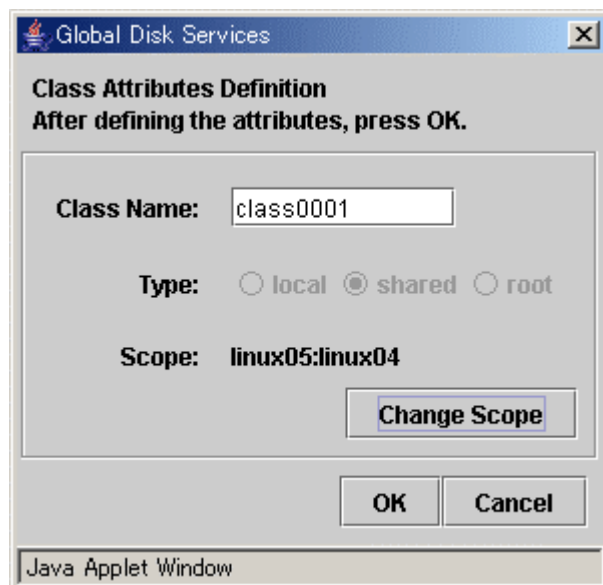
Click <Add>, and the message below appears.

Figure: Warning



Click <Yes> to continue, and <No> to cancel. Click <Yes>, and Class Attributes Definition screen appears.

Figure: 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>.



Note **Inputting Class Name**

Class name will be used for the device path name.
/dev/sfdsk/class_name/rdsk/volume_name [Linux2.4]
/dev/sfdsk/class_name/dsk/volume_name

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



Note **Creating Local Type Class with Cluster System**

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



See

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

This operation determines the class name.

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

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

4. Setting disk attributes

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

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

5. Completing class creation

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



Note **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.



See

For information on how to create a shared type class in cluster systems, see "5.2.2.2 Cluster System Class Configuration."

5.2.2.2 Cluster System Class Configuration

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

1. Setting [Type]

Sets the class type.

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

Selecting a shared physical disk, sets the type to "shared."

2. Displaying [Scope]

Displays connecting nodes that can share a class.

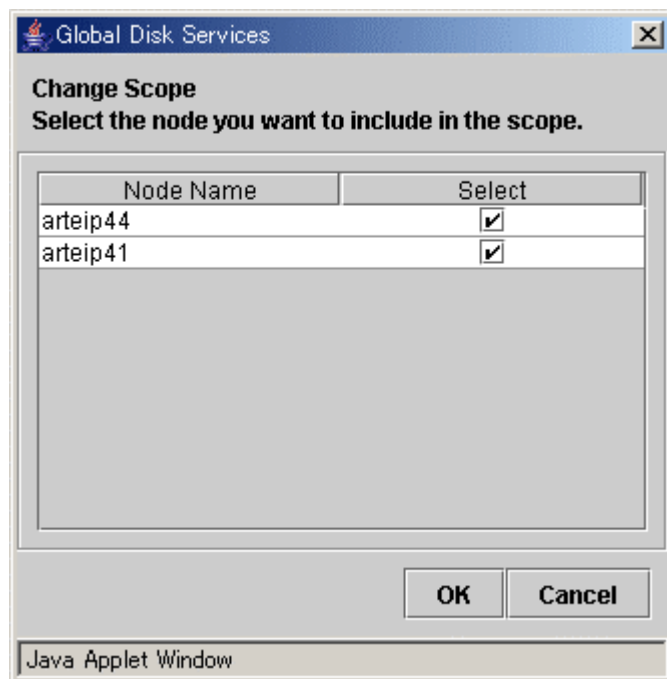
To change scope, click <Change Scope>.

3. <Change Scope> Button

Changes nodes connecting to class.

Click <Change Scope>, and Change Scope screen appears.

Figure: 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 **Class Resources**

- 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.



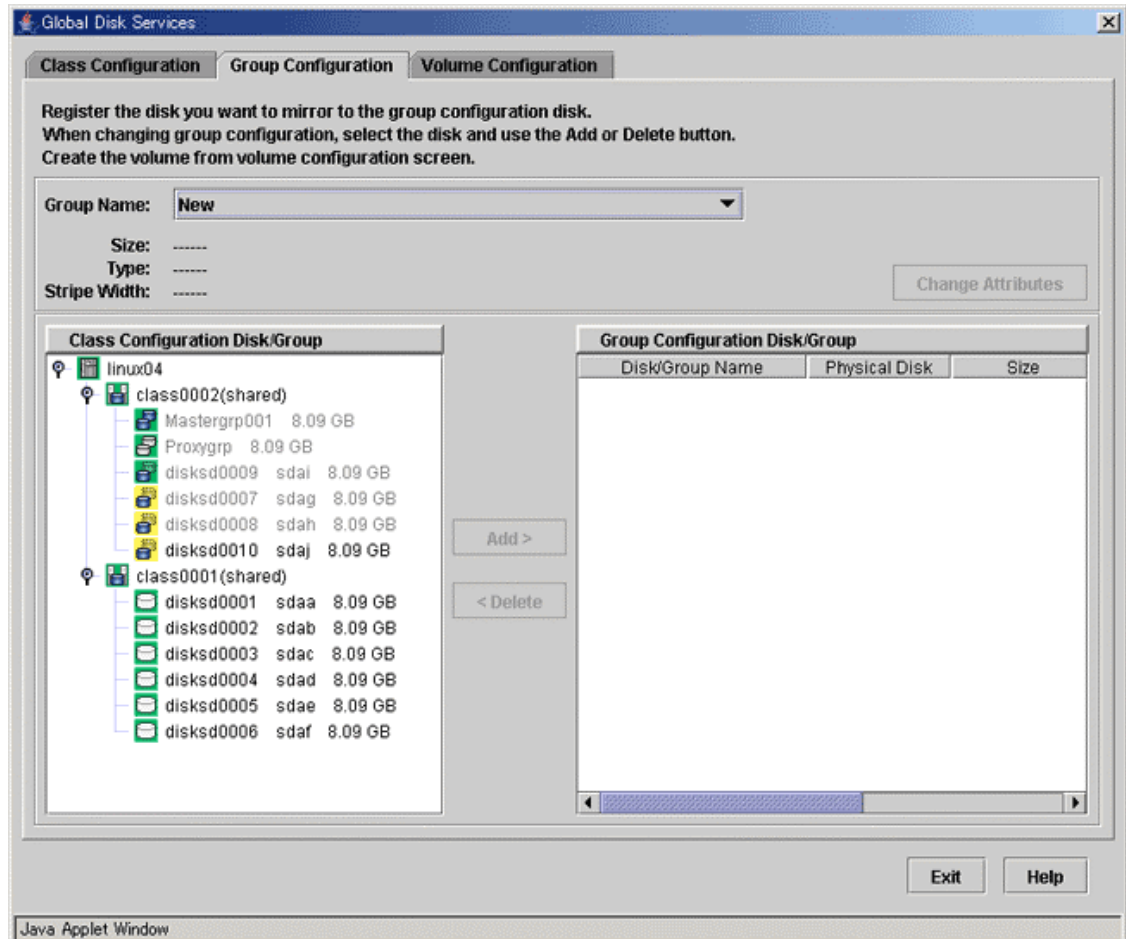
See

For information on installation and initial setting of cluster systems, refer to "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide."

5.2.2.3 Group Configuration

In this section, how to create a new group is explained.
In the [Settings] menu, select [Group Configuration]. Group Configuration screen appears.

Figure: Group Configuration



Follow the procedures below to create a new group.

- 1. Selecting group name**
In the [Group Name] list, select "New."

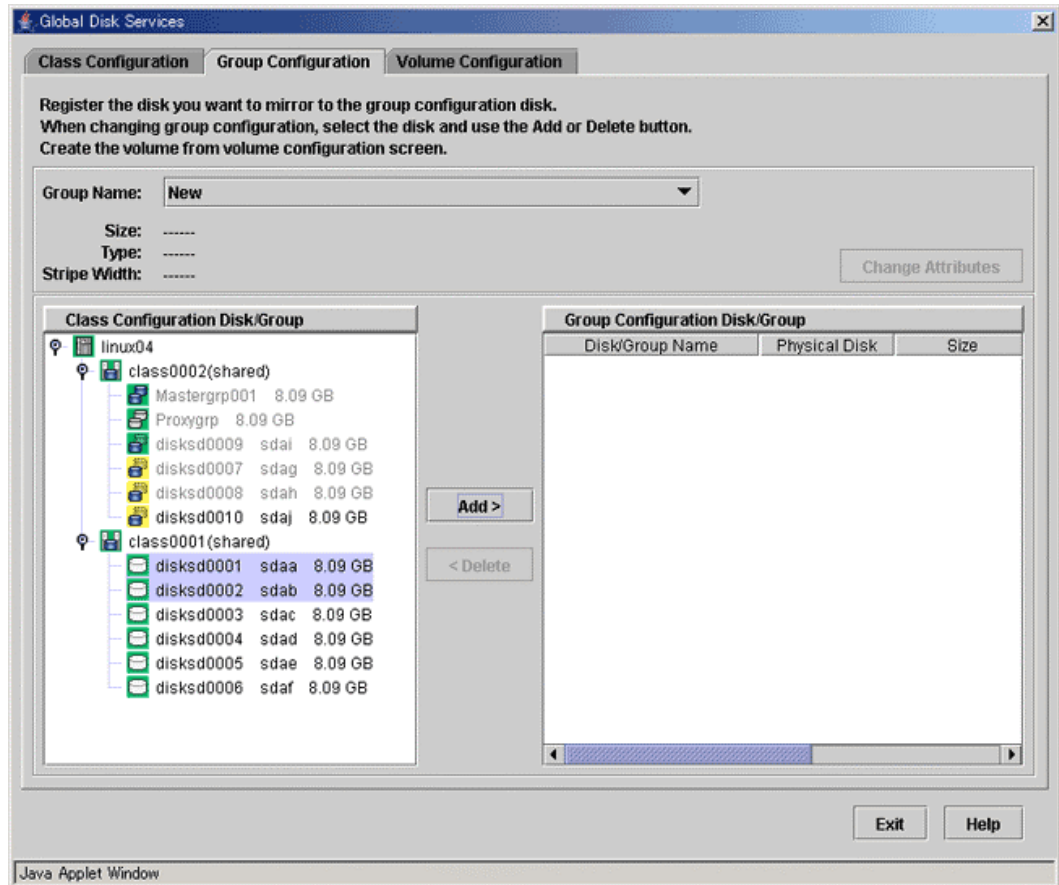
2. Selecting disk/lower level group

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: 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: Group Attributes Definition Screen



a) Setting [Group Name]

Enter the group name.
Change the default group name if necessary.

b) Setting [Type]

Set the group type.
Select "mirror" for mirroring, "stripe" for striping, and "concat" for concatenating. The default setting is "mirror."

c) Setting [Stripe Width]

You will be able to enter this field only when you select "stripe" for the [Type]. For the stripe width, you can specify a value of two raised to the power.
The default setting is "32."
After setting the attributes, click <Exit> and a new group will be created.
If you click <Cancel> in the Group Attributes Definition screen, connection of the disk itself will be canceled.



See

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

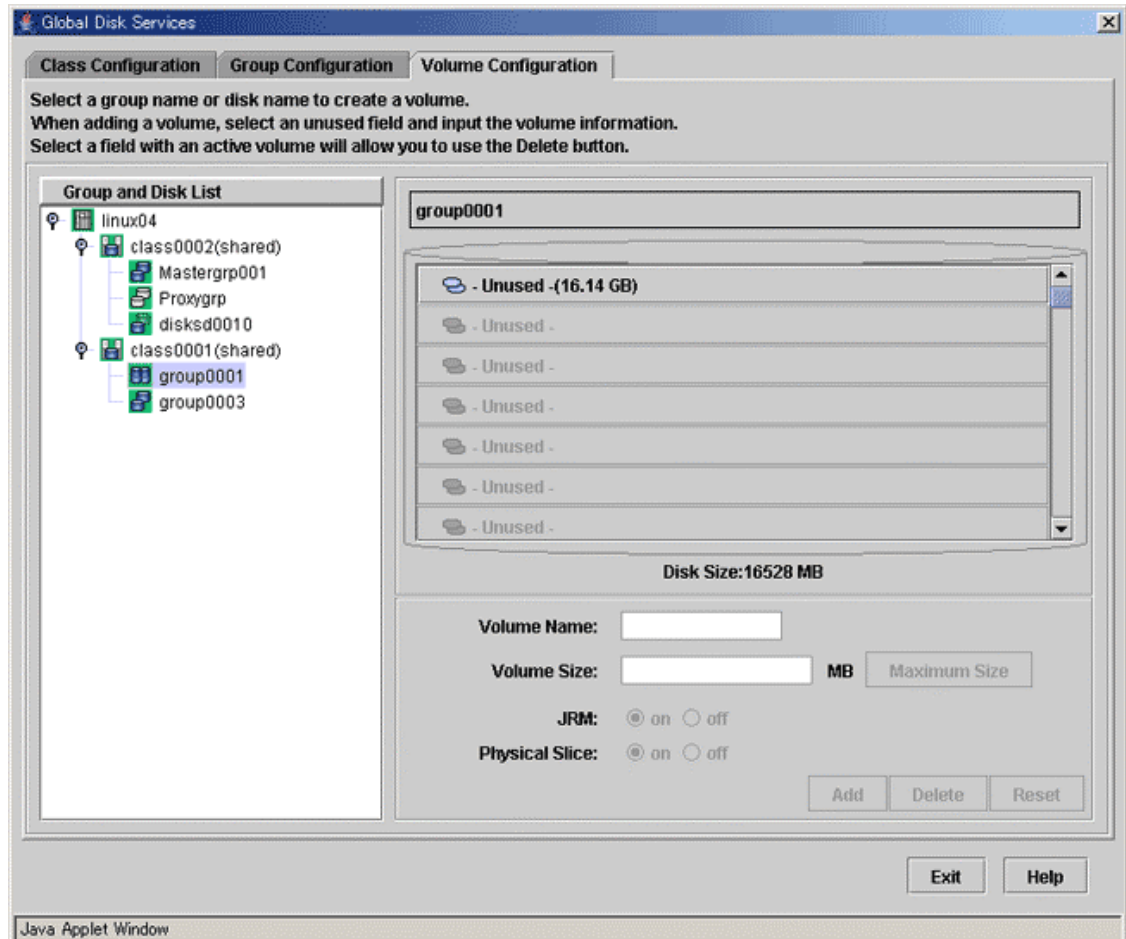
4. Completing group creation

After creating all groups, click <Exit> and close Group Configuration screen.

5.2.2.4 Volume Configuration

In this section, how to create a new volume is explained.
In the [Settings] menu, select [Volume Configuration]. The Volume Configuration screen appears.

Figure: Volume Configuration



Follow the procedures below to create a volume.

- 1. Selecting group/disk**
In the [Group and Disk List], select the group or the disk.
- 2. Selecting unused volume**
Click <Unused> with a volume icon, and select an unused volume.

3. Setting volume attributes

Selecting an unused volume will allow you to type in the volume attributes field (Volume Name, Volume Size, JRM, Physical Slice).

a) Setting [Volume Name]

Type the volume name.



See

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

b) Setting [Volume Size]

Type the volume size in MB units, using numbers only.



Note Size and Cylinder Boundary

When the specified size does not match the cylinder boundary, it will automatically adjust by rounding up. For details, see "A.2.6 Volume Size."

c) [Maximum Size] button

Sets the value in the [Volume Size] field to the maximum available size.

d) Setting [JRM]

Default is set to "on." Change the setting when you want to disable the just resynchronization copying feature.

When you select a stripe group or concatenation group in step 1, the setting will fail.

e) Setting [Physical Slice]

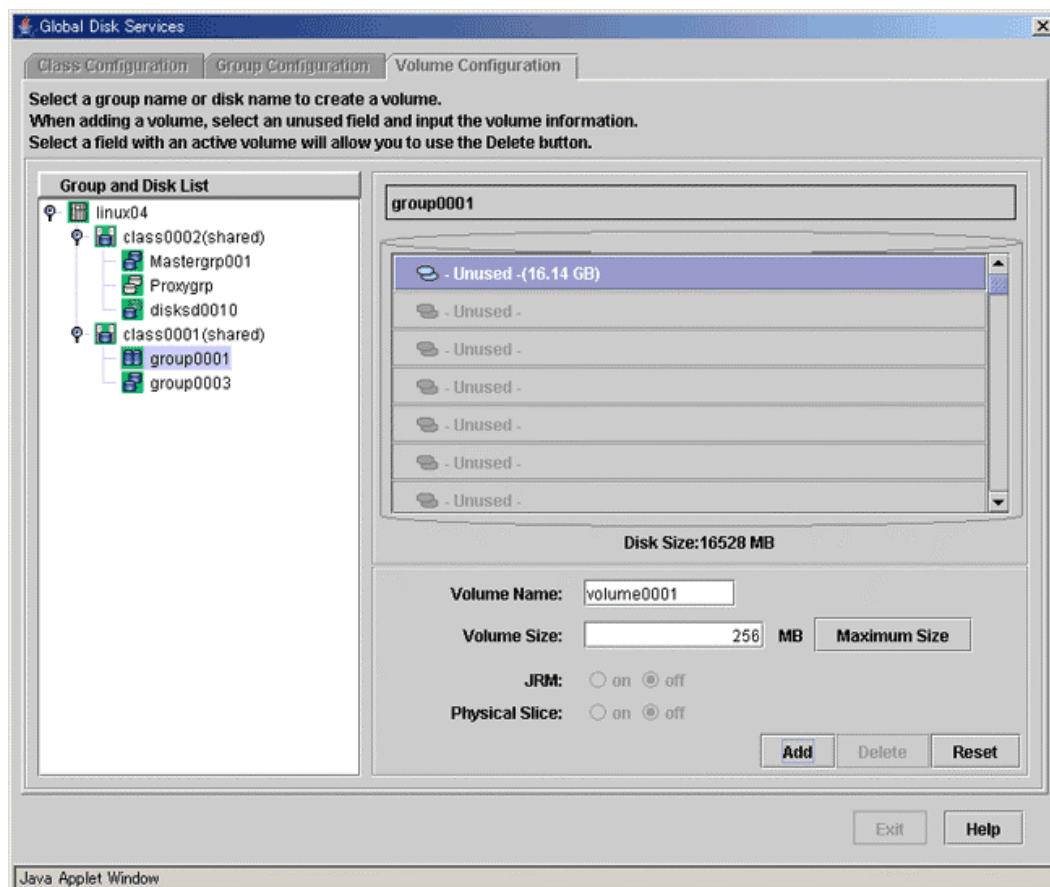
Sets the volume's physical slice attribute value. Default is the "on."

on : a volume which consists of physical slices is created.

off: a volume without physical slices is created.

When you select a stripe group or concatenation group in step 1, the setting will fail.

Figure: Setting Volume Attributes



4. Determining the new volume

After setting the attributes, click <Add>. A new volume will be created.

If you click <Reset>, creation of the new volume will be canceled.

After creating all volumes, click <Exit> and close Volume Configuration screen.

After creating the volume, the volume will be started. You can access to the volume using the following special files.

```
/dev/sfdsk/class_name/rdisk/volume_name [Linux2.4]  
/dev/sfdsk/class_name/dsk/volume_name
```



Note Shared Class Volume Operation

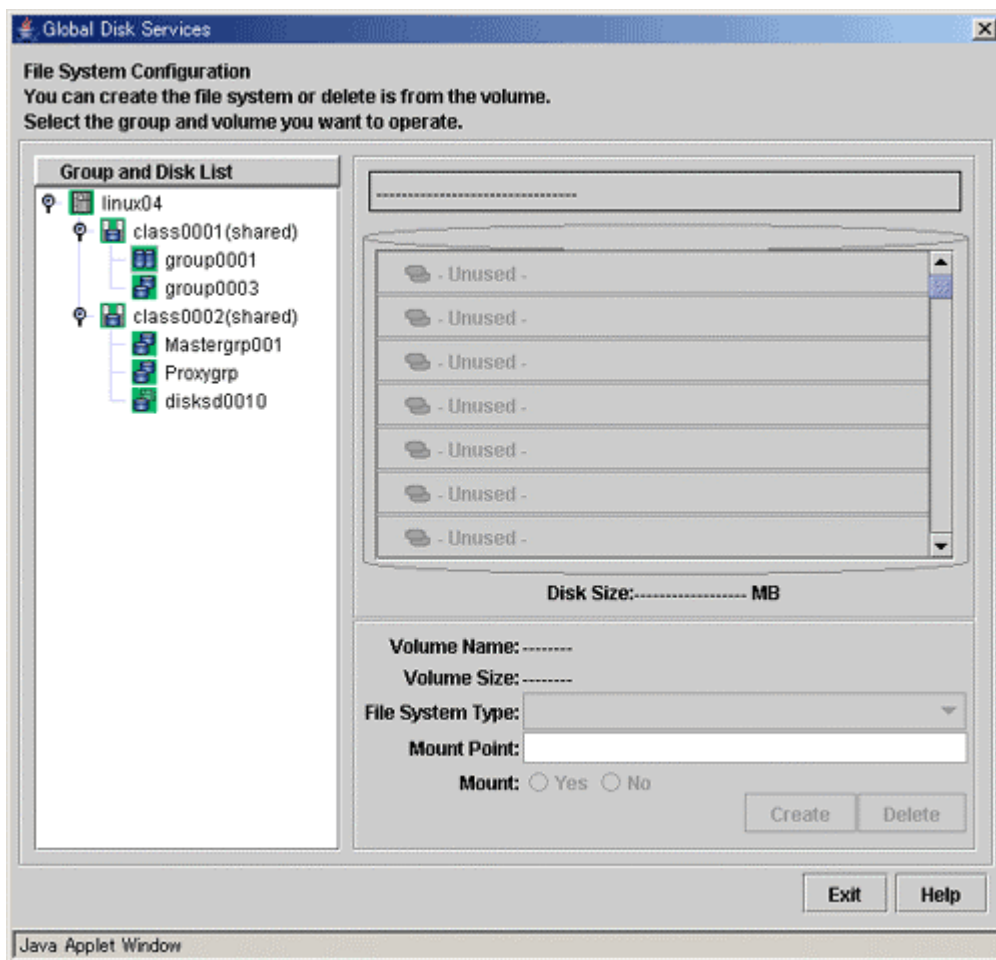
Volume created in a shared class cannot be used from other nodes immediately. If you wish to access from other nodes, you must activate the volume from the node you wish to access.

After activating the volume, the node which created the volume, and the node which activated the volume will gain access. However, since operating from two nodes could affect data consistency, you must be careful when operating the volume.

5.2.3 File System Configuration

In this section, how to create a file system in a volume is explained. In the [Settings] menu, select [File System Configuration]. File System Configuration screen appears.

Figure: File System Configuration



1. Selecting a group/disk

In the [Group and Disk List] field, select a group or disk with which you want to perform an operation.

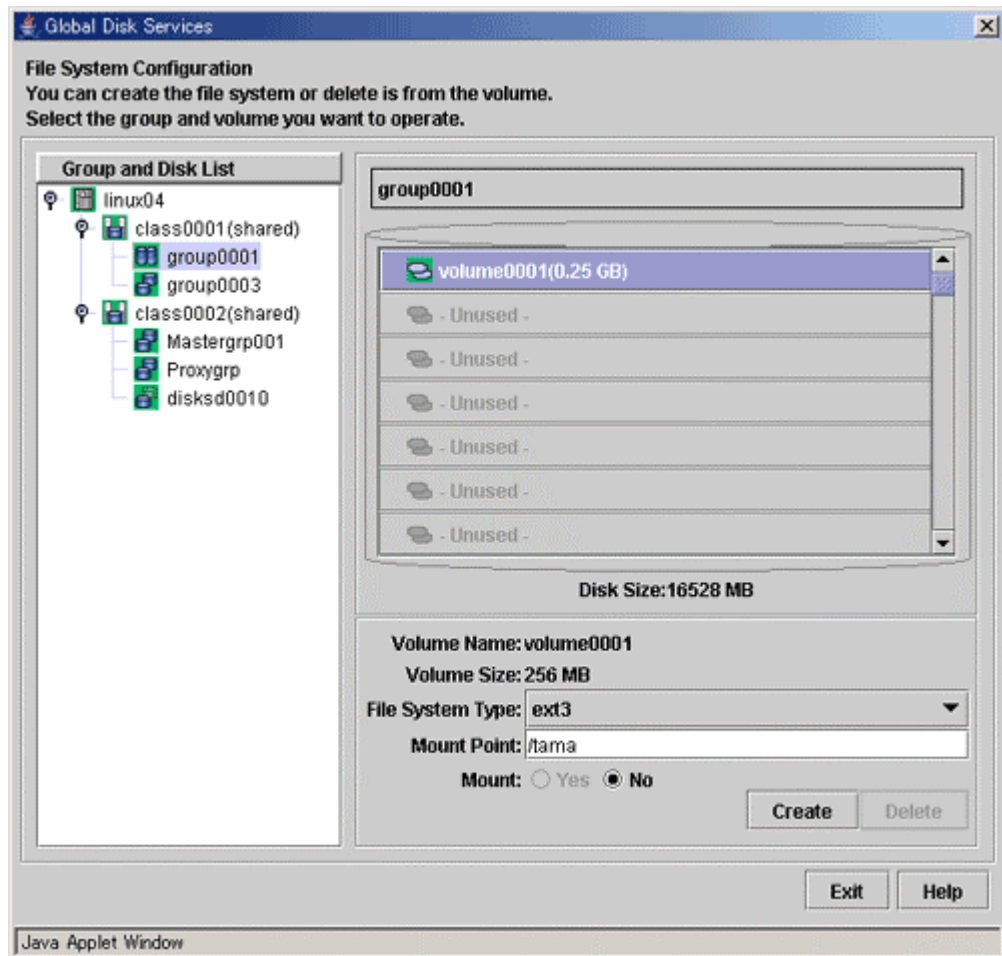
2. Selecting a volume

Select a volume in which you want to create a file system.

3. Setting file system attributes

Selecting a volume will allow you to type in the file system attributes field (File System Type, Mount Point, Mount).

Figure: Setting File System Attributes



a. Setting [File System Type]

Select the file system type.

b. Setting [Mount Point]

Type the mount point you want to set for /etc/fstab.

c. Setting [Mount]

Select "No".

Once file systems are created in step 4, mount information is added to the /etc/fstab file and "noauto" is specified in the fourth (mount option) field. After completing step 4, modify the mount information in the /etc/fstab file as needed.



Note Setting [Mount]

Do not select "Yes." Additionally, do not remove "noauto" from mount information added to the /etc/fstab file. For details, see "A.2.29 File System Auto Mount."

4. Creating the file system

After setting the attributes, click <Create>. A new file system will be created.
After creating all file systems, click <Exit>.



Note **In Cluster Systems**

When using a volume of a shared disk as a file system in a cluster system, certain settings are required after creating the file system. For details on how to set the file system created on a shared disk, see "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide."

5.2.4 Proxy Configuration

This section describes the procedures for relating proxy objects (proxy volumes or groups) to master objects (volumes or groups) in a system on which GDS Snapshot is installed.

The following two methods are available.

- **Join**
Relate a proxy to a master and join them. Synchronization copying from the master to the proxy is launched, and after the copying is complete, they are synchronized. When joining groups, proxy volumes are created in the proxy group and they are joined to their corresponding master volumes.
- **Relate**
Relate a proxy to a master and leave them parted. The statuses and contents of the master and the proxy remain unchanged.

5.2.4.1 Join

This subsection describes procedures for joining proxy volumes or proxy groups to volumes or groups in GDS Snapshot installed systems.

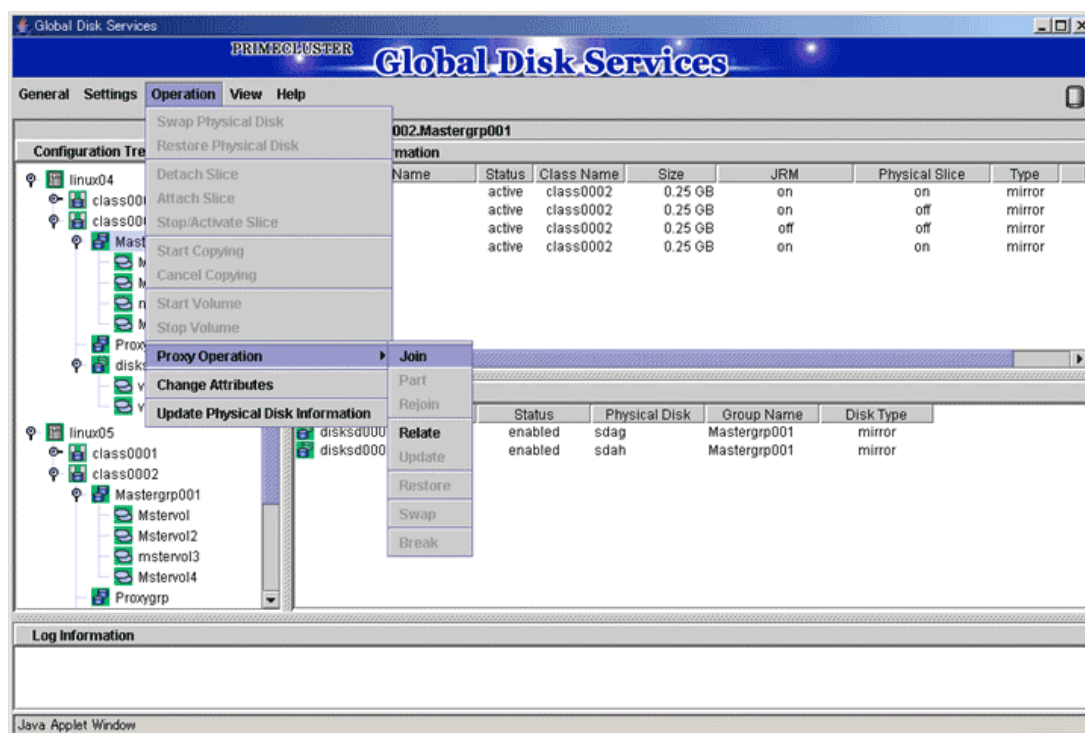
1. Selecting a master volume or a master group

Click an icon of a volume or a group to which a proxy is joined in the Main screen.

2. Selecting the [Join] menu

Select [Operation]:[Proxy Operation]:[Join] in the Main screen.

Figure: 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: 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>.

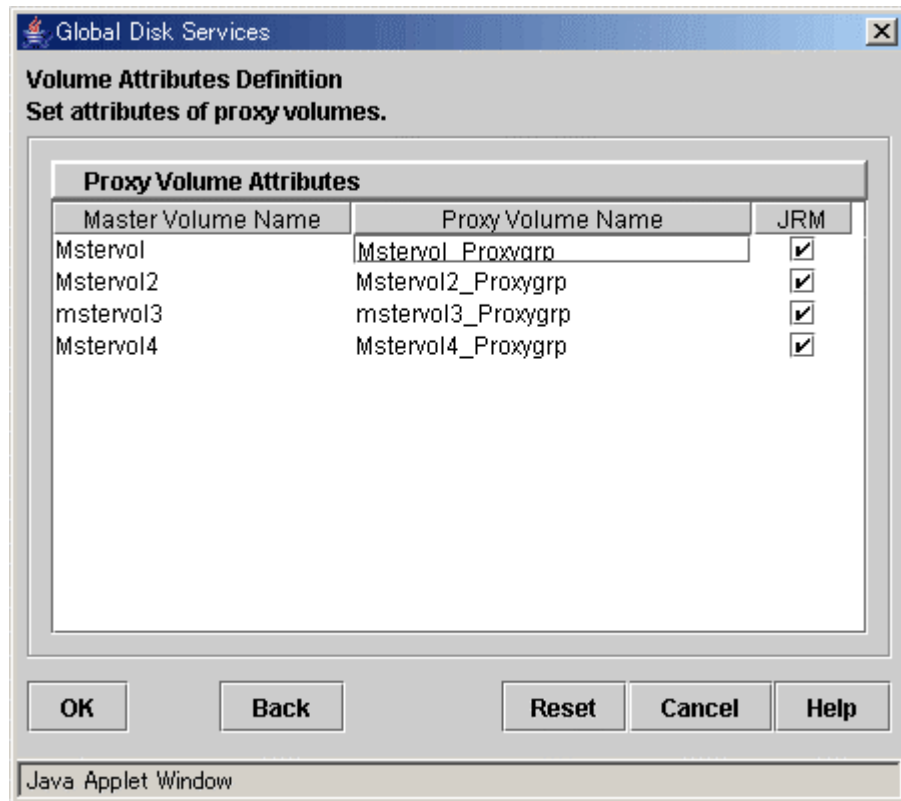
When joining groups, select a group to be joined to the master group and click <Next>.

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: Proxy Volume Attributes Definition



When a proxy group is joined to a master group, proxy volumes are created within the proxy group and joined to corresponding master volumes within the master group. In the Volume Attributes Definition dialog box, set attributes of such proxy volumes.

a) Proxy Volume Name

Assign volume names of proxy volumes. [Proxy Volume Name] shows default volume names. To change the default value, click and edit the volume name.



See

For the volume naming conventions, see "A.1.1 Object Name."



Information Automatically Generated Proxy Volume Names

Proxy volume names are automatically generated as "master volume name" + "_" (underscore) + "proxy group name" (e.g. volume0001_group0002).

b) JRM

Set JRM (Just Resynchronization Mechanism) for volumes. The default value is "on." To turn "off", uncheck the [JRM] box.



Information JRM for Volumes

The JRM setting in the Proxy Volume Attributes Definition dialog box is the "JRM for Volumes" of the proxy volume. Note that it is not the "JRM for Proxies." For details, see "A.2.12 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: Confirming "Join" (Volumes)



To continue the process, click <Yes>. Clicking <No> displays the Select Proxy dialog box shown in step 3 again.

Figure: 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: Information Message of "Join" Completion



Click <OK> to close the information message window.

5.2.4.2 Relate

This subsection describes the procedures for relating proxy volumes or proxy groups to volumes or groups in GDS Snapshot installed systems.



Information Data of Related Masters and Proxies

Even if masters and proxies are related, data of the master objects and the proxy objects remain unchanged.

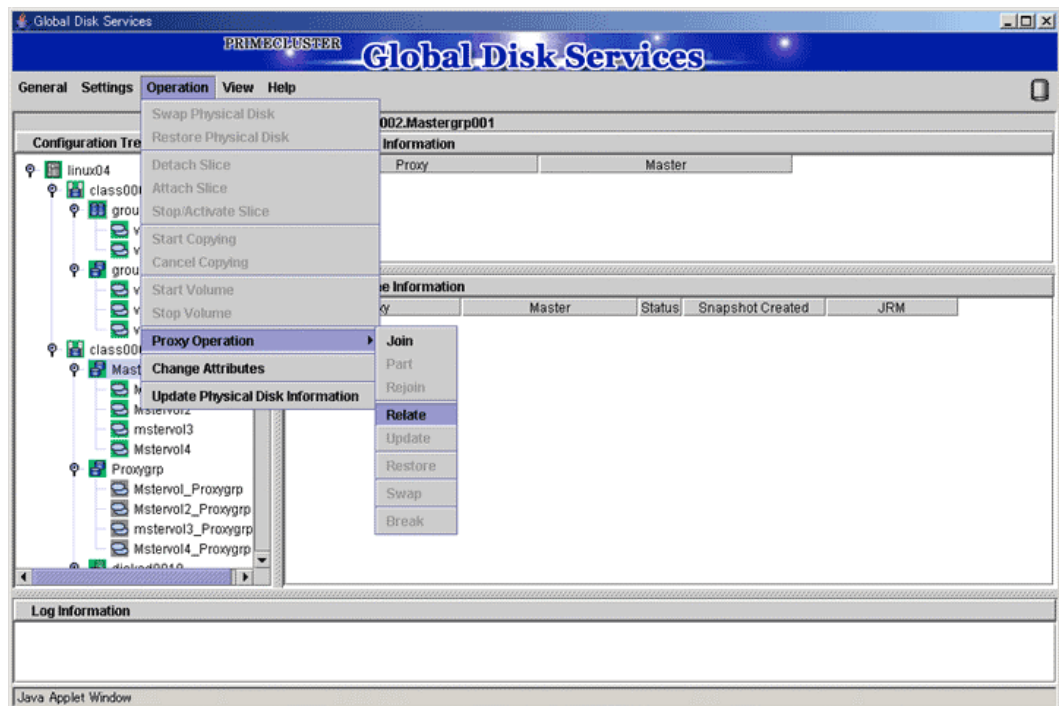
1. Selecting a master volume or a master group

Click the icon of a volume or a group to which a proxy is related in the Main screen.

2. Selecting the [Relate] menu

Select [Operation]:[Proxy Operation]:[Relate] in the Main screen.

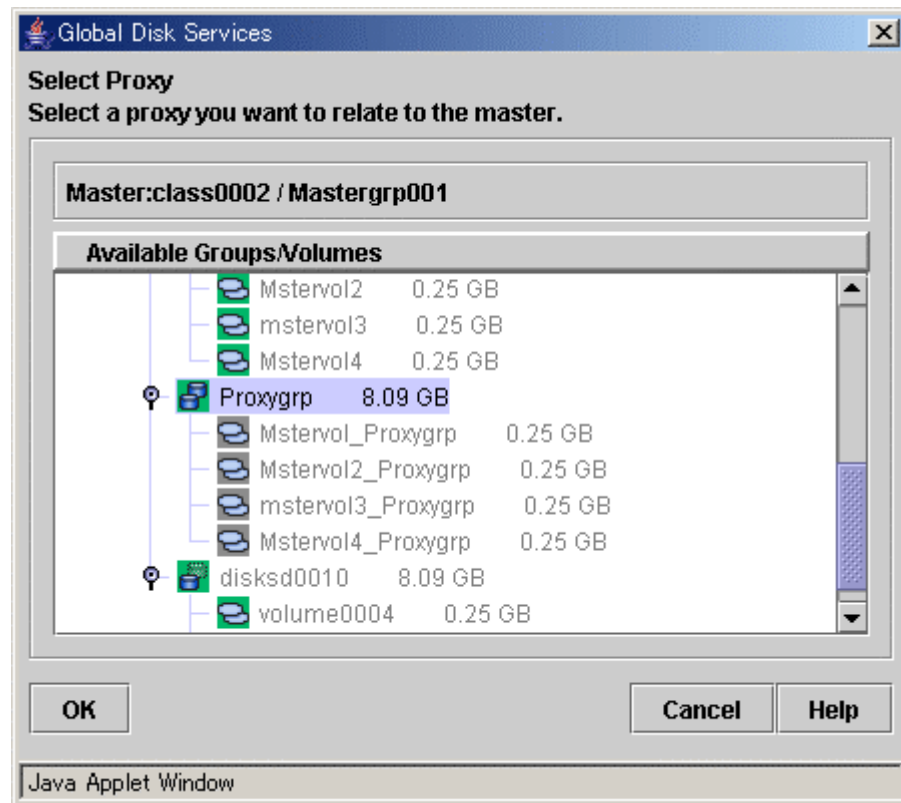
Figure: Relate



3. Selecting a proxy to be related

The Select Proxy dialog box appears.

Figure: 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: 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: Information Message of "Relate" Completion



Click <OK> to close the information message window.

5.3 Operation in Use

GDS operates monitoring and maintenance from the Main screen. This section explains the following operations:

- Viewing Configurations and Monitoring Statuses
- Backup
- Restore
- Disk Swap
- Disk Migration
- Copying Operation

5.3.1 Viewing Configurations and Monitoring Statuses

In the Main screen, you can view configurations and monitor statuses of objects.

5.3.1.1 Confirming SDX Object Configuration

Select [SDX Object] in the [View]:[Details] menu to view object configurations according to the following units.

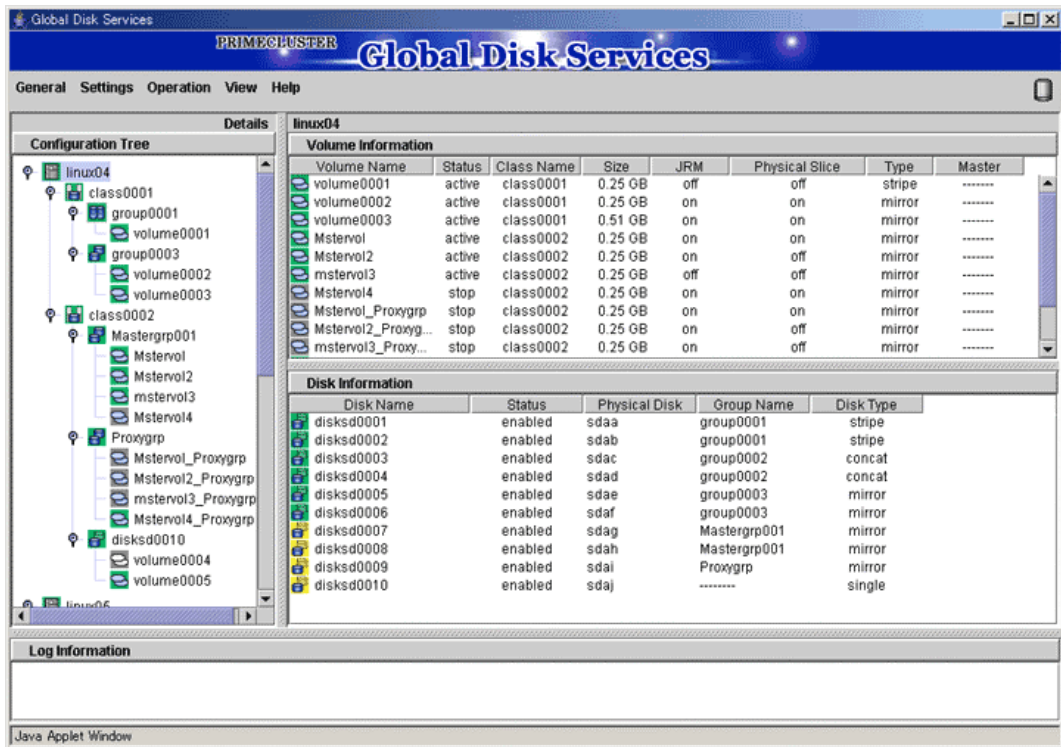
- Object configuration within a node
- Object configuration within a class
- Object configuration within a group
- Object configuration within a single disk
- Object configuration within a volume

For GDS Snapshot shadow objects, the object names, the status and so on are displayed in italics.

Object configuration within a node

Click the node icon in the Configuration Tree field, and all volumes and disks within the specified node appear.

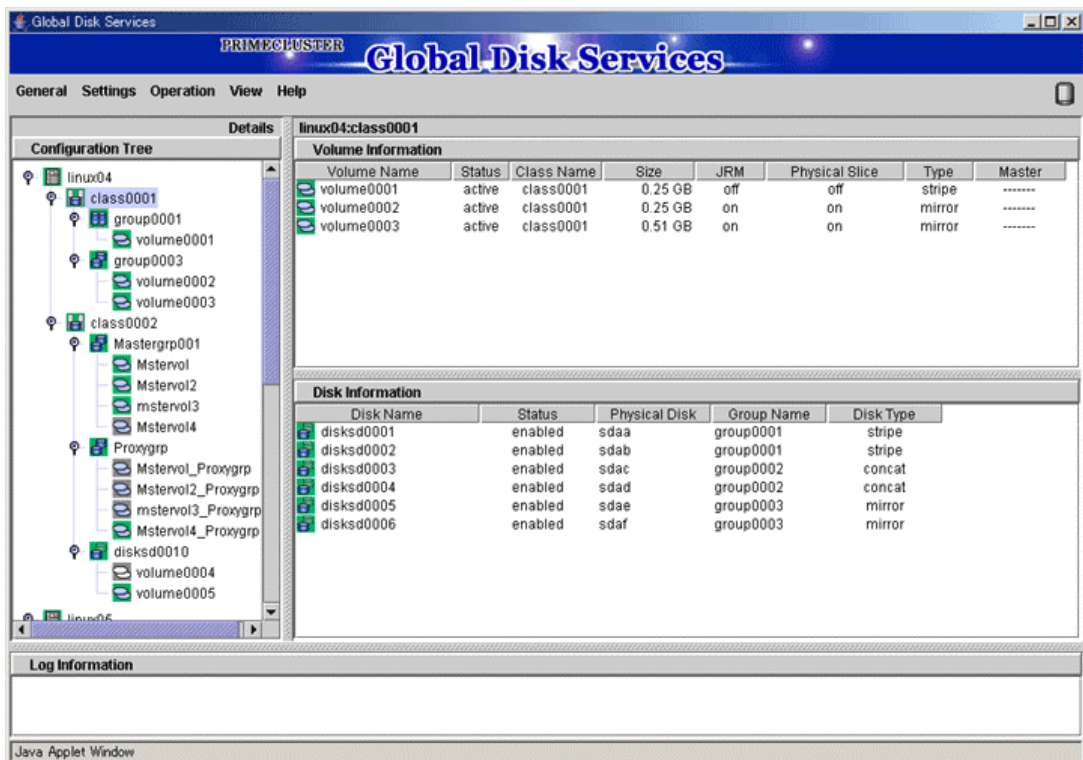
Figure: Main Screen (for SDX Objects of a Node)



Object configuration within a class

Click the class icon in the Configuration Tree field, and all volumes and disks within the class appear.

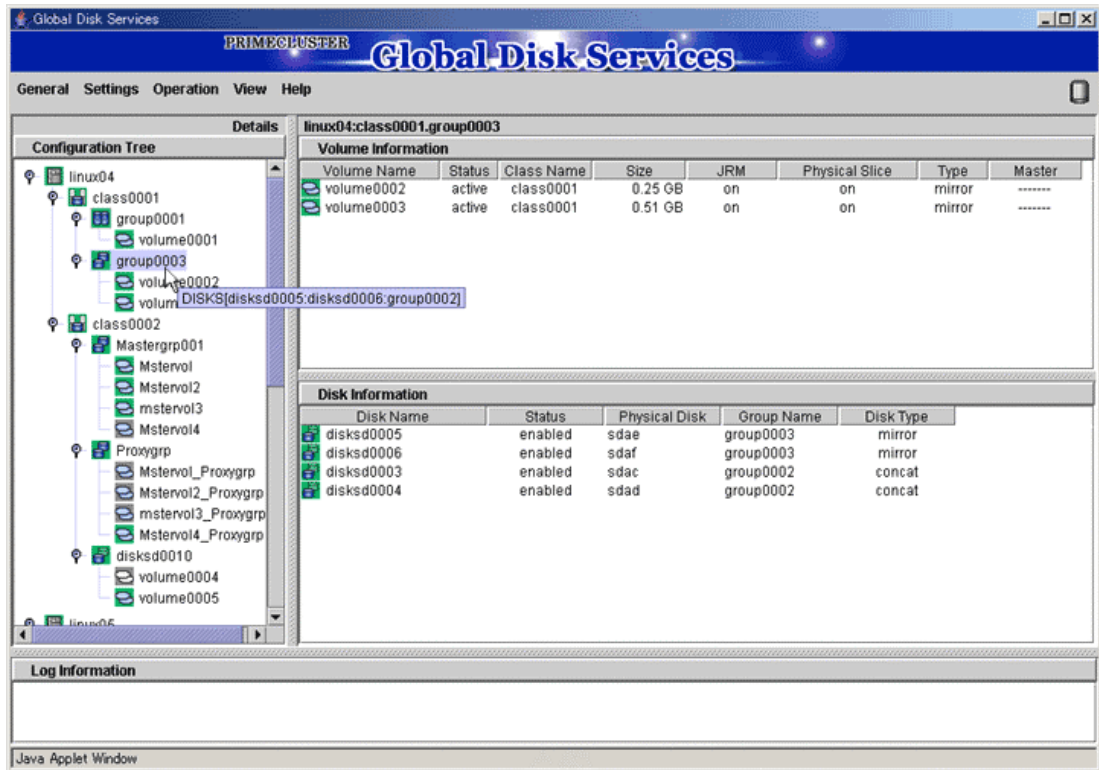
Figure: Main Screen (for SDX Objects of a Class)



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: Main Screen (for SDX Objects of a Group)



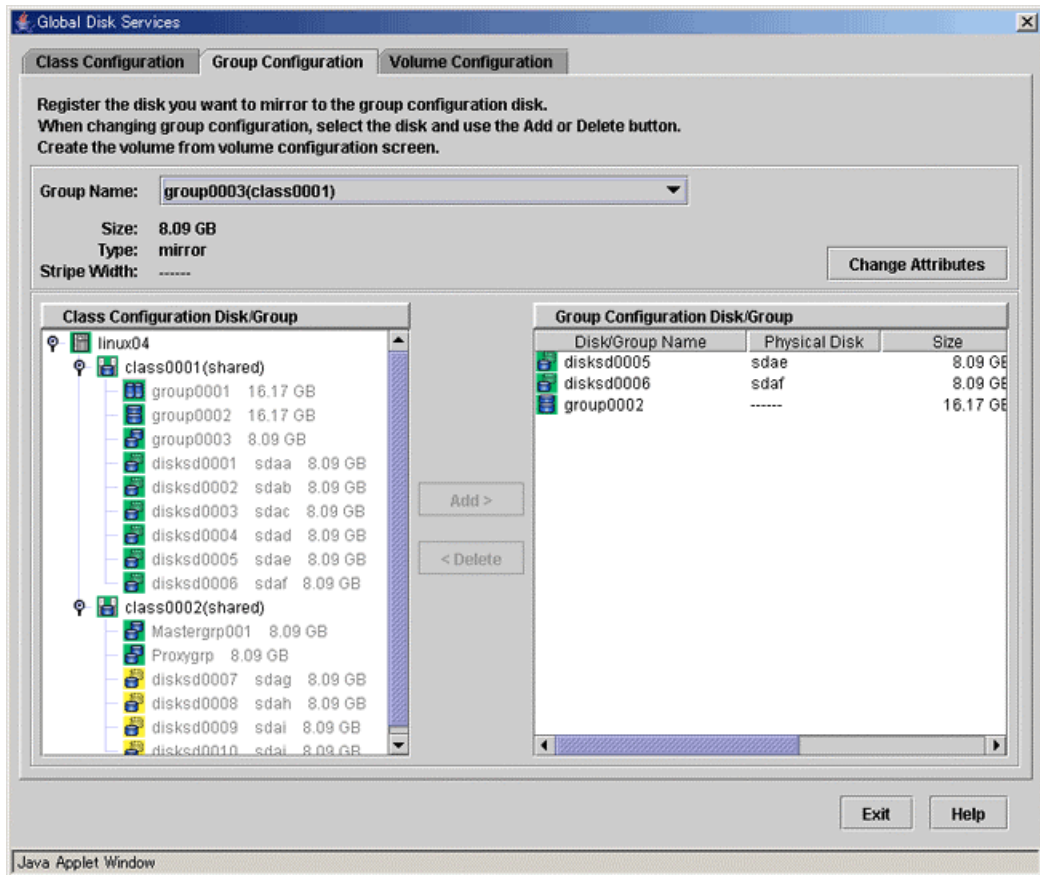
Groups displayed in the GDS configuration tree field are only the highest level groups. You can view the disks and lower level groups constituting their lower level groups in the Group Configuration screen.

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: 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: Main Screen (for SDX Objects of a Single Disk)

The screenshot displays the Global Disk Services (GDS) configuration interface. The window title is "Global Disk Services" and the interface includes a menu bar with "General", "Settings", "Operation", "View", and "Help".

The main area is divided into two panes:

- Configuration Tree:** A hierarchical tree view showing the configuration structure. The selected path is `linux04 > class0002 > disksd0010`. Other visible nodes include `class0001`, `group0001`, `group0003`, `class0002`, `Mastergrp001`, `Proxygrp`, and `linux05`.
- Details:** A pane showing the configuration details for the selected object, `linux04:class0002.disksd0010`. It contains two tables:

Volume Information Table:

Volume Name	Status	Class Name	Size	JRM	Physical Slice	Type	Master
volume0004	stop	class0002	0.25 GB	on	on	single	Mstervol2
volume0005	active	class0002	0.25 GB	on	on	single	-----

Disk Information Table:

Disk Name	Status	Physical Disk	Group Name	Disk Type
disksd0010	enabled	sdaj	-----	single

At the bottom of the interface, there is a "Log Information" section which is currently empty. The window footer indicates it is a "Java Applet Window".

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: Main Screen (for SDX Objects of a Volume)

The screenshot displays the 'Global Disk Services' interface. The title bar includes 'PRIMECLUSTER Global Disk Services'. The menu bar contains 'General', 'Settings', 'Operation', 'View', and 'Help'. The interface is divided into several sections:

- Configuration Tree:** A hierarchical tree view on the left. The selected path is 'linux04 > class0001 > group0001 > volume0001'. Other visible nodes include 'class0002', 'Mastergrp001', 'Proxygrp', and 'disksd0010'.
- Details:** A tabbed area on the right showing details for the selected volume: 'linux04:class0001.group0001.volume0001'. It contains two tables:
 - Slice Information:**

Slice Name	Status	Copying Status	Copy Type
group0001.volume0001	active	----	-----
 - Disk Information:**

Disk Name	Status	Physical Disk	Group Name	Disk Type
disksd0001	enabled	sdaa	group0001	stripe
disksd0002	enabled	sdab	group0001	stripe
- Log Information:** A section at the bottom of the details pane, currently empty.

The bottom of the window indicates it is a 'Java Applet Window'.

5.3.1.2 Viewing Proxy Object Configurations

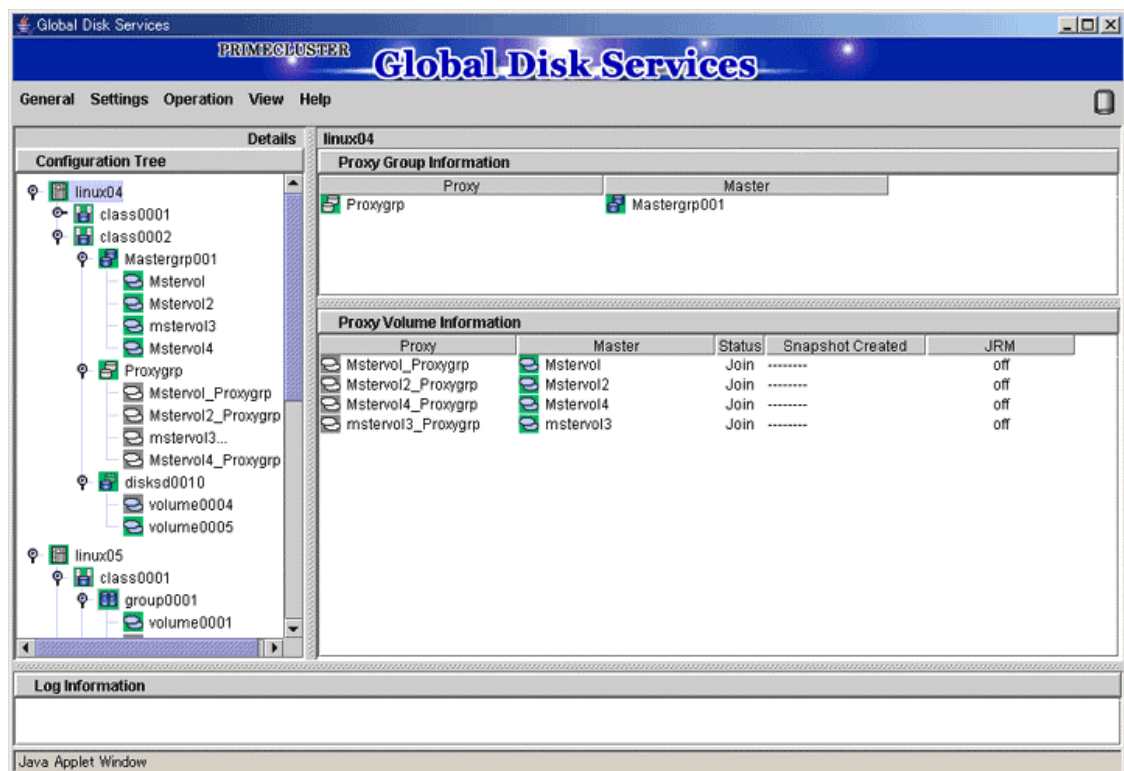
Select [Proxy Object] in the [View]:[Details] menu to view object configurations according to the following units.

- Proxy object configuration within a node
- Proxy object configuration within a class
- Proxy object configuration related to a group
- Proxy object configuration within a single disk
- Proxy object configuration related to a volume

Proxy object configuration within a node

Click the node icon in the Configuration Tree field, and all proxy groups and proxy volumes within the specified node appear.

Figure: 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: Main Screen (for Proxy Objects of a Class)

The screenshot displays the 'Global Disk Services' application window. The title bar includes 'PRIMECLUSTER' and 'Global Disk Services'. The menu bar contains 'General', 'Settings', 'Operation', 'View', and 'Help'. The main interface is divided into several sections:

- Configuration Tree:** A hierarchical tree on the left showing the system structure. The selected node is 'linux04: class0002'. Other nodes include 'linux04', 'class0001', 'class0002', 'Mastergrp001', 'Mstervol', 'Mstervol2', 'mstervol3', 'Mstervol4', 'Proxygrp', 'Mstervol_Proxygrp', 'Mstervol2_Proxygrp', 'mstervol3...', 'Mstervol4_Proxygrp', 'disksd0010', 'volume0004', 'volume0005', 'linux05', 'class0001', 'group0001', and 'volume0001'.
- Details:** A tabbed view showing information for the selected class 'linux04: class0002'.
 - Proxy Group Information:** A table with two columns: 'Proxy' and 'Master'. It lists 'Proxygrp' and 'Mastergrp001'.
 - Proxy Volume Information:** A table with columns: 'Proxy', 'Master', 'Status', 'Snapshot Created', and 'JRM'. It lists proxy volumes and their associated master volumes.
- Log Information:** A section at the bottom for viewing logs.

Proxy	Master	Status	Snapshot Created	JRM
Mstervol_Proxygrp	Mstervol	Join	-----	off
Mstervol2_Proxygrp	Mstervol2	Join	-----	off
Mstervol4_Proxygrp	Mstervol4	Join	-----	off
mstervol3_Proxygrp	mstervol3	Join	-----	off

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: Main Screen (for Proxy Objects of a Group)

The screenshot displays the 'Global Disk Services' configuration window. The title bar includes 'PRIMECLUSTER Global Disk Services'. The interface is divided into several sections:

- Configuration Tree:** A hierarchical tree on the left showing the system structure. The selected path is 'linux04 > class0002 > Mastergrp001'. Other visible nodes include 'linux04', 'class0001', 'class0002', 'Mstervol', 'Mstervol2', 'mstervol3', 'Mstervol4', 'Proxygrp', 'Mstervol_Proxygrp', 'Mstervol2_Proxygrp', 'mstervol3...', 'Mstervol4_Proxygrp', 'disksd0010', 'volume0004', 'volume0005', 'linux05', 'class0001', 'group0001', and 'volume0001'.
- Details:** The main area on the right, titled 'linux04:class0002.Mastergrp001'. It contains two sub-sections:
 - Proxy Group Information:** A table with two columns: 'Proxy' and 'Master'. It lists 'Proxygrp' and 'Mastergrp001'.
 - Proxy Volume Information:** A table with columns: 'Proxy', 'Master', 'Status', 'Snapshot Created', and 'JRM'. It lists four proxy volumes and their corresponding master volumes.
- Log Information:** A section at the bottom for viewing logs.

Proxy	Master	Status	Snapshot Created	JRM
Mstervol_Proxygrp	Mstervol	Join	-----	off
Mstervol2_Proxygrp	Mstervol2	Join	-----	off
Mstervol4_Proxygrp	Mstervol4	Join	-----	off
mstervol3_Proxygrp	mstervol3	Join	-----	off

Java Applet Window

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: Main Screen (for Proxy Objects of a Single Disk)

The screenshot displays the Global Disk Services (GDS) configuration window. The title bar includes 'PRIMECLUSTER Global Disk Services'. The interface is divided into several sections:

- Configuration Tree:** A hierarchical tree on the left showing the configuration structure. The selected path is `linux04 > class0002 > disksd0010`.
- Details:** A tabbed view on the right showing configuration details for the selected path. It includes:
 - Proxy Group Information:** A table with columns 'Proxy' and 'Master'. The 'Proxy' column contains 'volume0004' and the 'Master' column contains 'Mstervol2'.
 - Proxy Volume Information:** A table with columns 'Proxy', 'Master', 'Status', 'Snapshot Created', and 'JRM'. The 'Proxy' column contains 'volume0004', 'Master' contains 'Mstervol2', 'Status' contains 'Join', 'Snapshot Created' contains '-----', and 'JRM' contains 'off'.
- Log Information:** A section at the bottom for viewing logs.

The bottom of the window indicates it is a 'Java Applet Window'.

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: Main Screen (for Proxy Objects of a Volume)

The screenshot displays the Global Disk Services (GDS) configuration interface. The main window title is "Global Disk Services" and the subtitle is "PRIMECLUSTER". The interface is divided into several sections:

- Configuration Tree:** A hierarchical tree view on the left showing the configuration structure. The selected path is `linux04 > class0002 > Mastergrp001 > Mstervol`.
- Details:** A tabbed view showing configuration details for the selected volume. The current tab is "Proxy Volume Information".
- Proxy Volume Information:** A table showing the relationship between proxy and master volumes.
- Slice Information:** A table showing the status and copying status of slices within the volume.
- Log Information:** A section for viewing logs, currently empty.

Proxy Volume Information Table:

Proxy	Master	Status	Snapshot Created	JRM
Mstervol_Proxygrp	Mstervol	Join	-----	off

Slice Information Table:

Slice Name	Status	Copying Status	Copy Type
disksd0007.Mstervol	active	----	----
disksd0008.Mstervol	active	----	----
disksd0009.Mstervol_Proxygrp	copy	0.20 GB / 0.25 GB	OPC

5.3.1.3 Monitoring Object Status

You can monitor the object status from the Main screen. Object status will be updated at intervals specified in [Change Monitoring Intervals] in the [General] menu. You may also use [Update] in the [View] menu to update the status immediately. When an abnormality is detected with objects, a warning lamp (yellow/red) will flash. Clicking the flashing lamp will change the lamp to a lit-up lamp.



See

For details on warning lamps, see "4.1 Screen Configuration."

When the object status changes such as by detection of the failed disk, the icon color and description in the status field will change as well.



See

For details on icons, see "4.3 Icon Types and Object Status."

If a disk unit fails during operation, an icon of the disk in which an error is detected will turn red. Perform recovery work, following procedures described in "5.3.4 Disk Swap."

Figure: Object Status Monitoring in Main Screen

The screenshot shows the 'Global Disk Services' application window. The title bar includes 'PRIMECLUSTER Global Disk Services'. The menu bar has 'General', 'Settings', 'Operation', 'View', and 'Help'. A yellow warning lamp icon is visible in the top right corner of the window.

The main area is divided into two panes. The left pane shows a 'Configuration Tree' with a tree view of objects including 'linux04', 'class0001', 'group0001', 'group0003', 'class0002', 'Mastergrp001', and various 'Mstervol' and 'Proxygrp' objects.

The right pane shows 'Details' for 'linux04:class0001.group0003'. It contains two tables:

Volume Information								
Volume Name	Status	Class Name	Size	JRM	Physical Slice	Type	Master	
volume0002	active	class0001	0.25 GB	on	on	mirror	-----	
volume0003	active	class0001	0.51 GB	on	on	mirror	-----	
volume0007	warning	class0001	0.50 GB	on	on	mirror	-----	

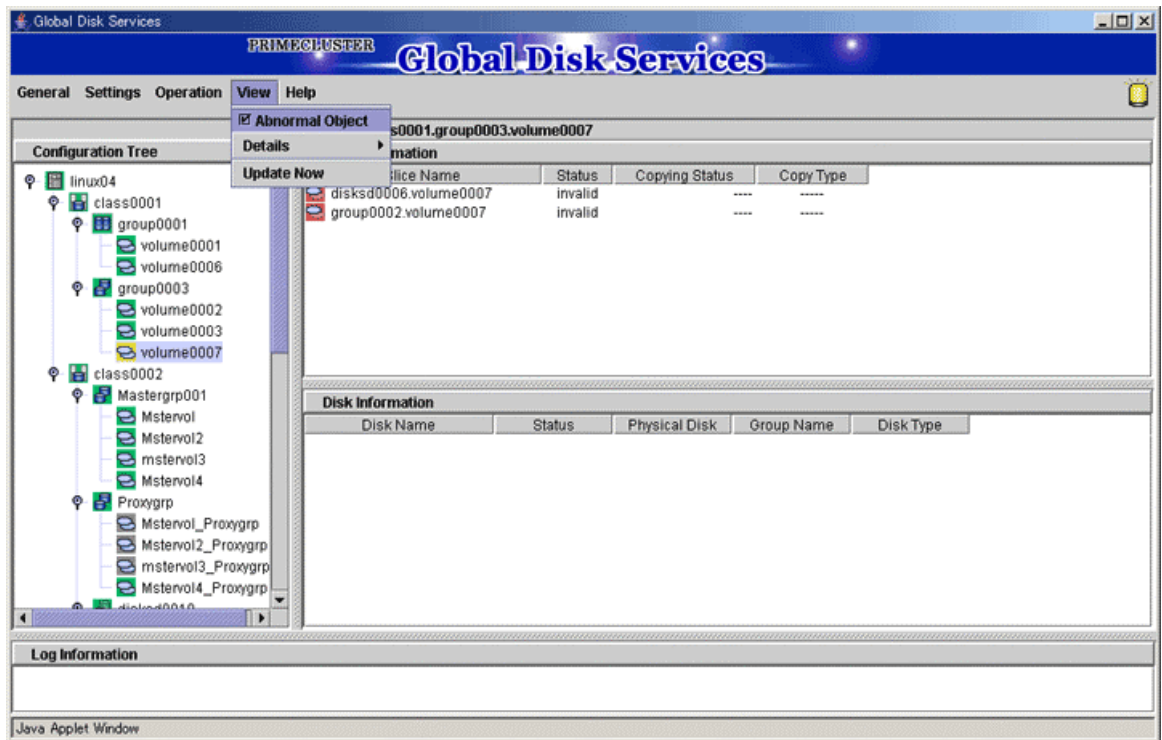
Disk Information					
Disk Name	Status	Physical Disk	Group Name	Disk Type	
disksd0005	enabled	sdae	group0003	mirror	
disksd0006	enabled	sdaf	group0003	mirror	
disksd0003	enabled	sdac	group0002	concat	
disksd0004	enabled	sdad	group0002	concat	

At the bottom of the main area is a 'Log Information' section, which is currently empty.

The bottom status bar indicates 'Java Applet Window'.

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: Main Screen when [View]: [Abnormal Object] is selected



Note **Abnormality Detected with GDS Management**

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

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

5.3.2 Backup

GDS is software that provides a highly reliable system environment at times of failure, allowing you to continue normal service.

However, using a mirrored system environment does not guarantee absolute safety.

Computer problems could occur due to hardware, software, operation error and poor environment.

A reliable method for working around to such problems is creating "backups." To minimize the damage caused by trouble, periodical backup is strongly recommended.



Information **Backing Up and Restoring a System Disk [Itanium]**

Among system disk volumes, volumes (such as /opt and /home) other than system volumes (/ , /usr, /var and swap area) can be backed up by following the procedures described in this section. For system volume backup and restore, see "6.1 Backing Up and Restoring a System Disk [Itanium]" and "6.2 Backing Up and Restoring a System Disk through an Alternative Boot Environment [Itanium]."

5.3.2.1 Backup (by Slice Detachment)

This section describes the procedure for creating backup of a mirror volume making use of a snapshot by detaching a slice. This method requires the following "slice operations."

- Detach Slice
- Attach Slice



See

- For snapshot by slice detachment, see "1.3.8 Snapshots by Slice Detachment."
- For the operation flow, see "5.1.3.1 Backup (by Slice Detachment)."



Note **Use Conditions on Snapshot by Slice Detachment**

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

This snapshot is also 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 "5.1.4 Restore."

Slice Detachment

In order to create a backup, you must temporarily exclude one of the slices from volume, and make it accessible as a separate volume.

The procedures are explained below.

1. Displaying the volume status including the slice

In the Main screen, display the volume containing the slice for which you want to create a backup. Click the icon and [Slice Information] appears.

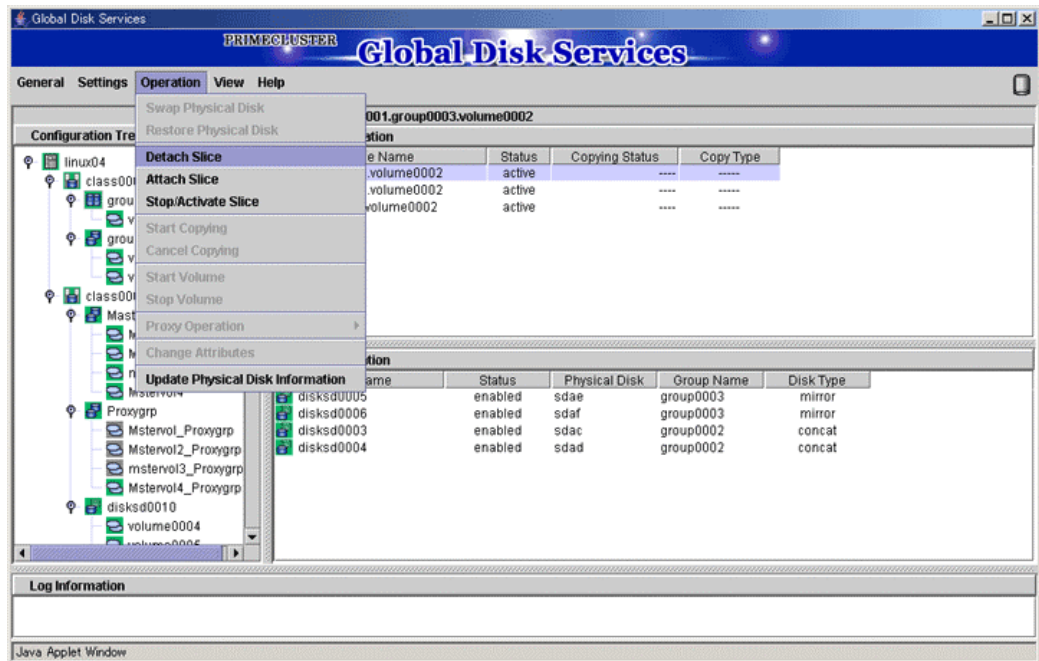
2. Selecting the slice to detach

In the Slice List, select the slice you want to detach by clicking its icon.

3. Selecting [Detach Slice] menu


In the Main screen [Operation] menu, select [Detach Slice].

Figure: Detach Slice



4. **Setting the environment of the detaching slice**
We will set the environment of the detaching slice.

Figure: Setting the environment of the detaching slice



a) **Access Mode**

Set the access mode of the detached slice.

The initial value is "Read and write possible."

When you specify "Read only possible," the detached mirror slice will be available for read only. Opening a read-only slice in write mode will result in an error.

Click <OK> after setting the environment. If you click <Cancel>, the detaching of the slice will be cancelled.




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: 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: 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.

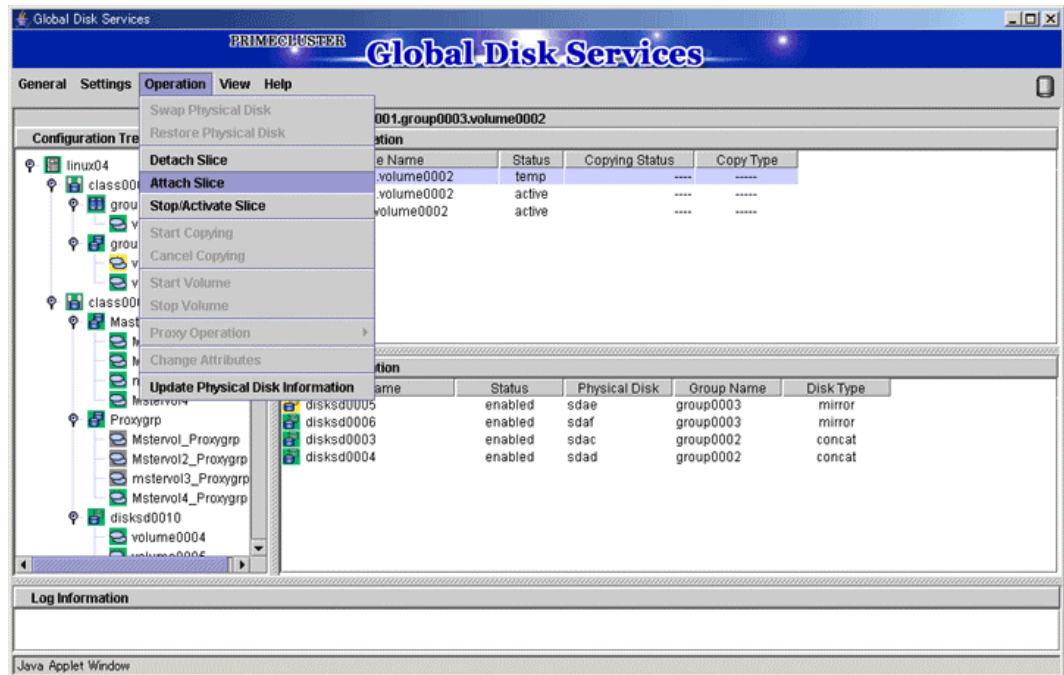
1. Selecting the slice to attach

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: Attach Slice



The confirmation screen shown below appears.

To continue the process, click <Yes>. Clicking <No> cancels the slice attachment process.

Figure: Confirming "Attach Slice"



3. Notifying the completion of Attach Slice

Click <OK>, and a message notifying the completion of Attach Slice appears.

Figure: Notifying Attach Slice Completion



Reliable Backup Procedures

Although the procedure above allows you to use the volume as it is after preparing backup, reliability is affected since one of the slices is excluded.

To ensure reliability, temporarily add a backup disk and perform synchronization copying. After its completion, exclude the slice.

For safe and reliable backup, follow the procedures below.

1. Register backup disk with class.
2. Connect backup disk to group.
3. After completing synchronization copying, stop service.
4. Detach slice.
5. Resume service.
6. Perform backup using access path for backup.
7. Attach slice.
8. Disconnect backup disk from group.
9. Remove backup disk from class.

Completing procedures 1 and 2 in advance will save the waiting time for synchronization copying in procedure 3, therefore reducing the time required for backup.



Note Perform [Detach Slice] after Stopping Services

To ensure integrity of backup data, always stop service before excluding the slice.

You may resume service once [Detach Slice] is complete. You do not have to suspend service during the actual backup process.

When excluding the slice without stopping service, run host programs such as fsck (in case of file system) as necessary.

Stop/Activate Slice

Stop Slice


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

1. **Selecting the slice to stop**
In the Slice Information Field, select the slice you want to stop by clicking the "temp" status slice icon.

2. Selecting [Stop/Activate Slice] menu

In the Main screen [Operation] menu, select [Stop/Activate Slice].

Figure: 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: Activate Slice



To continue the process, click <Yes>. Clicking <No> cancels the slice activation process.

5.3.2.2 Backup (by Synchronization)

This subsection describes the procedures for backing up volumes through use of snapshots of GDS Snapshot by synchronization. This method requires the following "Proxy Operations."

- Join
- Part
- Rejoin
- Break



See

- For snapshot by synchronization, see "1.5.1 Snapshot by Synchronization."
- For the operation flow, see "5.1.3.2 Backup (by Synchronization)."



Point Keep Proxies Parted If Possible

If proxies are kept parted, they can be used for master data restoration. Therefore, it is recommended to keep proxies parted if possible. For the restore procedures see "5.3.3 Restore."



Note "Part" Proxies after Stopping Services

To ensure integrity of backup data, always stop services before executing "Part." You may resume the services once "Part" is complete. It is unnecessary to stop services when backing up data to tape and so on.

For details, see "A.2.20 Ensuring Consistency of Snapshot Data."



Note Snapshot by Synchronization Use Conditions

See the following sections for points of concern.

- "A.1.8 Exception to Proxy Configuration"
- "A.1.9 Number of Proxy Volumes"
- "A.1.10 Proxy Volume Size"
- "A.1.11 Proxy Group Size"



Information Use Conditions for Copy Functions of Disk Units

See "A.2.16 Using the Advanced Copy Function in a Proxy Configuration" and "A.2.19 Using EMC TimeFinder or EMC SRDF in a Proxy Configuration [Linux2.6]."

Join

Join a volume for backup (a proxy volume) to a volume to be backed up (a master volume).

To back up all the volumes within a group simultaneously, join a group for backup (a proxy group) to a group to be backed up (a master group).

For the "Join" procedures see "5.2.4.1 Join."

Part

Make sure that synchronization copying from a master to a proxy is complete in the Main screen, and part the proxy from the master.

Follow the procedures below.

1. Selecting a proxy to be parted

Click a master volume icon on the GDS Configuration Tree in the Main screen.

To back up all the master volumes within a master group simultaneously, click the master group icon.

Select [View]:[Details]:[Proxy Object] to view all the proxy objects related to the master object selected on the GDS Configuration Tree, in the Object Information Field. In this field, click an icon of a proxy volume or a proxy group used for backup.



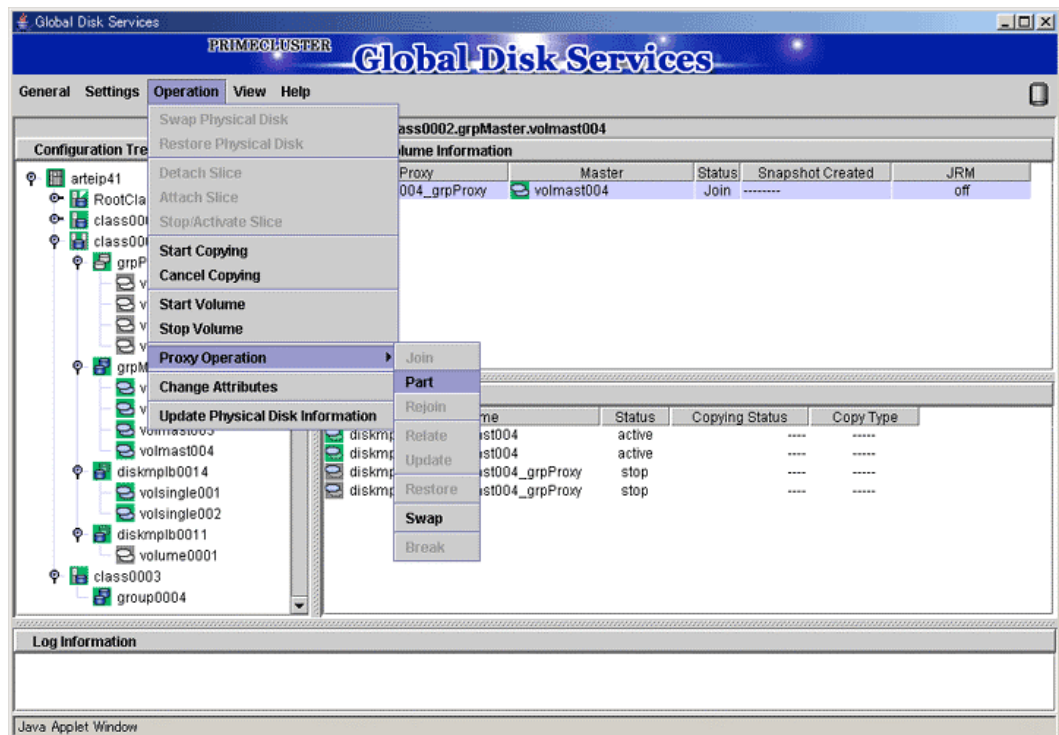
Note Proxy Objects That Can Be Parted

Proxy volumes can be parted if they are joined and if copy is complete.

2. Selecting the [Part] menu

Select [Operation]:[Proxy Operation]:[Part].

Figure: Part



3. Setting the environment for parting the proxy

The Part Proxy dialog box appears.

Figure: 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.2.16 Using the Advanced Copy Function in a Proxy Configuration"
- "A.2.17 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.2.12 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: 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.

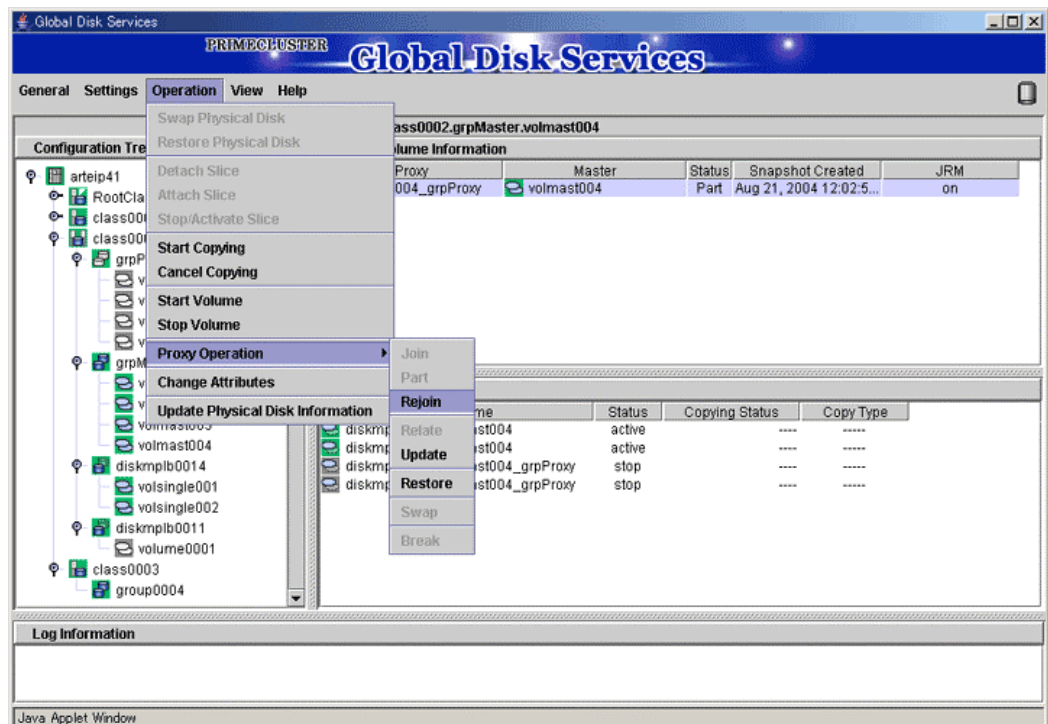
To back up all the master volumes in a master group simultaneously, click the master group icon.

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: Rejoin



3. Confirming

A confirmation dialog box appears asking you whether to rejoin the proxy.

Figure: 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: Information Message of "Rejoin" Completion



Click <OK> to close the information message window.

Break

If no more backup is necessary, break the relationship between the master and the proxy. For "Break Proxy" procedures see "5.5.6 Breaking a Proxy."

5.3.2.3 Backup (by OPC)

This subsection describes the procedures for backing up volumes through use of snapshots of GDS Snapshot by OPC. This method requires the following "Proxy Operations."

- Relate
- Update
- Break



See

- For snapshot by OPC, see "1.5.3 Instant Snapshot by OPC."
- For the operation flow, see "5.1.3.3 Backup (by OPC)."



Note "Update" Proxies after Stopping Services

To ensure integrity of backup data, always stop services before executing "Update." You may resume the services once "Update" is complete. It is unnecessary to stop services when backing up data to tape and so on.

For details, see "A.2.20 Ensuring Consistency of Snapshot Data."



Note Instant Snapshot by OPC Use Conditions

See the following sections for points of concern.

- "A.1.8 Exception to Proxy Configuration"
- "A.1.9 Number of Proxy Volumes"
- "A.1.10 Proxy Volume Size"
- "A.1.11 Proxy Group Size"
- "A.2.16 Using the Advanced Copy Function in a Proxy Configuration"
- "A.2.17 Instant Snapshot by OPC"

Relate

Relate a volume for backup (a proxy volume) to a volume to be backed up (a master volume).

To back up all the volumes within a group simultaneously, relate a group for backup (a proxy group) to a group to be backed up (a master group).

For the "Relate Proxy" procedures see "5.2.4.2 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.



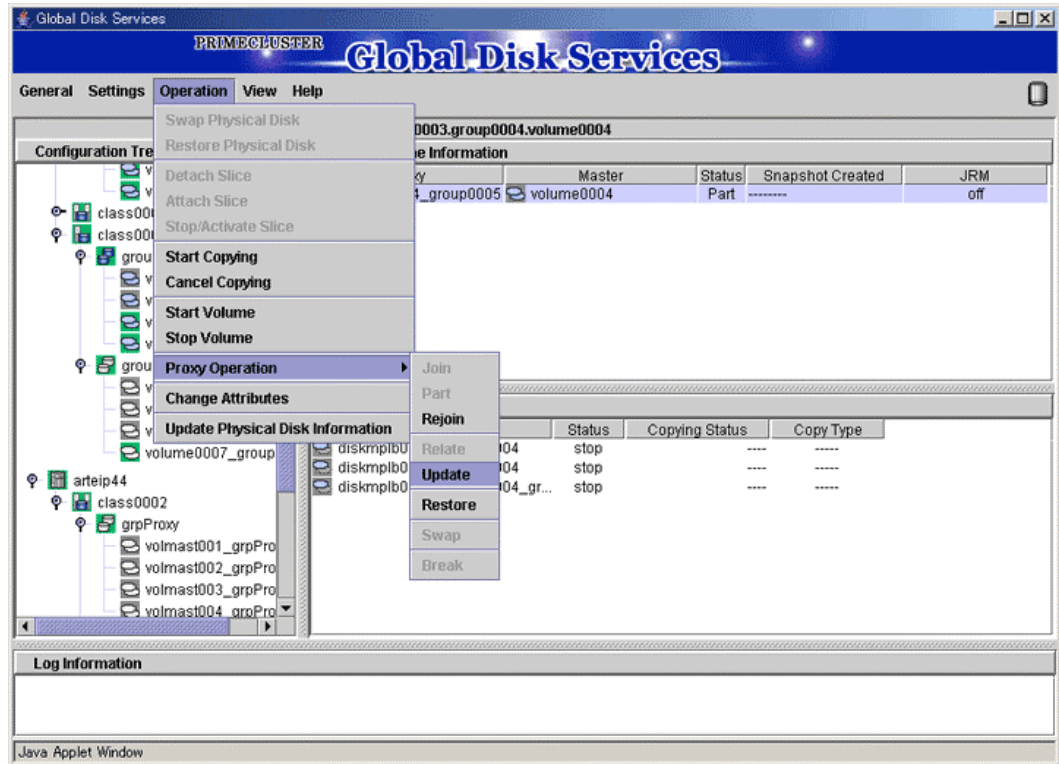
Note Proxy Objects That Can Be Updated

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.

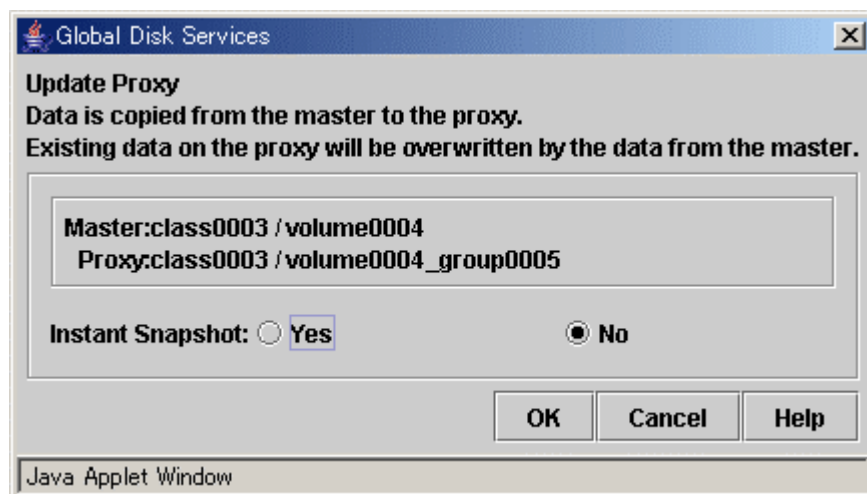
Figure: Update



3. Setting the environment for updating the proxy

The Update Proxy screen appears.

Figure: 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: 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: Information Message of "Update" Completion (Yes to Instant Snapshot)



If "Yes" to [Instant Snapshot] was selected when setting the update environment in step 3, you may start the proxy volume and execute backup immediately without waiting until copying is complete.

Break

If no more backup is necessary, break the relationship between the master and the proxy.

For "Break Proxy" procedures see "5.5.6 Breaking a Proxy."

5.3.3 Restore

This subsection describes the procedures for restoring volumes through use of GDS Snapshot proxy volumes. This method requires the following "Proxy Operations."

- Restore

Stop services using a volume to be restored (a master volume) to stop the master volume and then perform the following procedures.



See

- For restore with proxy volumes, "1.5.4 Instant Restore."
- For the operation flow, see "5.1.4 Restore."



Note System Volume Restoration [Itanium]

The system volumes currently running as file systems such as /, /usr, and /var cannot be stopped, and such volumes cannot be restored through this procedure. For the system volume restoration methods, see "6.1 Backing Up and Restoring a System Disk [Itanium]" or "6.2 Backing Up and Restoring a System Disk through an Alternative Boot Environment [Itanium]."

Restore

Copy (overwrite) data from a proxy to a master.
Follow the procedures below.

1. Selecting a proxy as a restore copy source

Click an icon of a master volume to be restored on the GDS Configuration Tree in the Main screen.

To restore all the master volumes within a master group simultaneously, click the master group icon.

Select [View]:[Details]:[Proxy Object] to view all the proxy objects related to the master object selected on the GDS Configuration Tree, in the Object Information Field. In this field, click an icon of a proxy volume (or a proxy group) as a restore copy source.



Point Backup Generation

Proxy volume data is a replica of master volume data at the moment of snapshot creation. For snapshot creation time, check [Snapshot Created] time in the proxy volume information field.



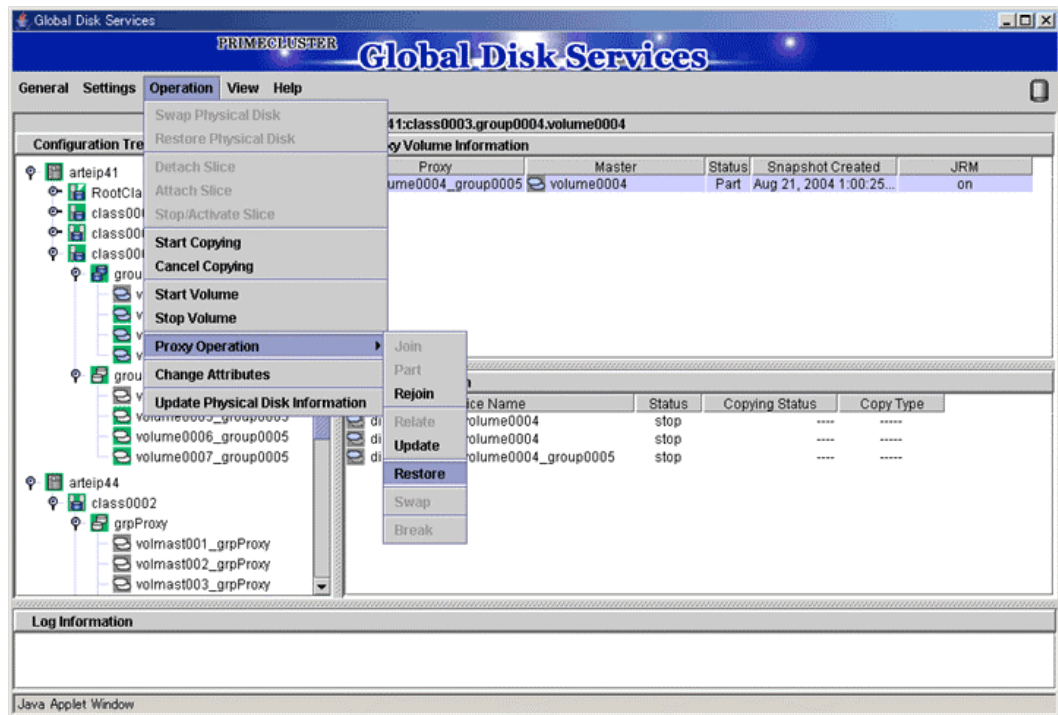
Note Proxy Volumes That Can Be Restore Copy Sources

Proxy volumes can be restore copy sources if they are parted and in "active" or "stop" status. However, when selecting "Yes" to "Rejoin" for setting the restore environment in step 3, copy sources must be proxy volumes that are parted and are in "stop" status.

2. Selecting the [Restore] menu

Select [Operation]:[Proxy Operation]:[Restore] in the Main screen.


Figure: Restore



3. Setting the environment for restoring the master

The Restore Master screen appears.

Figure: Setting the Restore Environment



Set the environment for restoring the master.

a) Rejoin

Specify whether to rejoin the master and the proxy.

The default value is "Yes." The master and the proxy will be joined, and after copy is complete they will be synchronized. If the OPC function is unavailable, select "Yes."

To leave the master and the proxy parted and execute restore with the OPC function, set this option to "No."



[See](#)

For the OPC function use conditions, see the following sections.

- "A.2.16 Using the Advanced Copy Function in a Proxy Configuration"
- "A.2.17 Instant Snapshot by OPC"



Note Master Volumes with Mirroring Multiplicity of Two and Higher

Even if the device supports the OPC function, selecting "Yes" to [Rejoin] disables the OPC function. To use the OPC function for restoration, select "No" to [Rejoin]. However, slices other than the OPC copy destinations will be excluded from mirroring and the data statuses will be invalid. To restore the master volume mirroring status, select the master volume and execute [Operation]:[Start Copying] in the Main screen. If not executing [Start Copying], resynchronization copying automatically starts when the master volume starts.

b) Instant Restore

Specify whether to apply instant restore

The default value is "No." Wait until copying from the proxy volume to the master volume is complete and start the master volume.

To start the master volume immediately without waiting until copying from the proxy volume to the master volume is complete, set this option to "Yes."



Note If "Yes" to [Rejoin] and "Yes" to [Instant Restore] Are Selected

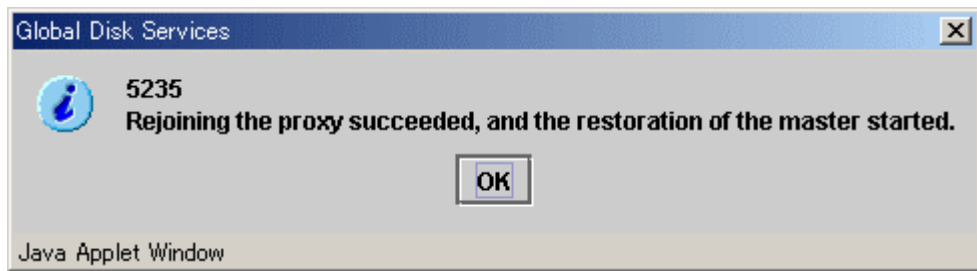
Even if synchronization copying from the proxy to the master is in progress, the master volume can be started and accessed. Note, however, that the master and the proxy are joined and data written to the master is also written to the proxy. To prevent proxy data from being updated, wait until copying is complete and execute "Part" before starting the master volume.

After the settings are complete, click <OK>. Clicking <Cancel> cancels the restore process.

4. Information message of the completion


A message window appears informing you that the restore process is complete.

Figure: 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: 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: 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: Information Message of "Restore" Completion (No to Rejoin/Yes to Instant Restore)



If "No" to [Rejoin] and "Yes" to [Instant Restore] were selected when setting the restore environment in step 3, the master and the proxy are left parted. You may start the master volume immediately without waiting until OPC copying from the proxy to the master is complete. If the OPC function is unavailable, the restore process fails.

Break

If no more backup is necessary after the restore, break the relationship between the master and the proxy.

For "Break Proxy" procedures see "5.5.6 Breaking a Proxy."

5.3.4 Disk Swap

When a disk unit abnormality occurs, contact your local customer support to swap the disk units. When using GDS, 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 "F.1.12 Disk Unit Error."



Note Notes on Physical Disk Swapping

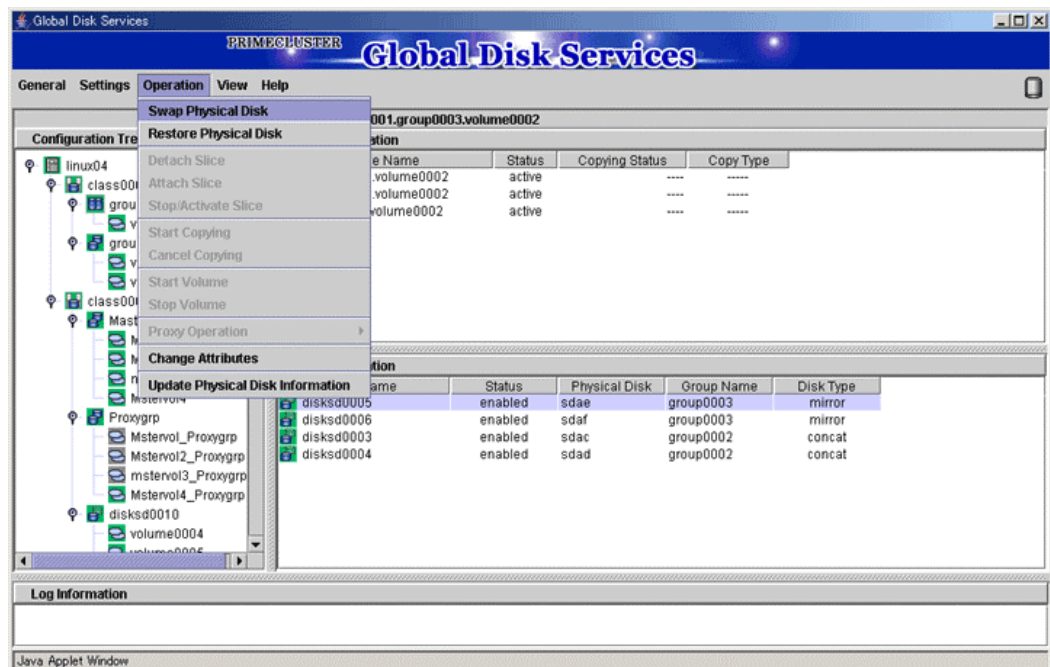
See "A.2.14 Swapping Physical Disks."

Swap Physical Disk

In order to swap the disk units, you must take the physical disk offline. The procedures are explained below.


1. **Displaying the status of physical disk**
In the Main screen, display the physical disk to be swapped. Click the icon and select the physical disk.
2. **Selecting [Swap Physical Disk]**
In the Main screen [Operation] menu, select [Swap Physical Disk].

Figure: Swap Physical Disk



The confirmation screen shown below appears.
To continue the process, click <Yes>. Clicking <No> cancels the physical disk swapping process.

Figure: Confirming Swap Physical Disk



3. Requesting the swap of physical disks

Click <Yes>, and a message notifying the offline status appears.
Click <OK>, and request your local customer support to swap the disk unit.

Figure: Notifying Offline Completion



Restore Physical Disk

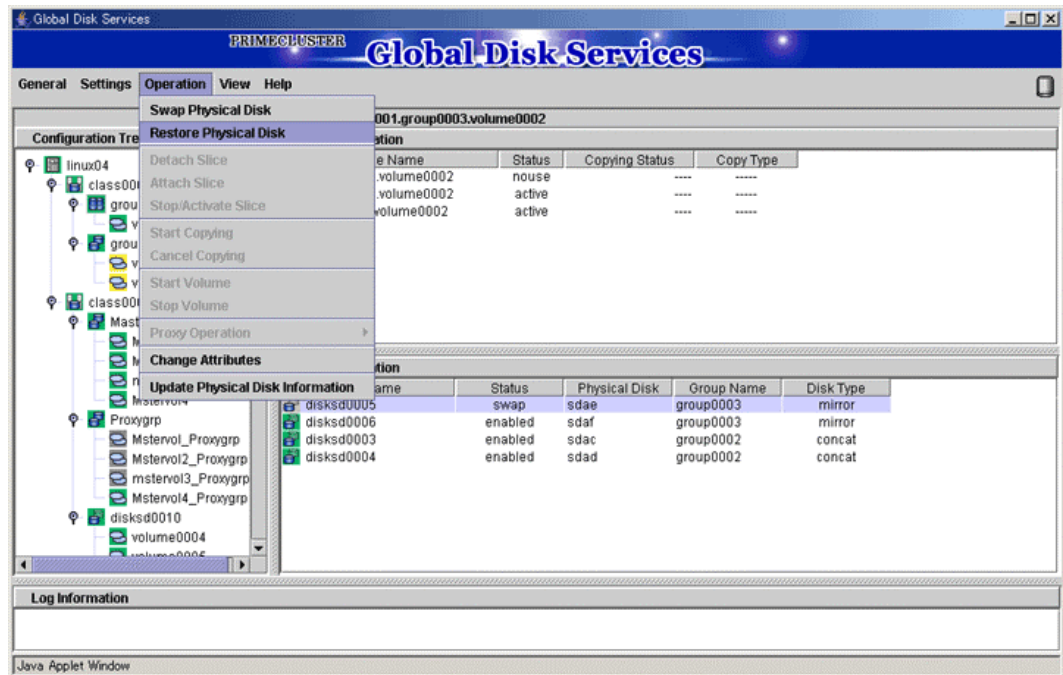
After swapping the disk units, you must put the swapped physical disk back online.
The procedures are explained below.

1. Selecting the physical disk to restore

Select the physical disk you want to restore.


2. **Selecting [Restore Physical Disk] menu**
In the Main screen [Operation] menu, select [Restore Physical Disk].

Figure: Restore Physical Disk




The confirmation screen shown below appears.
To continue the process, click <Yes>. Clicking <No> cancels the physical disk restore process.

Figure: Confirming Restore Physical Disk



3. **Notifying online status**
Click <OK>, and a message notifying the online status appears.

Figure: Notifying Online Completion



5.3.5 Disk Migration

This subsection describes the procedures for transferring volumes to other disks through use of GDS Snapshot proxy volumes. This method requires the following "Proxy Operations."

- Join
- Swap Slice
- Break



See

- For disk migration with proxy volumes, see "1.5.5 Online Disk Migration."
- For the operation flow, see "5.1.6 Disk Migration."



Note Proxy Volume Use Conditions

See the following sections for points of concern.

- "A.1.8 Exception to Proxy Configuration"
- "A.1.9 Number of Proxy Volumes"
- "A.1.10 Proxy Volume Size"
- "A.1.11 Proxy Group Size"

Join

Join a destination volume (a proxy volume) to a volume for disk migration (a master volume).

To perform disk migration on all the volumes within a group simultaneously, join a destination group (a proxy group) to a group for disk migration (a master group).

For "Join Proxy" procedures see "5.2.4.1 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).



Note Possible Proxy Objects for Swapping

"Swap Slices" is possible with proxy objects if they are joined and copy is complete.



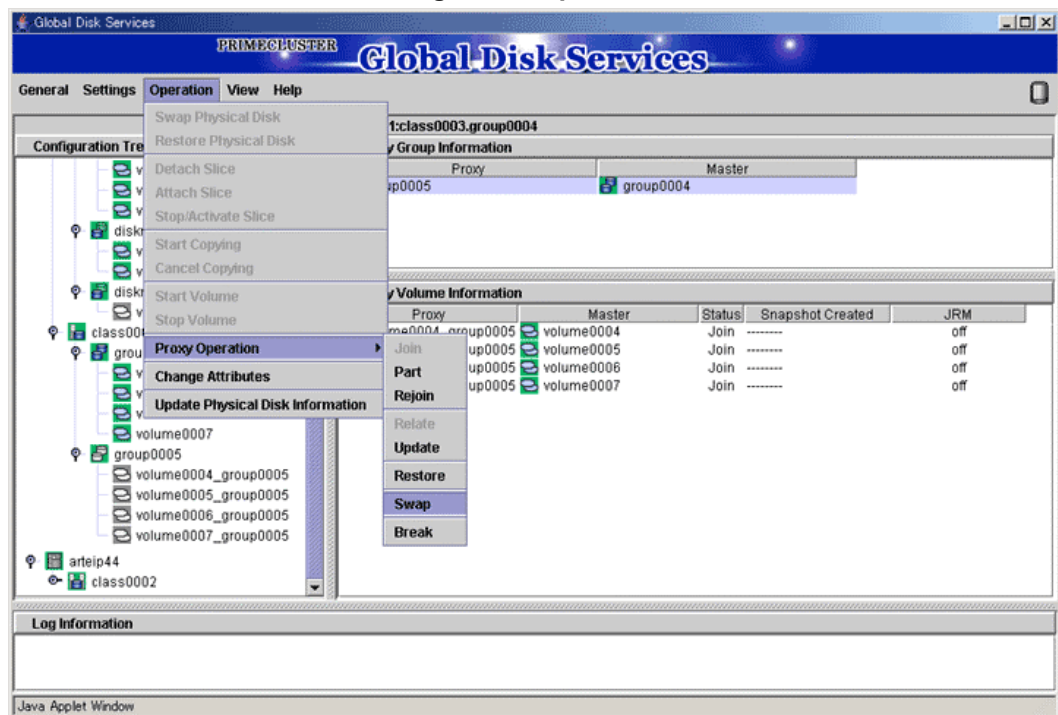
Note If There Is a Session by Disk Unit's Copy Function

If there is a session by disk unit's copy function between the master and the proxy, slice swapping fails. For existing sessions, use the `sdxinfo -S -e` long command and check the FUNC field of the results. If the FUNC field for any slice of the master or the proxy shows a value other than the asterisk (*), a session exists between the master and the proxy. In this event, cancel the session with the `sdxproxy Cancel` command to swap the slices. For details, see "D.6 sdxinfo - Display object configuration and status information" and "D.14 sdxproxy - Proxy object operations."

2. Selecting the [Swap Slice] menu

Select [Operation]:[Proxy Operation]:[Swap Slice] in the Main screen.


Figure: Swap Slice



3. Confirming

A confirmation screen appears asking you whether to swap the slices.

Figure: Confirming "Swap Slice"

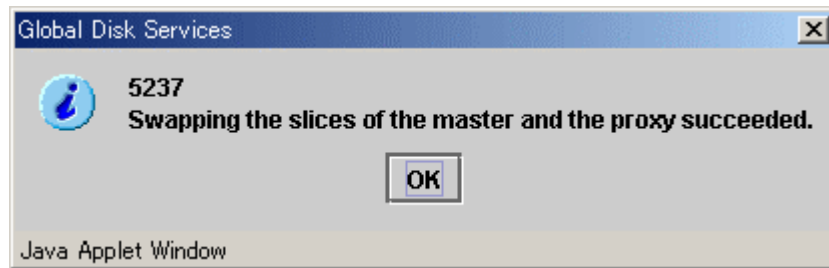


To continue the process, click <Yes>. Clicking <No> cancels the slice swapping process.

4. Information message of the completion

A message window appears informing you that the swap process is complete.

Figure: Information Message of "Swap Slice" Completion



Click <OK> to close the information message window.

Break

Break the relationship between the master and the proxy.

For "Break Proxy" procedures see "5.5.6 Breaking a Proxy."

5.3.6 Copying Operation

The Copying Operation function controls synchronization copying of mirror volumes. GDS provides the following copying operation.

- Start Copying
- Cancel Copying

Start Copying

Synchronization copying will be performed after attaching the slice that is in "invalid" or "copy-stop" status as a result of [Cancel Copying] operation.

For slices in "copy-stop" status, copying will resume from the point where copying was interrupted.

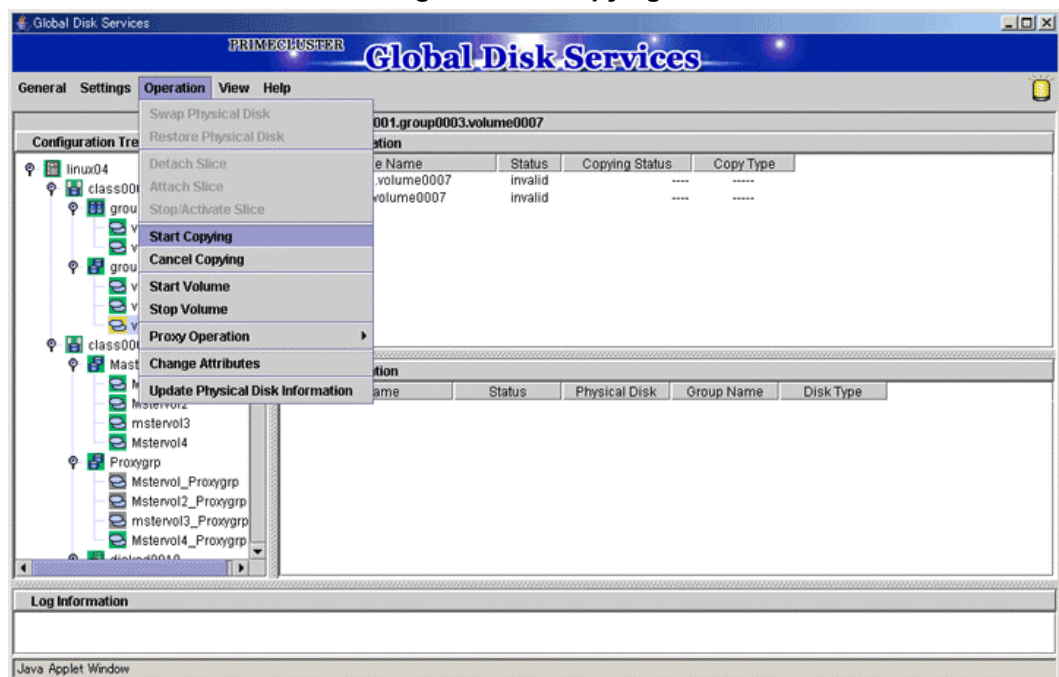
1. Select volume for synchronization copying

In the GDS Configuration tree field or Volume Information field, select the volume you want to copy by clicking the icon.

2. Selecting [Start Copying] menu


In the Main screen [Operation] menu, select [Start Copying].

Figure: Start Copying



The confirmation screen shown below appears.
To continue the process, click <Yes>. Clicking <No> cancels the copy start process.


Figure: Confirming Start Copying



3. [Start Copying] completion screen

Click <OK>, and the message notifying the [Start Copying] completion screen shown below appears.

Figure: 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.

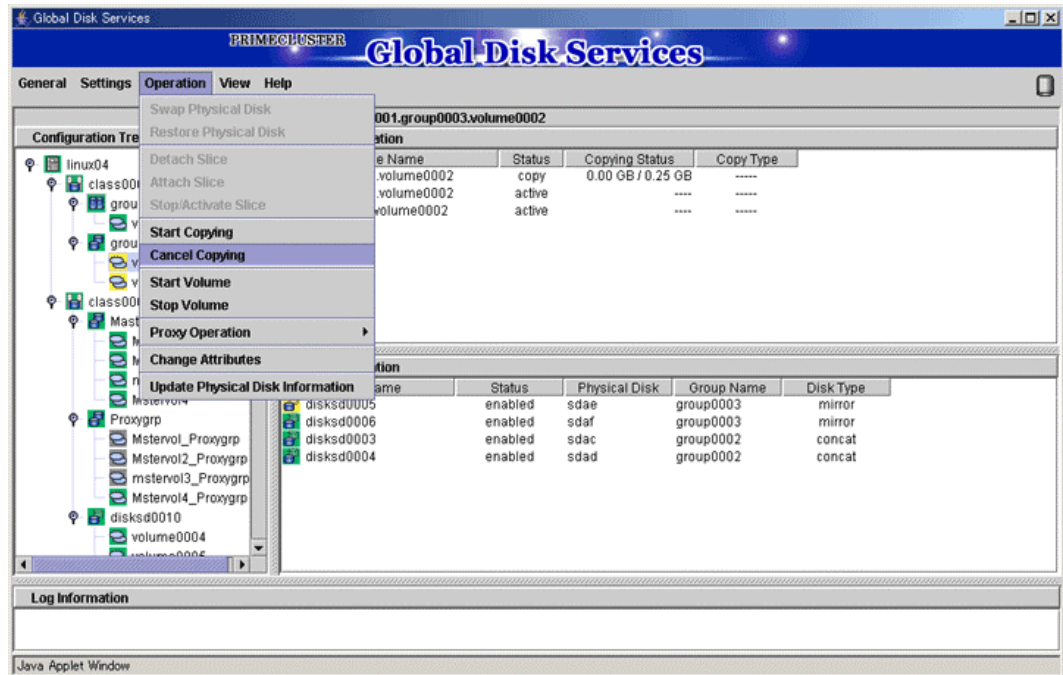
The slice whose copying has been stopped will be in "invalid" status. Perform [Start Copying] operation to resume its normal status

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: Cancel Copying



The confirmation screen shown below appears.
To continue the process, click <Yes>. Clicking <No> cancels the copy cancellation process.


Figure: Confirming Cancel Copying



3. **[Cancel Copying] completion screen**

Click <OK>, and the message notifying the [Cancel Copying] completion screen shown below appears.

Figure: Notifying Cancel Copying Completion



Note

[Cancel Copying] operation unavailable

When more than one volume is selected, [Cancel Copying] operation cannot be performed.

5.4 Changes

Class configurations, group configurations, and volume configurations can be changed through two types of operation: reconfiguration and attribute change. This section explains the changing procedures for each configuration.

5.4.1 Class Configuration

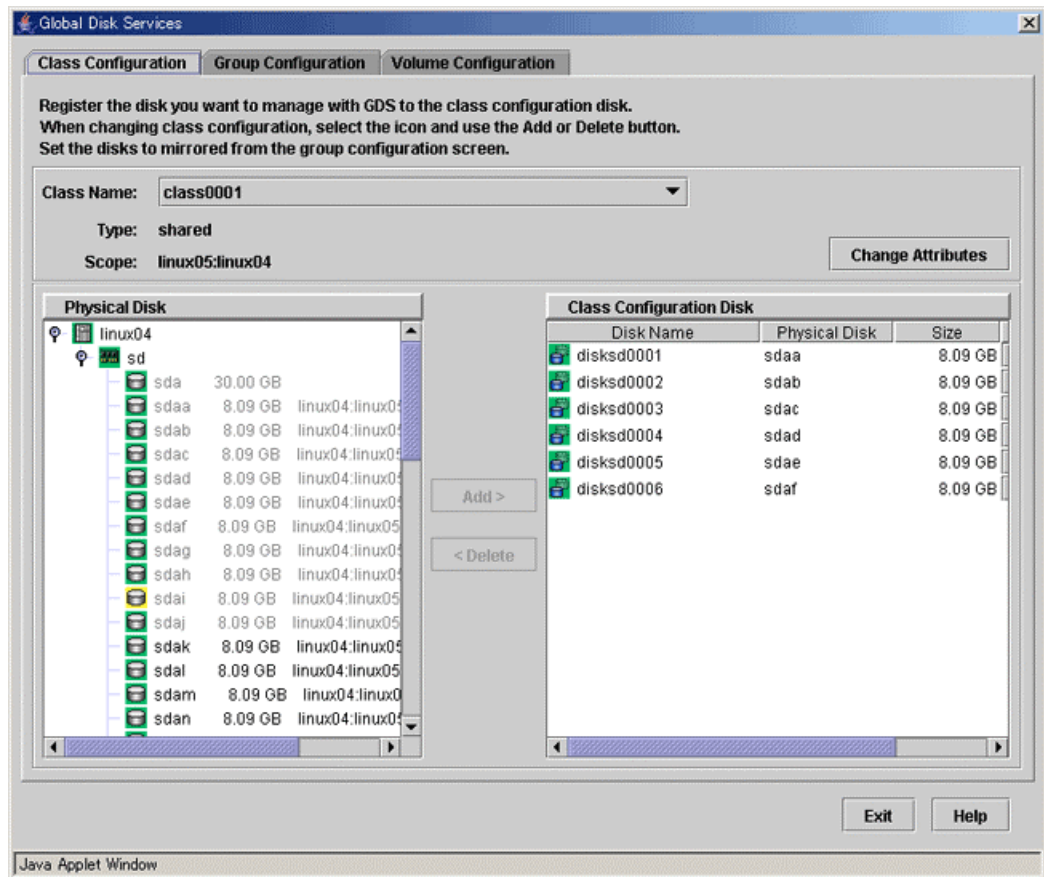
The procedures for changing class configuration are explained below.

Changing Class Configuration

1. Displaying Class Configuration screen

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

Figure: Class Configuration



2. Selecting class to change

Select the class you want to change from Class Configuration screen [Class Name].

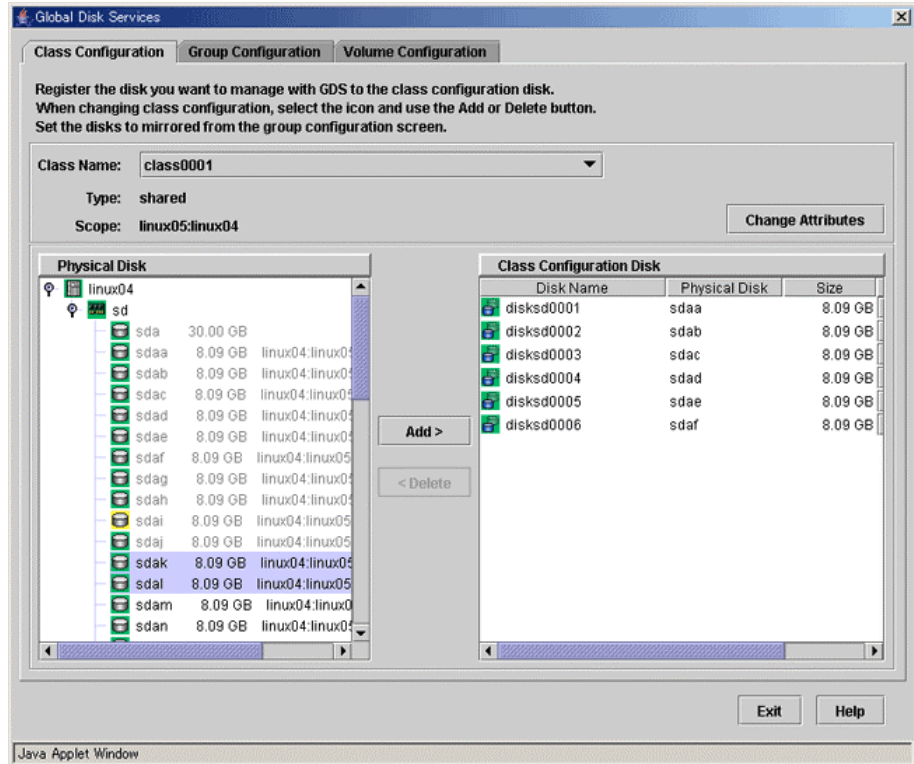
3. Changing configuration

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>.
- 3- 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: Physical Disk Registration



b) Removing a disk

- 1- Select the disk you want to remove from the [Class Configuration Disk] field.
- 2- Click <Delete>.

4. Completing the change

If you have no other changes to make, click <Exit>.

Changing the Class Attributes

Change the class attributes using the following procedures.



Note Preconditions for Changing Class Attributes

- When the class has active volumes, the following class attributes cannot be changed. To change these class attributes, first stop all the volumes within the class on all the nodes within the class scope.
 - Class name
 - Type (from “shared” to “local”)
 - Scope (node deletion)
- When the class has proxy objects, the class attributes cannot be changed. To change attributes of such a class, break the proxy objects within the class.

1. Invoking the Class Attributes Definition screen

Invoke the Class Attributes Definition screen using one of the following methods.

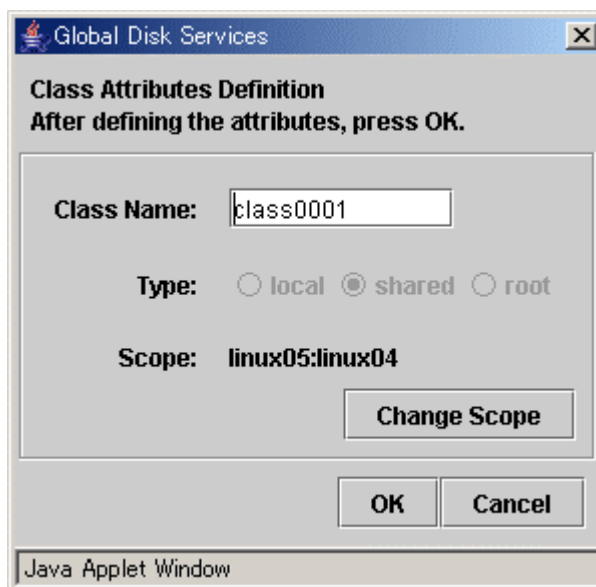
a) Operation menu in the Main screen

Click the target class's icon on the GDS configuration tree in the Main screen, and select [Operation]:[Change Attributes] to invoke the Class Attributes Definition screen.

b) Change Attributes button in the Class Configuration screen

Select the target class from the [Class Name] in the Class Configuration screen, and click the <Change Attributes> button to invoke the Class Attributes Definition screen.

Figure: Class Attributes Definition Screen



The screenshot shows a dialog box titled "Global Disk Services" with a close button (X) in the top right corner. The main title is "Class Attributes Definition" and the instruction below it is "After defining the attributes, press OK." The dialog contains the following fields and controls:

- Class Name:** A text input field containing "class0001".
- Type:** Three radio buttons labeled "local", "shared", and "root". The "shared" radio button is selected.
- Scope:** A text input field containing "linux05:linux04".
- Change Scope:** A button located below the Scope field.
- OK** and **Cancel:** Two buttons at the bottom of the dialog.

At the bottom of the dialog, there is a status bar that reads "Java Applet Window".

2. Changing attributes

a) For a single node

You can only change the class name.

b) For a cluster

You can change the following attributes:

- Class Name
- Type
- Scope



See

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

3. Implementing changes

Click <OK> to implement the attributes change, or <Cancel> to cancel.



Note **Changes in Special File Path Names**

Changing the class name will also change the special file path name used to access the volume.

Changing the Disk Attributes

When changing the disk attributes, there are two procedures to choose from:

- Using Main screen Operation menu.
- Using the Class Configuration Setting screen.



Note 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

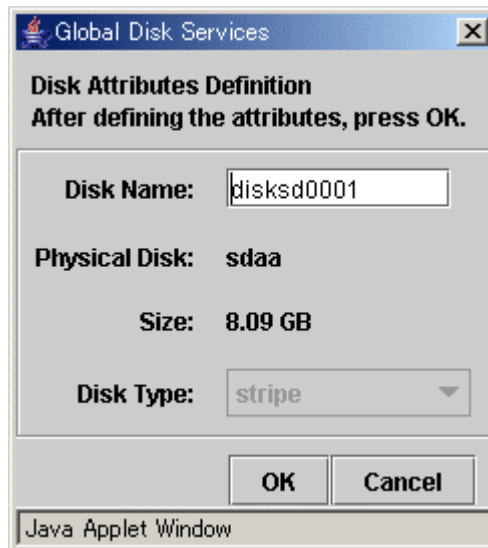
1. Selecting disk to change

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: Disk Attribute Definition Screen



3. Changing attributes

You can change the following attributes.

- Disk Name
- Disk Type



See

For details on assigning a disk name, see "A.1.1 Object Name."

4. Applying changes

Click <OK> to apply changes in the attributes change, or <Cancel> to cancel.

Using the Class Configuration screen

1. Displaying the disk to change attributes

Select the disk of which you want to change attributes in the [Class Configuration Disk] field.

2. Changing attributes

You can change the following attributes in the [Class Configuration Disk] field.

- Disk Name
Double-click the disk name and edit.
- Disk Type
Scroll to the right to get the [Disk Type] column, and select the disk type from the list.



See

For details on assigning a disk name, see "A.1.1 Object Name."

5.4.2 Group Configuration

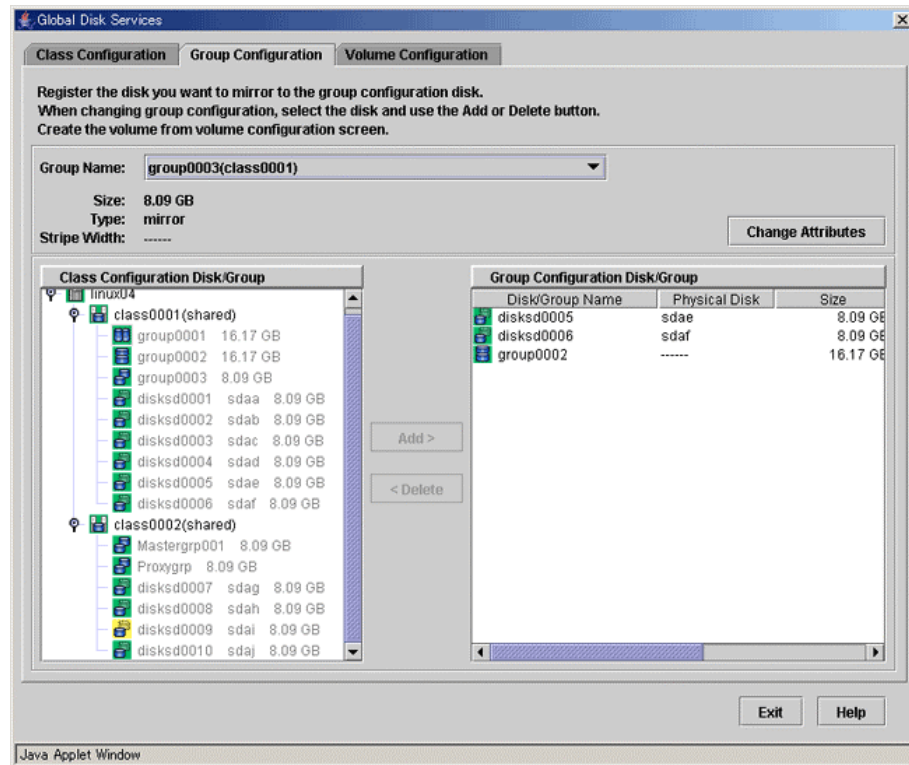
The procedures for changing group are explained below.

Changing Group Configuration

1. Displaying Group Configuration screen

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

Figure: Group Configuration



2. Selecting group to change

Select the group you want to change from Group Configuration screen [Group Name].

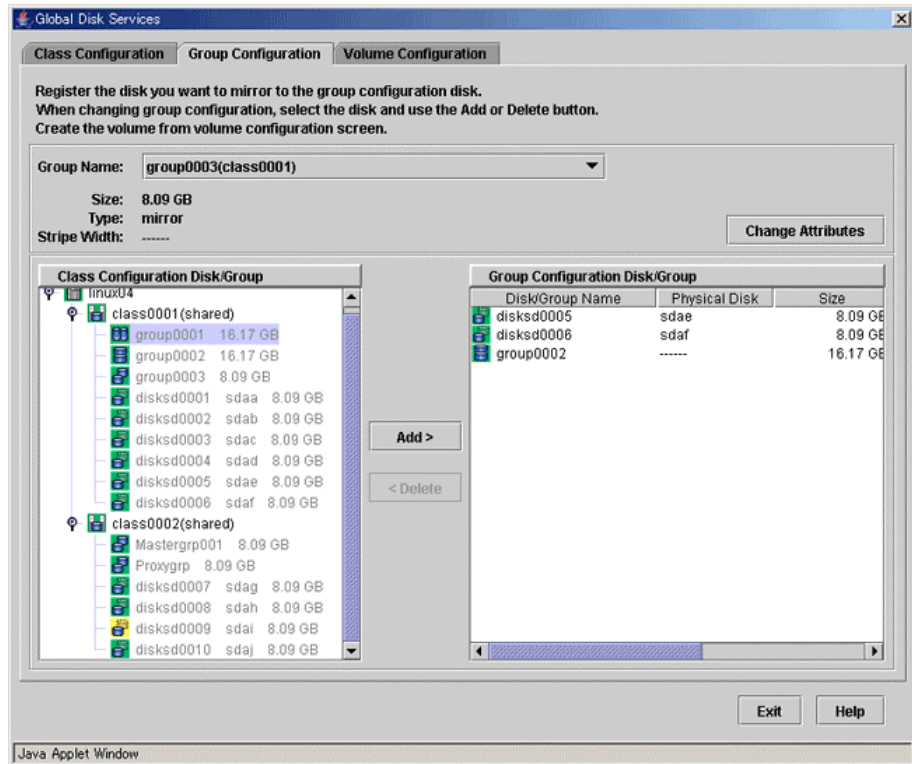
3. Changing configuration

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: 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: 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.



See

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

5.4.3 Volume Configuration

The procedures for changing volume configuration are explained below.

Changing Volume Configuration

Change volume attributes using the following procedures.



Note Preconditions for Changing Volume Attributes

- When the volume is active, the following volume attributes cannot be changed. To change these volume attributes, first stop the volume on all the nodes within the class scope.
 - Volume name
 - Physical slice attribute
- When a volume's slice is detached temporarily, the physical slice attribute of the volume cannot be changed. To change this attribute, first attach the temporarily detached slice to the volume.

1. Selecting volume to change

Display the you want to change in the Main screen Configuration Tree field. Click the icon and select the volume you want to change.

2. Displaying Volume Attributes Definition screen

In the Main screen [Operation] menu, select [Change Attributes]. Volume Attributes Definition screen appears.

Figure: Volume Attributes Definition Screen

Global Disk Services

Volume Attributes Definition
After defining the attributes, press OK.

Volume Name: volume0001

Size: 0.25 GB

JRM: on off

Physical Slice: on off

OK Cancel

Java Applet Window

3. Changing attributes

You can change the following attributes:

- Volume Name
- JRM (on/off)
- Physical slice attribute (on/off)



See

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

4. Applying changes

Click <OK> to apply changes in the attributes change, or <Cancel> to cancel.



Note **Changes in Special File Path Name**

Changing the volume name will also change the special file path name used to access the volume.

5.5 Removals

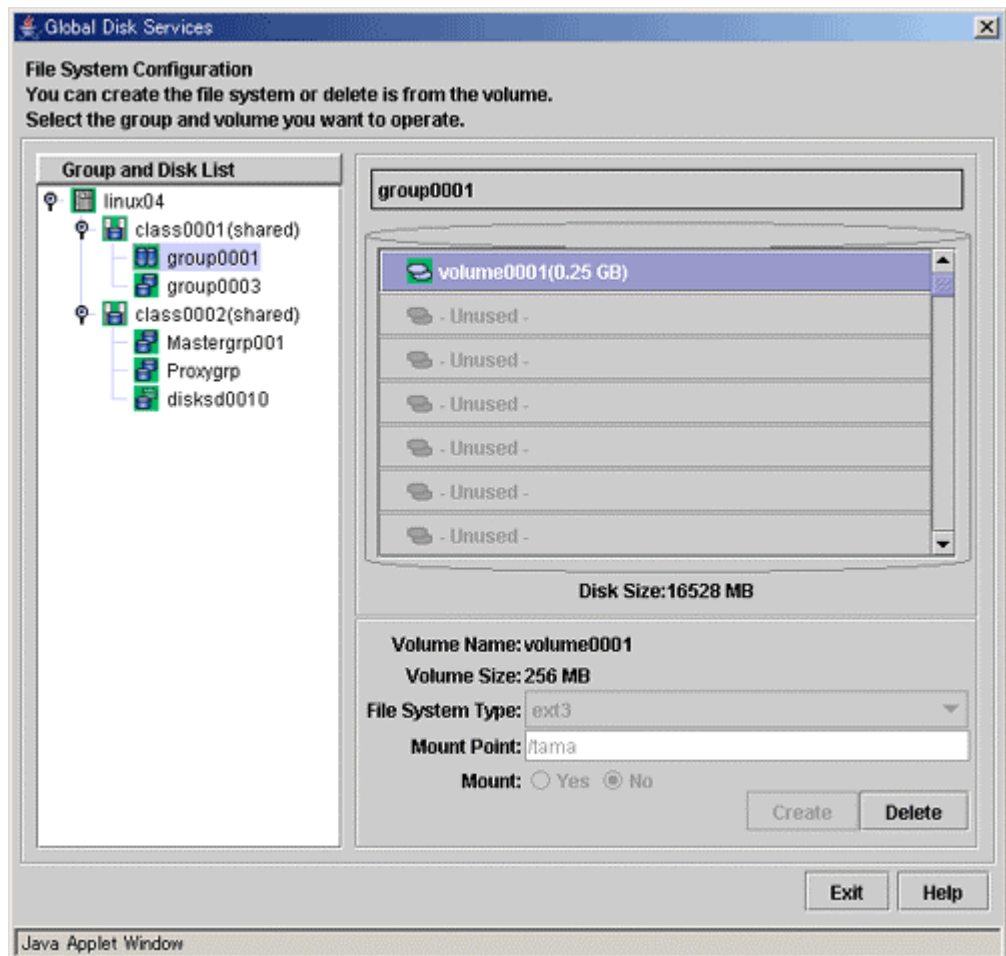
When file system is not created, you can start the unmirroring process from removing the volume.

5.5.1 Removing a File System

The procedures for removing a file system are explained below.

- 1. Displaying the File System Configuration screen**
In the Main screen [Settings] menu, select [File System Configuration]. File System Configuration screen appears.
- 2. Selecting group/disk**
In the Group/Disk List, select a group or disk with which you want to perform an operation.
- 3. Removing file system**
Select the volume from which you want to remove the file system, and click <Delete>. When removing more than one file system, select the next volume, and repeat the same process.

Figure: Removing File System



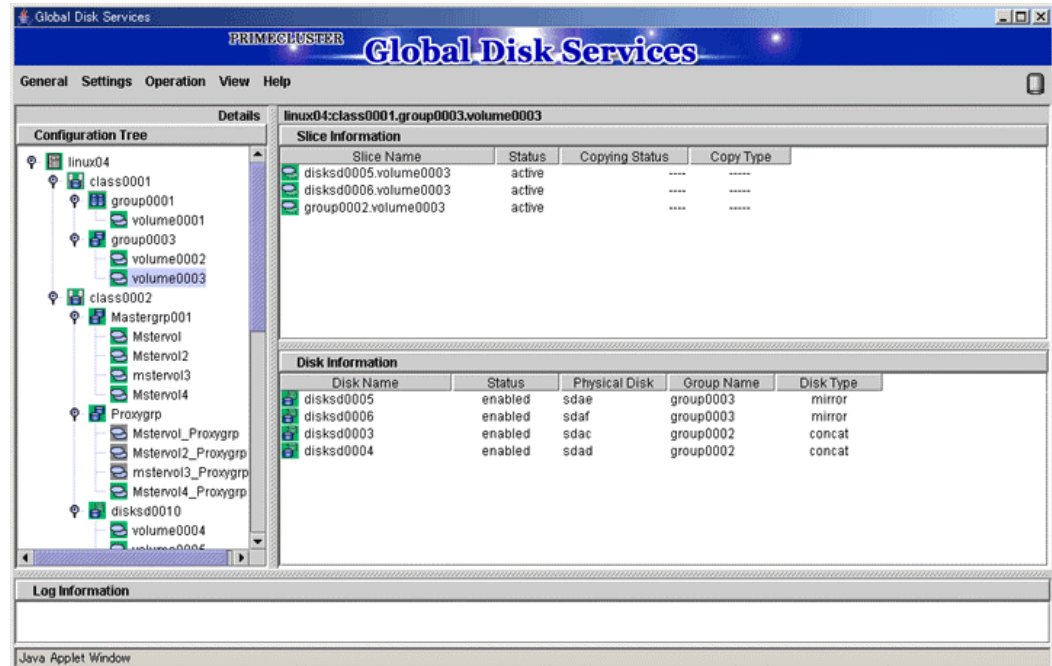
5.5.2 Removing a Volume

The procedures for removing a mirror volume are explained below.

1. Confirming the volume status

A volume containing a temporarily detached slice cannot be removed. When there is a temporarily detached slice, you must attach the slice before removing the volume. Volume status can be confirmed in the Main screen.

Figure: Displaying Volume Information



2. Attaching a temporarily detached slice

When there is a temporarily detached slice, go to [Operation] menu and select [Attach Slice]. The detached slice will be attached again. For information on Attach Slice operation, see "5.3.2.1 Backup (by Slice Detachment)."

3. Removing volume

The procedures for removing a volume are explained below.

(1) Displaying the Volume Configuration screen

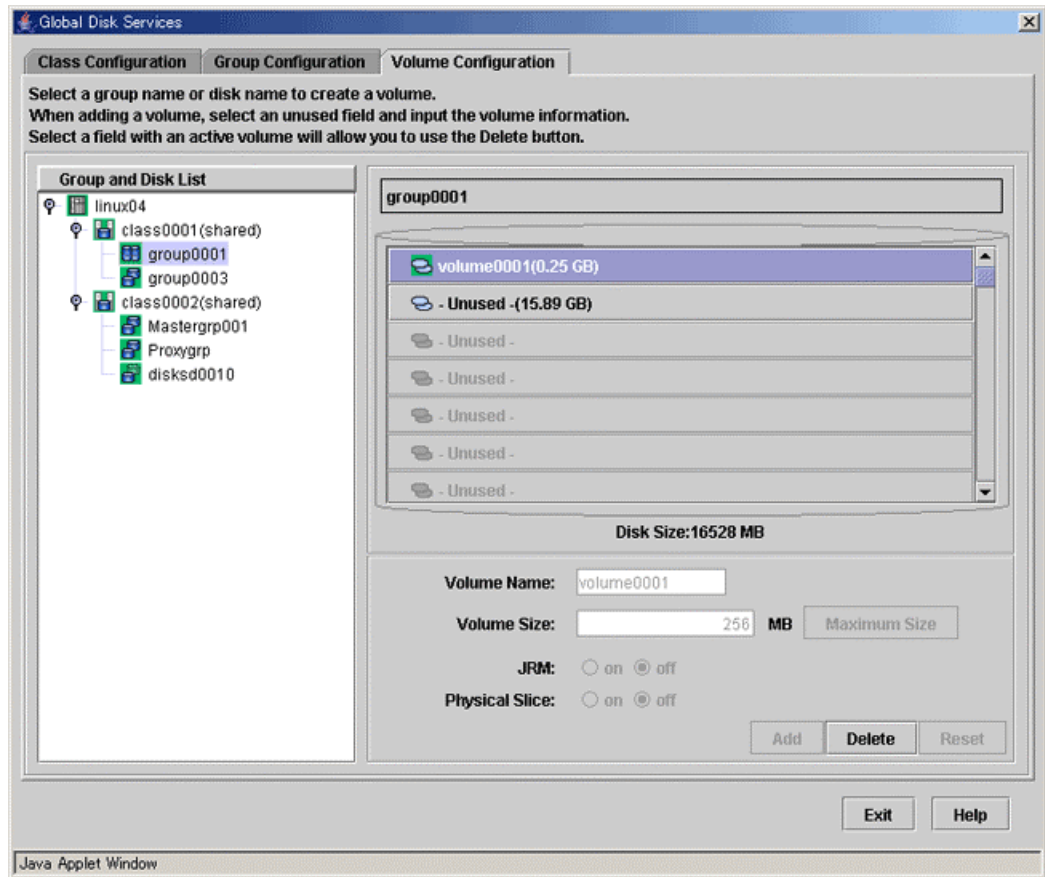
In the Main screen [Settings] menu, select [Volume Configuration]. Volume Configuration screen appears.

(2) Removing a volume

Select the volume you want to remove, and click <Delete>.

When removing more than one volume, select the next volume, and repeat the same process.

Figure: Removing Volume



5.5.3 Removing a Group

Disconnecting all disks/lower level groups registered with a group will automatically remove the group. The procedures are explained below.

1. Removing all volumes within the group

If there is even a single volume, you cannot remove the group.

Remove all volumes within the group by following the procedures described in “5.5.2 Removing a Volume.”

2. Disconnecting disks/lower level groups from group

Disconnect all disks/lower level groups by following the procedures below.

(1) Displaying Group Configuration screen

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

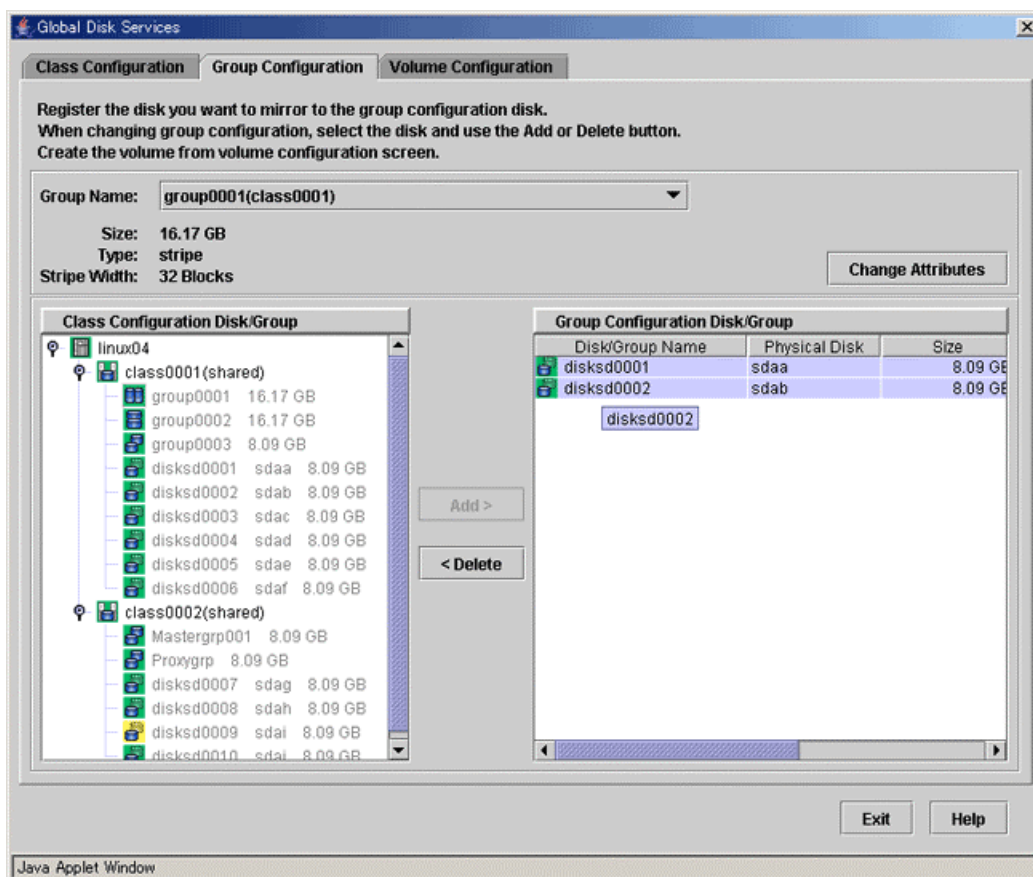
(2) Selecting group to remove

Select the group you want to remove from Group Configuration screen [Group Name] field.

(3) Disconnecting disk/lower level group

Select the disk/lower level group you want to remove from [Group Configuration Disk/Group] field, and click <Delete>.

Figure: Disconnecting Disk/Lower level group



5.5.4 Removing a Class

Removing all disks within a class will automatically remove the class. The procedure is explained below.

1. Removing all groups within the class

Remove all groups within the class by following the procedure described in “5.5.3 Removing a Group.”

2. Removing all disks from class

Disconnect all disks by following the procedure below.

(1) Displaying Class Configuration screen

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

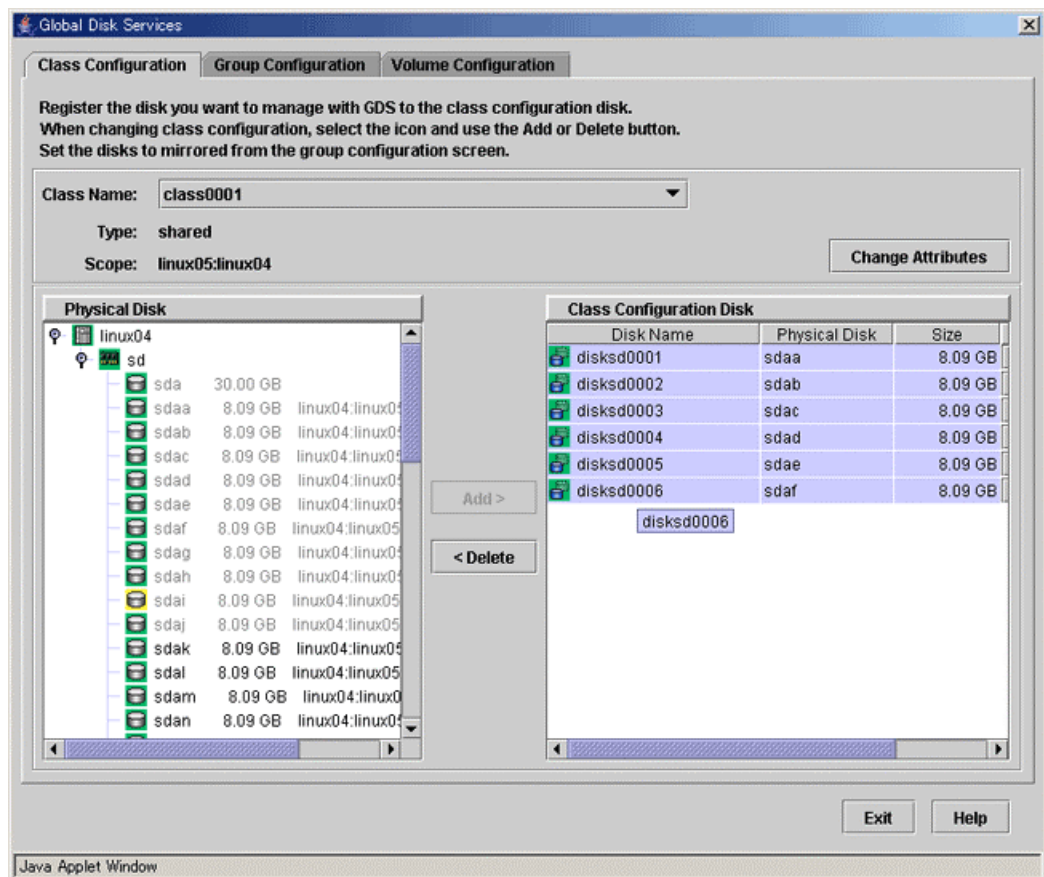
(2) Selecting class to remove

Select the class you want to remove from Class Configuration screen [Class Name] field.

(3) Removing disk

Select the disk you want to remove from [Class Configuration Disk] field, and click <Delete>.

Figure: Removing Disk



5.5.5 Unmirroring the System Disk [Itanium]

In this section, how to unmirror a system disk is explained.
In the [Settings] menu, select [Unmirror System Disk].



Note When You Cannot Unmirror the System Disk

If system disks are under the following conditions, the system disks cannot be unmirrored. In these situations, restore the disk status first and unmirror the system disks.

- A disk has been disconnected with [Swap Physical Disk].
- A disk is in disabled status.
- All disks in each group contain a slice that is not active.
- There is a proxy volume.
- There is a volume to which a proxy volume is joined.



Note For Safe Unmirroring of System Disk

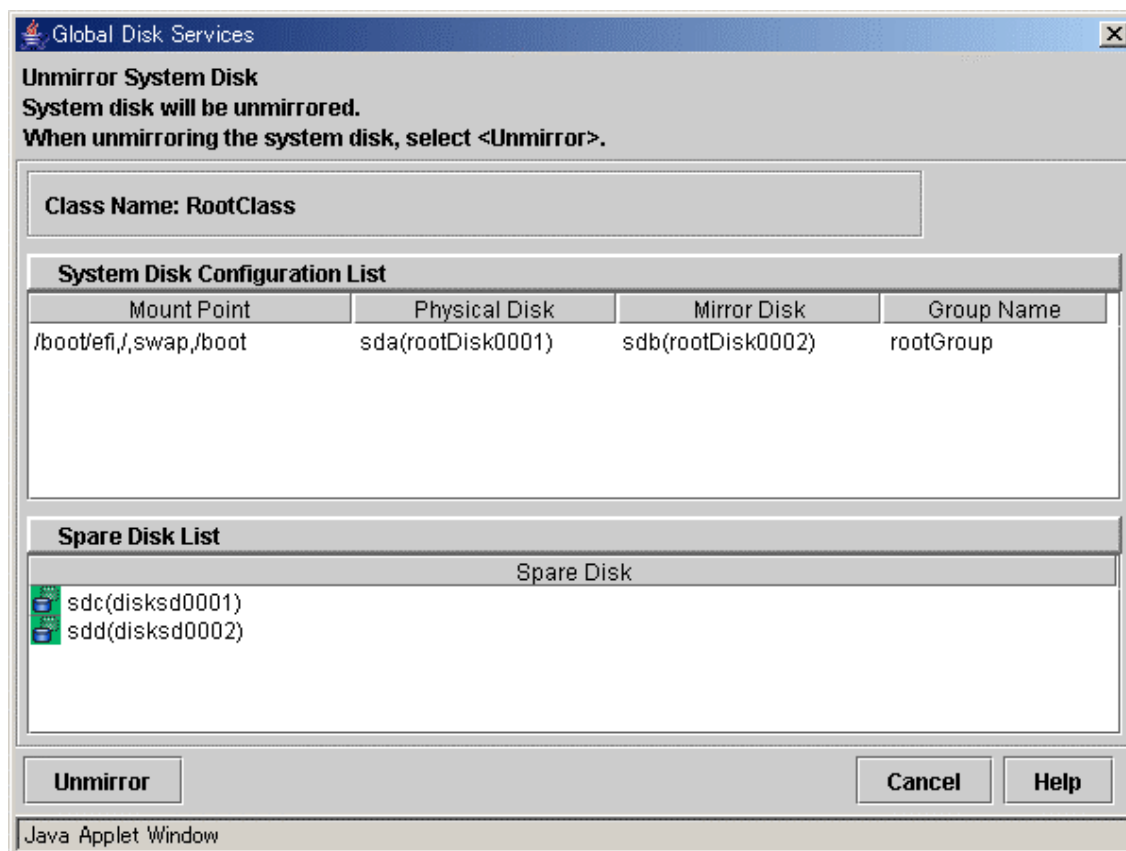
To safely unmirror the system disk, exit all active applications before proceeding to cancel system disk settings.

After unmirroring the system disk is complete, promptly reboot the system in multi-user mode.

1. Confirming system disk configuration

System disk mirroring configurations are displayed.

Figure: Unmirror System Disk



If unmirroring is performed, the disk displayed in the [Mirror Disk] field is disconnected and the disk displayed in the [Physical Disk] field will be used as the system disk.

If the root class includes only system disks with their mirror disks and spare disks, the entire root class is removed. Here, the spare disks displayed in the spare disk field are also removed.


If the root class includes objects other than system disks, such as single disks and mirror groups, those settings are retained and only the unmirroring of system disks is executed. Here, the spare disk field does not display any information and spare disks are not be 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: Unmirroring System Disk Confirmation Screen

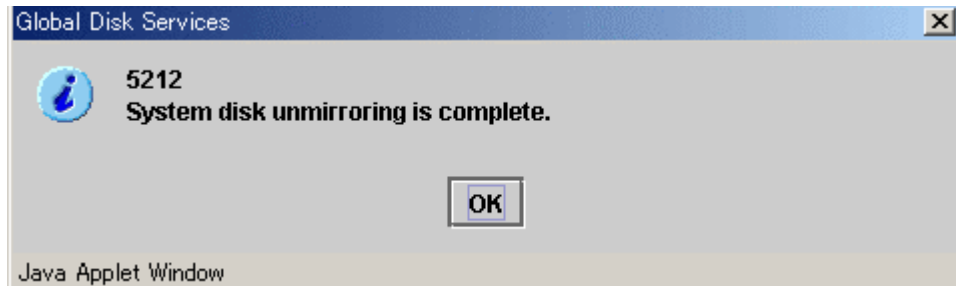


To continue the process, click <Yes>. Clicking <No> cancels the system disk unmirroring process.

3. Unmirroring System Disk Completion screen


If you click <Yes> on the Unmirroring System Disk Confirmation screen, the screen below appears.

Figure: Unmirroring System Disk Completion Screen



Confirm that unmirroring of system disk is complete and click <OK>.

Figure: 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.

5.5.6 Breaking a Proxy

This subsection describes the procedures for breaking the relationships between masters and proxies in GDS Snapshot installed systems.

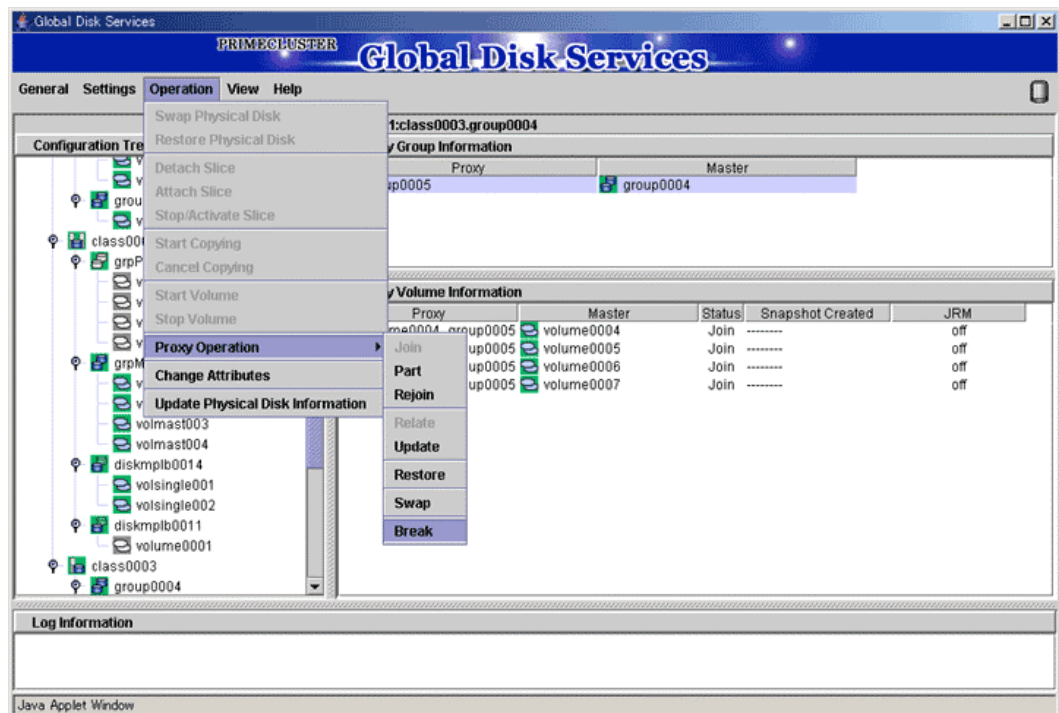
1. Selecting a master for break

Click an icon of a master object for break on the GDS Configuration Tree in the Main screen. Select [View]:[Details]:[Proxy Object] to view all the proxy objects related to the master selected on the GDS Configuration Tree, in the Object Information Field. In this field, click an icon of a proxy object for break.

2. Selecting the [Break] menu

Select [Operation]:[Proxy Operation]:[Break] in the Main screen.


Figure: Break



3. Confirming

A confirmation dialog box appears asking you whether to break the proxy.

Figure: 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: Information Message of "Break" Completion



Click <OK> to close the information message window.

Chapter 6 Backing Up and Restoring

The GDS mirroring function secures data from disk failures by preserving replicas of data on multiple disks. However, the mirroring function is not capable of protecting data against accidental erasure by the user or data crash due to an application malfunction. In addition, data can be lost due to multiple disk breakdown or large-scale disaster. To recover data when these troubles occur, data back up is mandatory. Be sure to back up data on a regular basis.

This chapter discusses the operating procedures for backing up and restoring GDS volumes to provide useful information concerning operational design for backup and restore.

6.1 Backing Up and Restoring a System Disk [Itanium]

Mirroring system disks will protect the data when a physical disk on one side fails. However, data must be restored from backup data created in advance if the data is damaged due to a critical failure caused by multiple breakdown and so on or by an operation mistake.

This section discusses the method of backing up data on a system disk to tape and the method of restoring data back to the system disk from tape. You must follow different restore procedures depending on whether or not the system can be booted.



Note

Data backed up before system disk mirroring cannot be restored back to the mirrored system disk. If system disk mirroring is configured, perform system disk backup using this procedure.



See

For backing up and restoring volumes (e.g. /opt, /home) other than system volumes (/, /usr, /var, /boot, /boot/efi, swap area) among volumes in the root class, see "6.3 Backing Up and Restoring Local Disks and Shared Disks."

6.1.1 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.



Note

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/System/dsk/rootVolume on / type ext3 (rw)
/dev/sfdsk/System/dsk/varVolume on /var type ext3 (rw)
/dev/sfdsk/System/dsk/usrVolume on /usr type ext3 (rw)
/dev/sfdsk/System/dsk/bootVolume on /boot type ext3 (rw)
/dev/sfdsk/System/dsk/efiVolume on /boot/efi type vfat (rw)
...
# swapon -s
Filename                                Type           ...
/dev/sfdsk/System/dsk/swapVolume        partition      ...
```

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 sdxinfo 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.

```
# sdxinfo -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 sdxinfo 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.

```
# sdxinfo -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.

```
# ls -l /sys/block/sda/device
lrwxrwxrwx 1 root root 0 Jun 1 2005 /sys/block/sda/device ->\
../../../../devices/pci0000:02/0000:02:1f.0/0000:06:02.0/host0\
target0:0:0:0:0:0
# ls -l /sys/block/sdb/device
lrwxrwxrwx 1 root root 0 Jun 1 2005 /sys/block/sdb/device ->\
../../../../devices/pci0000:02/0000:02:1f.0/0000:06:02.0/host0\
target0:0:2/0:0:2:0
```

For the ls 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.

6.1.2 Backing Up

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.



See

- When using GDS Management View, see "5.4.2 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 "6.1.4 Restoring (When the System Cannot Be Booted)," restore the backup target volumes.

a2) Shut down the system.

```
# shutdown -h now
```

a3) Turn on the power of the node, and insert the OS installation CD into the CD-ROM drive.

a4) From boot devices displayed in the boot option selector window of the EFI boot manager, select the CD-ROM device, and boot the system in rescue mode.

With Red Hat Enterprise Linux AS (v.4 for Itanium), boot the system using the following procedure.

For details, see the OS manual.

```
EFI Boot Manager ver 1.10

Please select a boot option

Root1
Root2
DVD/Acpi(PNP0A03,0)/Pci(1D|1)/Usb(0,0) ← Select
...

Use ↑ and ↓ to change option(s). Use Enter to select an option
```

- When “ELILO boot:” is displayed, enter “linux rescue.”
- In the Choose a Language window, select “English.”
- In the Keyboard Type window, select “us.” However, it may be necessary to choose another option depending on the keyboard.
- In the Setup Networking window, select “Yes” to perform network configuration, or select “No” not to. If “Yes” is selected, the IP address setup window will appear. Following the instructions of this window, specify the IP address.
- In the Rescue window, select “Skip.”

a5) Check the backup target physical slice name.

Check the backup target physical disk name as follows.

```
# ls -l /sys/block/sd*/device | grep 0000:06:02.0 | grep 0:0:0
lrwxrwxrwx 1 root root 0 Jun  1 2005 /sys/block/sda/device ->\
../../../../devices/pci0000:02/0000:02:1f.0/0000:06:02.0/host2/\
target0:0:0/0:0:0
```

For the grep command arguments, specify physical disk information on the backup target disk (Root1 in this example) confirmed as described in "6.1.1 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 "6.1.1 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
/var	sda2
/usr	sda3
/boot	sda4
/boot/efi	sda5

a6) Back up data of a file system to a tape medium.

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 a5).



See

For details on the backup methods, see the manuals of file systems to be backed up and used commands.

a7) Exit the rescue mode, and start the system.

With Red Hat Enterprise Linux AS (v.4 for Itanium), exit the rescue mode using the following command. For details, see the OS manual.

```
# exit
```

a8) When step a1) was performed, reconnect the disk disconnected in that step to the group.

```
# sdxdisk -C -c System -g Group1 -d Root2
```

Confirm that disk Root2 is connected to group Group1 (Group1 is indicated in the GROUP field of the Root2 line).

```
# sdxinfo -D -c System
OBJ   NAME   TYPE   CLASS  GROUP  DEVNAM  DEVBLKS  DEVCONNECT STATUS
----  -
disk  Root1   mirror System  Group1 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 "5.4.2 Group Configuration."

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.

```
# init 1
```

b3) 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/System/dsk/rootVolume on / type ext3 (rw)
```

In this example, the device special file for the volume of the root (/) file system is /dev/sfdsk/System/dsk/root Volume.

b4) Back up data of the file system to a tape medium.

The following example shows the procedure for backing up data to a tape medium of tape device /dev/st0 with the dump(8) command.

```
# dump 0uf /dev/st0 /dev/sfdsk/System/dsk/rootVolume
```

For the dump command's argument, specify the device special file of the volume displayed in step b3).



See

For details on the backup methods, see the manuals of file systems to be backed up and used commands.

b5) Reboot the system.

```
# shutdown -r now
```

6.1.3 Restoring (When the System Can Be Booted)

1) Exit all running application programs. If higher security is required, you should make a backup of the system disk in advance. For details on the backup procedure, see "6.1.2 Backing Up."

2) Disconnect disks other than the disk that will be the restore destination from the group to have only one disk connected to the group.

The command line shown below is an example of one used to disconnect disk Root2 from group Group1 to use disk Root1 as the restore destination when Root1 and Root2 are connected to Group1 and mirrored.

```
# sdxdisk -D -c System -g Group1 -d Root2
```

Confirm that only one disk is connected to group Group1 (Group1 is indicated in the GROUP field of only one disk).

```
# sdxinfo -D -c System
```

OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	DEVCONNECT	STATUS
disk	Root1	mirror	System	Group1	sda	35368272	node1	ENABLE
disk	Root2	undef	System	*	sdb	35368272	node1	ENABLE



Information

If disk Root1 has an INVALID slice, disconnect the Root1.



See

When using GDS Management View, see "5.4.2 Group Configuration."

3) Shutdown the system.

```
# shutdown -h now
```

4) Turn on the power of the node, and insert the OS installation CD into the CD-ROM drive.

5) From boot devices displayed in the boot option selector window of the EFI boot manager, select the CD-ROM device, and boot the system in rescue mode.

With Red Hat Enterprise Linux AS (v.4 for Itanium), boot the system using the following procedure.

For details, see the OS manual.

```
EFI Boot Manager ver 1.10
```

```
Please select a boot option
```

```
Root1
```

```
Root2
```

```
DVD/Acpi(PNP0A03,0)/Pci(1D|1)/Usb(0,0) ← Select
```

```
...
```

```
Use ↑ and ↓ to change option(s). Use Enter to select an option
```

- When "ELILO boot:" is displayed, enter "linux rescue."
- In the Choose a Language window, select "English."
- In the Keyboard Type window, select "us." However, it may be necessary to choose another option depending on the keyboard.
- In the Setup Networking window, select "Yes" to perform network configuration, or select "No" not to. If "Yes" is selected, the IP address setup window will appear. Following the instructions of this window, specify the IP address.
- In the Rescue window, select "Skip."

6) Check the restore destination physical slice name.
Check the restore destination physical disk name.

```
# ls -l /sys/block/sd*/device | grep 0000:06:02.0 | grep 0:0:0
lrwxrwxrwx 1 root root 0 Jun  1  2005 /sys/block/sda/device ->\
../devices/pci0000:02/0000:02:1f.0/0000:06:02.0/host2/\
target0:0:0:0:0:0
```

For the grep command arguments, specify physical disk information on the restore destination disk (Root1 in this example) confirmed as described in "6.1.1 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 "6.1.1 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

7) 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 ext3, and a temporary mount point is the /work directory.

```
# mkdir /work
# mkfs.ext3 /dev/sda2
# mount -t ext3 /dev/sda2 /work
# cd /work
# restore rf /dev/st0 .
# cd /
# umount /work
```

For the mkfs.ext3(8) command's and mount(8) command's arguments, specify the physical slice displayed in step 6).



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.

8) Exit the rescue mode, and boot the system.
With Red Hat Enterprise Linux AS (v.4 for Itanium), exit the rescue mode using the following command.
For details, see the OS manual.

```
# exit
```

9) 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     DEVNAM     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 "5.4.2 Group Configuration."

6.1.4 Restoring (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 selector window of the EFI boot manager, select the CD-ROM device, and boot the system in rescue mode.
With Red Hat Enterprise Linux AS (v.4 for Itanium), boot the system using the following procedure.
For details, see the OS manual.

```
EFI Boot Manager ver 1.10

Please select a boot option

Root1
Root2
DVD/Acpi(PNP0A03,0)/Pci(1D|1)/Usb(0,0) ← Select
...

Use ↑ and ↓ to change option(s). Use Enter to select an option
```

- When "ELILO boot." is displayed, enter "linux rescue."
- In the Choose a Language window, select "English."
- In the Keyboard Type window, select "us." However, it may be necessary to choose another option depending on the keyboard.
- In the Setup Networking window, select "Yes" to perform network configuration, or select "No" not to. If "Yes" is selected, the IP address setup window will appear. Following the instructions of this window, specify the IP address.
- In the Rescue window, select "Skip."

- 3) Check the restore destination physical slice names.
Check the restore destination physical disk names.

```
# ls -l /sys/block/sd*/device | grep 0000:06:02:0 | grep 0:0:0
lrwxrwxrwx 1 root root 0 Jun 1 2005 /sys/block/sda/device ->\
.../devices/pci0000:02/0000:02:1f.0/0000:06:02.0/host2\
target0:0:0:0:0:0
# ls -l /sys/block/sd*/device | grep 0000:06:02:0 | grep 0:2:0
lrwxrwxrwx 1 root root 0 Jun 1 2005 /sys/block/sdb/device ->\
.../devices/pci0000:02/0000:02:1f.0/0000:06:02.0/host2\
target0:0:2:0:0:2:0
```

For the grep command arguments, specify physical disk information on the restore destination disks (Root1 and Root2 in this example) confirmed as described in "6.1.1 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 "6.1.1 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
/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 `ext3`, and a temporary mount point is the `/work` directory.

```
# mkdir /work
# mkfs.ext3 /dev/sda2
# mount -t ext3 /dev/sda2 /work
# cd /work
# restore rf /dev/st0 .
# cd /
# umount /work
```

For the `mkfs.ext3(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, and boot the system.

With Red Hat Enterprise Linux AS (v.4 for Itanium), exit the rescue mode using the following command. For details, see the OS manual.

```
# exit
```


6.2 Backing Up and Restoring a System Disk through an Alternative Boot Environment [Itanium]

This section discusses the method of backing up system disks with the GDS Snapshot function and creating an alternative boot environment by use of the backup copies, and the method of restoring the system disks through the alternative boot environment.

GDS Snapshot can collect snapshots of system disks (replications at a certain point) in the disk areas for backup (proxy volumes) during the system operation. You can configure an alternative boot environment to allow the system to be booted through such a proxy volume in order that the system operates in the alternative boot environment even if the system cannot be booted normally due to a failed system disk or damaged data. After switching to the alternative boot environment, you can also restore the original boot environment simply by restoring backup disk data to the original system disks and rebooting the system.

In this manner, operation down time while backing up and restoring system disks and time required for recovery from a system disk failure can be reduced drastically.

6.2.1 System Configuration

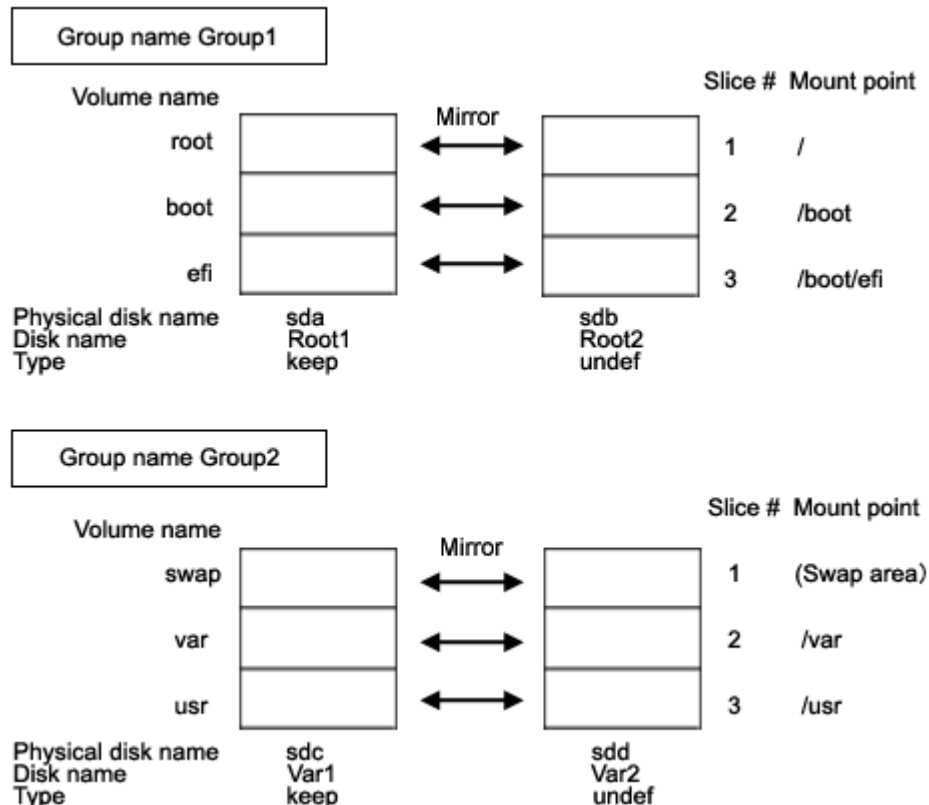
In preparation for alternative boot environment creation, mirror system disks. This sub-section provides an example of mirroring system disks in the configuration as shown in the figure below.



Information

System disk mirroring is not required. A configuration without 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: 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: Proxy Group Configuration

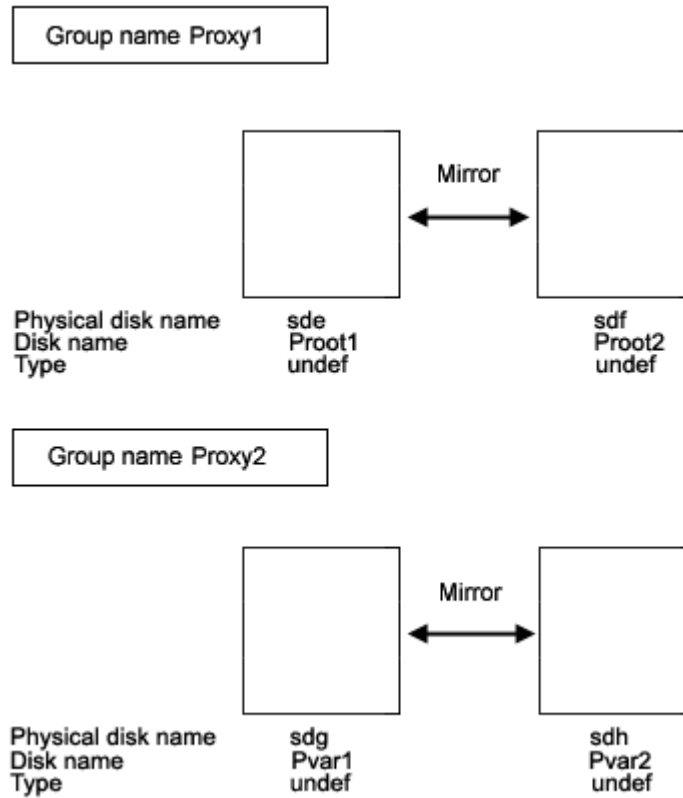
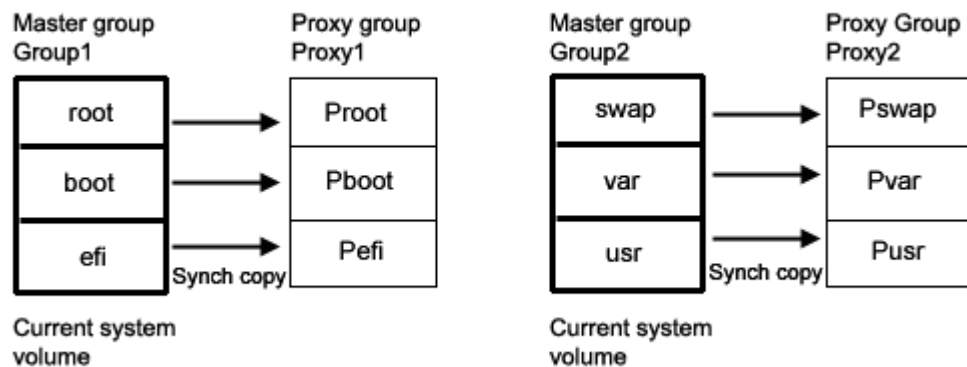


Figure: Joining Proxy Groups



6.2.2 Summary of Backup

Snapshots of data in system disks must be collected in proxy groups during the system operation. Additionally, the environment allowing for booting through the proxy groups should be configured for errors due to a failed system disk and damaged data.

Figure: Backup

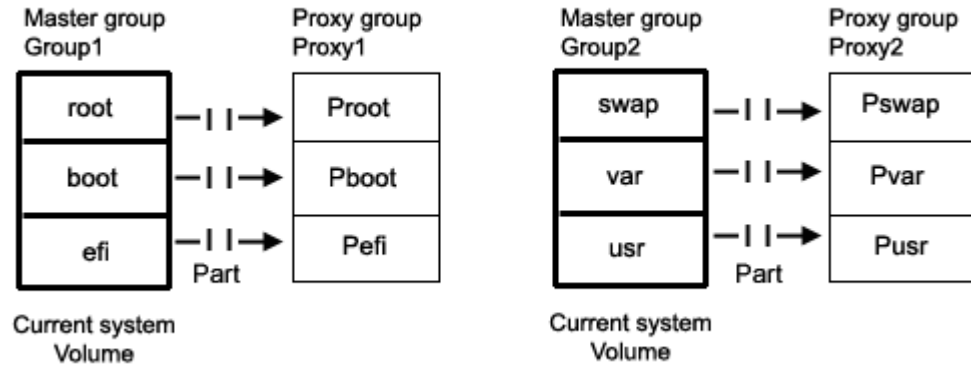



Figure: Backup Schedule



6.2.3 Summary of Restore

If booting the system becomes unavailable due to a failed system disk or damaged data, you can switch the environment to the alternative boot environment created on proxy groups to enable the system to continue operating. You can restore data back to the original system disks by joining the disk areas for backup as the master and the original system disks as the proxy.

Figure: Restore

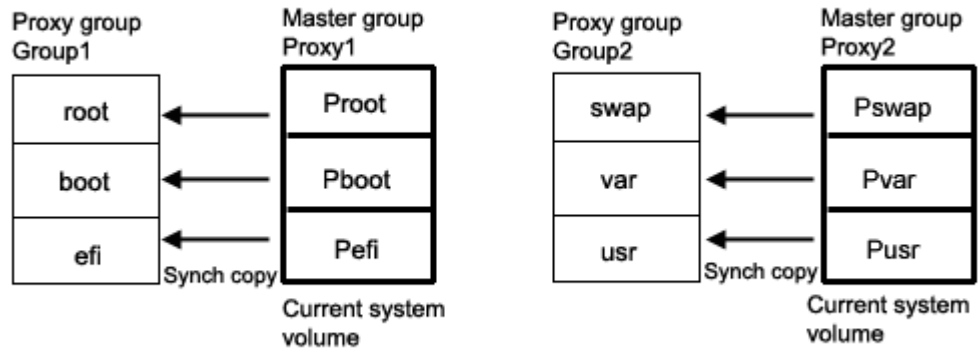
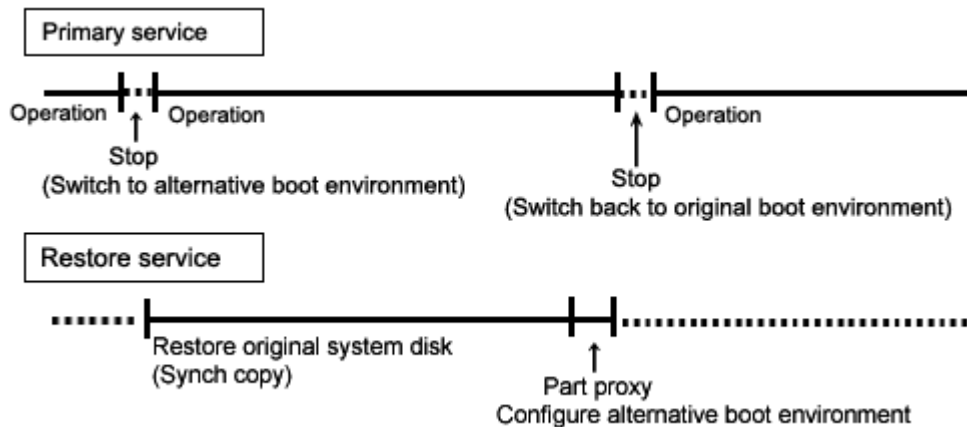


Figure: Restore Schedule



6.2.4 Summary of Procedure

Figure: Outline of the Configuration Procedure




Figure: Outline of the Backup Procedure

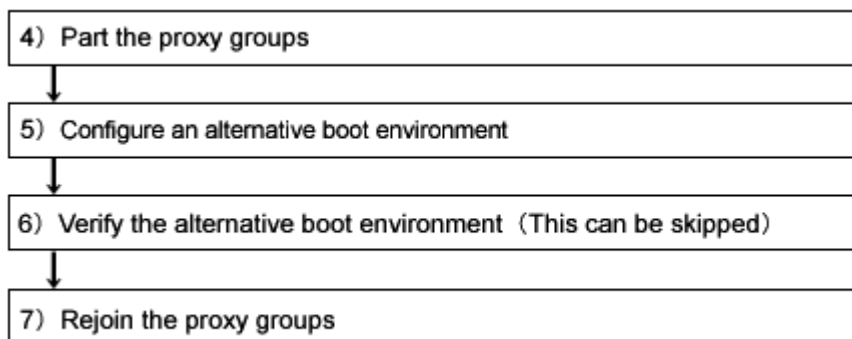



Figure: Outline of the Restore Procedure



6.2.5 Configuring an Environment

1) Mirroring system disks

In preparatory for alternative boot environment creation, mirror system disks. This sub-section describes the procedures for mirroring system disks in the configuration as shown in "6.2.1 System Configuration."



See

For details on GDS Management View, see "5.2.1 System Disk Settings [Itanium]."

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 "6.2.1 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 "5.4.1 Class Configuration."

2-2) Connect the disks to groups respectively.

```
# sdxdisk -C -c System -g Proxy1 -d Proot1,Proot2  
# sdxdisk -C -c System -g Proxy2 -d Pvar1,Pvar2
```



See

For details on GDS Management View, see "5.2.2.3 Group Configuration."

3) Joining the proxy groups

Copy data in system disks into the backup disks by joining a group of the backup disks (proxy group) to a group of the system disks (master group). The following describes the procedure for joining proxy groups in the configuration as shown in "6.2.1 System Configuration."

3-1) Join the proxy groups.

```
# sdxproxy Join -c System -m Group1 -p Proxy1 -a
root=Proot:on,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 "5.2.4.1 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.

6.2.6 Backing Up

4) Parting the proxy groups

Once synchronization copying is completed, the master group and the proxy group become synchronized. Snapshots of a master group can be collected in a proxy group by parting the synchronized master group and proxy group.

4-1) Secure consistency of the file systems.

To secure consistency of snapshot file systems, restrain update of the file systems. However, file systems such as /, /usr, and /var are necessary for the system to operate, they cannot be unmounted during system operation. You can follow the procedure as below in order to reduce write occurrences to system disks and write occurrences to system disks that are not updated yet.

- a) Boot the system in single user mode. (This can be skipped.)
- b) Exit all active application programs writing in the system disks. (This can be skipped.)
- c) 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) Part the proxy groups.

```
# sdxproxy Part -c System -p Proxy1,Proxy2
```



See

When using GDS Management View, see "Part Proxy" in "5.3.2.2 Backup (by Synchronization)."

4-3) When the system was booted in single user mode in a) of step 4-1), reboot it in multi-user mode.

4-4) When application programs were exited in b) of step 4-1), launch the application programs.

5) Configuring an alternative boot environment

Make booting available through the proxy volumes in preparation for an error due to a failed system disk or damaged data.

5-1) Set the access mode attribute of the proxy volumes to rw (read and write).

If the access mode attribute of proxy volumes created in the proxy group is ro (read only), it must be changed to rw (read and write). The access mode attribute can be viewed in the MODE field output by the sdxinfo -V -e long command. If the access mode attribute is already set to rw (read and write), executing the following commands is not necessary.

```
# sdxvolume -F -c System -v Proot,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.

```
# 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 selector window, select one of the devices in the alternative boot environment output in the message as shown in step 5-3).

```
EFI Boot Manager ver 1.10

Please select a boot option

Root1
Root2
Proot1 ← Select
Proot2
...

Use ↑ and ↓ to change option(s). Use Enter to select an option
```

6-2) 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.

6-3) Return to the original boot environment.

From boot devices displayed in the EFI boot manager's boot option selector window, select one of the devices in the alternative boot environment output in the message as shown in step 5-3).

```
EFI Boot Manager ver 1.10

Please select a boot option

Root1 ← Select
Root2
Proot1
Proot2
...

Use ↑ and ↓ to change option(s). Use Enter to select an option
```

6-4) Stop the proxy volumes.

To protect data in the alternative boot environment from illegal write access, the proxy volumes should be inactivated.

```
# sdxvolume -F -c System -v Proot,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 copying when the proxy groups are rejoined in step 7-1) if the system is rebooted after step 4-3) or step 6).

7-1) Rejoin the proxy groups.

```
# sdxproxy Rejoin -c System -p Proxy1,Proxy2
```



See

When using GDS Management View, see "Part Proxy" in "5.3.2.2 Backup (by Synchronization)."

7-2) Confirm that synchronization copying is complete.

```
# sdxinfo -S -c System
```

OBJ	CLASS	GROUP	DISK	VOLUME	STATUS
slice	System	Group1	Root1	root	ACTIVE
slice	System	Group1	Root2	root	ACTIVE
slice	System	Group1	Root1	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).

6.2.7 Restoring

8) Switching to the alternative boot environment

If the system cannot be booted due to a failed system disk or damaged data, switch the environment to the alternative boot environment created in the proxy volume to allow the system to continue operating

8-1) Boot in the alternative boot environment.

From boot devices displayed in the EFI boot manager's boot option selector window, select one of the devices in the alternative boot environment output in the message as shown in step 5-3).

```
EFI Boot Manager ver 1.10

Please select a boot option

Root1
Root2
Proot1 ← Select
Proot2
...

Use ↑ and ↓ to change option(s). Use Enter to select an option
```

8-2) 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.

8-3) Break up the former boot environment according to need.

To break up the former boot environment, break the master and proxy relationships, remove the master volumes, and remove groups and disks from the mater 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
```



See

When using GDS Management View, see "5.5.6 Breaking a Proxy," "5.5.2 Removing a Volume," "5.5.3 Removing a Group," and "5.4.1 Class Configuration."

9) Restoring system disks

After rebooting the system in the alternative boot environment, restore backup disk data back to the original system disks.

9-1) Cancel the master and proxy relationship.

```
# sdxproxy Break -c System -p Proxy1  
# sdxproxy Break -c System -p Proxy2
```



See

When using GDS Management View, see "5.5.6 Breaking a 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 "5.5.2 Removing a 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 "5.3.4 Disk Swap."

9-3-1) Exclude the disk to be swapped from the GDS management to make it exchangeable.

```
# sdxswap -O -c System -d Root1
```

9-3-2) Swap physical disk 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 "5.2.4.1 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 “6.2.6 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.

	Slice #	Mount point	Slice #	Mount point
Physical disk name sda	1	/	1	(Swap area)
	2	/boot	2	/var
	3	/boot/efi	3	/usr
				Physical disk name sdc

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 "5.5.6 Breaking a Proxy," "5.5.2 Removing a Volume," "5.5.3 Removing a Group," and "5.4.1 Class Configuration."

11) Unmirroring system disks in an alternative boot environment



See

When using GDS Management View in steps 10-2) and after, see "5.5.5 Unmirroring the System Disk [Itanium]."

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 cancelled normally.

Using the mount(8) command or the sdxinfo command, verify that the system disk mirroring was cancelled properly.

11-6) Cancel the system disk management.

```
# sdxvolume -F -c System -v Proot,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
```

6.3 Backing Up and Restoring Local Disks and Shared Disks

This section discusses the methods of backing up and restoring local disks and shared disks in a system where GDS Snapshot has not been installed. Among volumes in the root class, volumes (e.g. /opt, /home) other than system volumes (/usr, /var, swap area) can also be backed up and restored following the procedures described here. However, volumes in the root class cannot be backed up with the “6.3.2 Online Backup (by Slice Detachment)” method.

The following is an example of backing up and restoring volume Volume1 in class Class1.

6.3.1 Offline Backup

1) Stopping the services

1a) With a shared volume used in a cluster application

1a-1) Exit all cluster applications.

1a-2) Activate the volume on a node on which backup is conducted.

```
# sdxvolume -N -c Class1 -v Volume1
```

1b) With a volume not used by a cluster application

1b-1) Stop the services using the volume.

1b-2) When the volume is used as a file system, unmount the file system. In the following example, the mount point is /mnt1.

```
# cd /  
# umount /mnt1
```

2) Backing Up

Back up volume data. In the following 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 data with the dd(1) command

```
# dd if=/dev/sfdsk/Class1/dsk/Volume1 of=/dev/st0 bs=32768
```


3) Resuming the services

3a) With a shared volume used in a cluster application

3a-1) Inactivate the volume on the node where backup was conducted.

```
# sdxvolume -F -c Class1 -v Volume1
```

3a-2) Launch cluster applications.

3b) With a volume not used by a cluster application

3b-1) When the volume is used as a file system, mount the file system. The following shows examples when the mount point is /mnt1.

- For the ext3 file system

```
# mount -t ext3 /dev/sfdsk/Class1/dsk/Volume1 /mnt1
```

3b-2) Resume the services.

6.3.2 Online Backup (by Slice Detachment)

For mirror volumes, data can be backed up through snapshot by slice detachment during the services operation.



See

For use conditions on snapshots by detaching a slice, see "A.2.10 Creating a Snapshot by Slice Detachment."

To secure consistency of data in a detached slice, the services must be stopped temporarily when detaching the slice.



See

For securing consistency of snapshot data, see "A.2.20 Ensuring Consistency of Snapshot Data."

1) Stop the services

1a) With a shared volume used by a cluster application
Exit the cluster application.

1b) With a volume not used by a cluster application

1b-1) Stop the services using the volume.

1b-2) When the volume is used as a file system, unmount the file system. In the following example, the mount point is /mnt1.

```
# cd /  
# umount /mnt1
```

2) Detaching the slice

Detach the slice from the mirror volume. The following shows an example of detaching the slice from disk Disk1.

```
# sdxslice -M -c Class1 -d Disk1 -v Volume1
```



See

When using GDS Management View, see "Slice Detachment" in "5.3.2.1 Backup (by Slice Detachment)."

3) Resuming the services

3a) With a shared volume used by a cluster application
Execute the cluster application.

3b) With a volume not used by a cluster application

3b-1) When the volume is used as a file system, mount the file system. In the following example, the mount point is /mnt1.

- For the ext3 file system

```
# mount -t ext3 /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 ext3 file system

```
# fsck -t ext3 -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
```



[See](#)

When using GDS Management View, see “Attach Slice” of “5.3.2.1 Backup (by Slice Detachment).”

6.3.3 Restoring

1) Stopping the services

1a) With a shared volume used by a cluster application.
Exit the cluster application.

1b) With a volume not used by a cluster application

1b-1) Stop the services using the volume.

1b-2) When the volume is used as a file system, unmount the file system. In the following example, the mount point is /mnt1.

```
# cd /  
# umount /mnt1
```

2) Restoring

Restore the volume data. The following shows an example of restoring data from a tape medium of tape device /dev/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
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 ext3 file system

```
# mount -t ext3 /dev/sfdsk/Class1/dsk/Volume1 /mnt1
```

3b-2) Resume the services.

6.4 Online Backup and Instant Restore through Proxy Volume

This subsection describes the procedures for online backup and instant restore of local volumes and shared volumes through use of snapshots by GDS Snapshot proxy volumes.

There are two online backup methods: one uses the function as "1.5.1 Snapshot by Synchronization" and the other uses the function as "1.5.3 Instant Snapshot by OPC." By cooperation with the Advanced Copy function (EC or OPC) of the Fujitsu Storage Systems ETERNUS, or a copy function (EMC TimeFinder or EMC SRDF) of EMC's Symmetrix storage unit, backup and restore processing that does not impose the load on the server or SAN becomes possible.

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

Operation	Disk Unit's Copy Function	Master Configuration (*1)	Proxy Configuration (*1)	Feature
Backup(Sync)	-	single mirror stripe concat	single mirror stripe concat	With synchronized masters and proxies, snapshots are created simply by parting the proxies. By use of the JRM for proxies, high-speed resynchronization is executable for snapshot re-creation. Striping and concatenating masters are possible. Mirroring, striping and concatenating proxies are possible.
	EC REC TimeFinder SRDF	single mirror	single	With synchronized masters and proxies, snapshots are created simply by parting the proxies. By use of a disk unit's copy function, snapshot creation that does not impose the load on the server or SAN is possible. By use of a disk unit's incremental copy function, high-speed resynchronization that does not impose the load on the server or the SAN is performed for snapshot re-creation.
	OPC ROPC	single mirror	single	With synchronized masters and proxies, snapshots are created simply by parting the proxies. By use of a disk unit's copy function, snapshot creation that does not impose the load on the server or SAN is possible.
Backup(OPC)	OPC ROPC	single mirror	single	There is no need to synchronize masters and proxies before creating snapshots, and backups can be created any time without scheduling in advance. By use of a disk unit's copy function, snapshot creation that does not impose the load on the server or SAN is possible.
Restore	-	single mirror stripe concat	single mirror stripe concat	By use of the JRM for proxies, high-speed restore is possible. Striping and concatenating masters are possible. Striping and concatenating proxies are possible.
	OPC ROPC	single mirror	single mirror	By use of a disk unit's copy function, restore that does not impose the server or SAN is possible.

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

Configuration	Description
single	This configuration means any one of the following objects. 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	This configuration means either the following objects. 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
stripe	This configuration means either the following objects. A mirror group to which one or more stripe groups are connected A volume created in a mirror group to which one or more stripe groups are connected
concat	This configuration means either the following objects. 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 the Server/SAN," "A.2.16 Using the Advanced Copy Function in a Proxy Configuration," "A.2.17 Instant Snapshot by OPC" and "A.2.19 Using EMC TimeFinder or EMC SRDF in a Proxy Configuration [Linux2.6]."

6.4.1 Online Backup (by Synchronization)

In a system employing GDS Snapshot, backups can be created using snapshots from proxy volumes during the services operation.

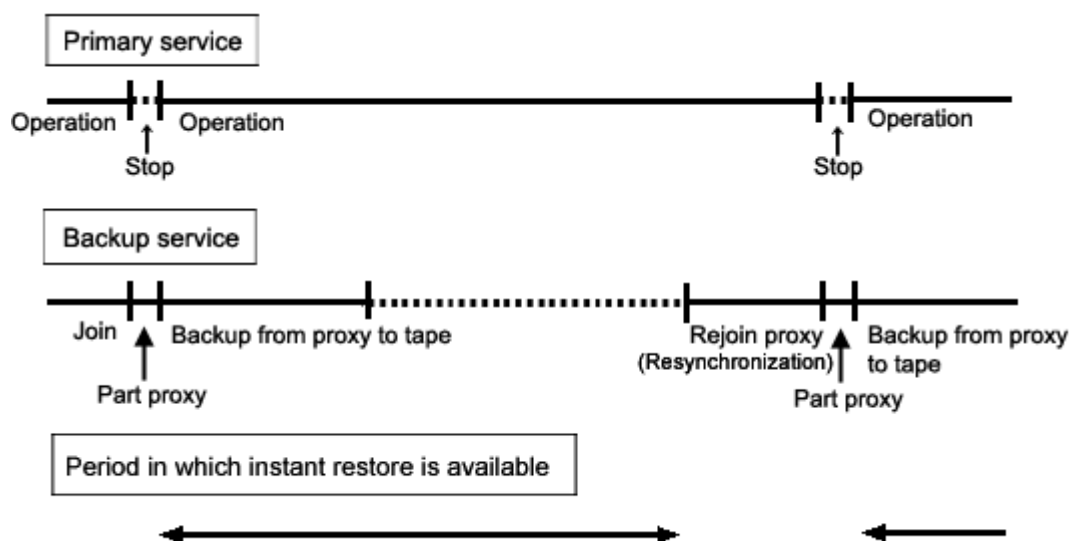
However, to secure consistency of snapshot data, the services must be stopped temporarily when creating a snapshot (when parting a proxy volume from a master volume).



See

For securing consistency of snapshot data, see "A.2.20 Ensuring Consistency of Snapshot Data."

Figure: Backup Schedule



Note Instant Restore

If an error occurs in master volume data while the master volume and a proxy volume are joined, the same error occurs in data of the proxy volume and instant restore becomes unavailable. If this happens, you must restore data from tape. After you execute online backup by parting a proxy volume, you are recommended to leave the proxy volume parted until just before re-executing online backup.



Note Snapshot through Disk Unit's Copy Functions

If groups are hierarchized, or if proxies are in mirroring configuration with the multiplicity of two and higher, copy functions of disk units are unavailable for copying data from masters to proxies. For details, see "A.2.16 Using the Advanced Copy Function in a Proxy Configuration." and "A.2.19 Using EMC TimeFinder or EMC SRDF in a Proxy Configuration [Linux2.6]."



Note When Using EMC TimeFinder or EMC SRDF [Linux2.6]

- In the procedures described below, volumes are used for snapshot operations. However, when using TimeFinder or SRDF, use groups for similar operations.
- When using TimeFinder, if standard devices that compose the master group and BCV devices that compose the proxy group are established, cancel the BCV pairs prior to joining the master and the proxy.
- When using SRDF, have source (R1) devices that compose the master group and target (R2) devices that compose the proxy group split prior to joining the master and the proxy.
- After relating the master and the proxy, do not perform TimeFinder or SRDF operations for BCV pairs or SRDF pairs that compose the master and proxy using the SYMCLI command and so on before canceling the relation.

Procedure

1) Joining a proxy volume

To prepare for creating a snapshot, relate and join proxy volume Volume2 as the copy destination of master volume Volume1 to the master volume. This example shows the procedure when Volume1 and Volume2 belong to Class1. Execute the following commands on a node that belongs to the scope of class Class1.

1-1) Stop proxy volume Volume2. If Class1 is a shared class, specify the “-e allnodes” option to stop Volume2 on all nodes.

```
# sdxvolume -F -c Class1 -v Volume2
```



Information

When using GDS Management View, select a proxy volume and execute [Operation]:[Stop Volume] in the Main Screen.

1-2) Relate and join proxy volume Volume2 to master volume Volume1.

```
# sdxproxy Join -c Class1 -m Volume1 -p Volume2
```

After returning from the command, synchronization copying from Volume1 to Volume2 is executed.



Information Relating and Joining a Pair of Groups

If the proxy group includes volumes, remove the volumes before executing the sdxproxy Join command, and also specify the -a option for this command.
Example)

Relate and join proxy group Group2 to master group Group1.
Assign volume names Proxy1 and Proxy2 to the proxy volumes automatically created in Group2 corresponding to volumes Volume1 and Volume2 in Group1.

```
# sdxproxy Join -c Class1 -m Group1 -p Group2 \  
-a Volume1=Proxy1:on,Volume2=Proxy2:on
```



See

When using GDS Management View, see "5.2.4.1 Join."

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
```



See

When using GDS Management View, see "Part" in "5.3.2.2 Backup (by Synchronization)." For setting of the part environment, select "No" to "Instant Snapshot."

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 ext3 file system on master volume Volume1 is mounted on mount point /DATA.

```
# mount -t ext3 /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.



See

For details on the backup methods, see the manuals of file systems to be backed up and used commands.

6a) When backing up data with the dd(1) command

```
# dd if=/dev/sfdsk/Class1/dsk/Volume2 of=/dev/st0 bs=32768
```

6b) When backing up the ext3 file system with the tar(1) command

6b-1) Before mounting

Check and repair consistency of the ext3 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 ext3 /dev/sfdsk/Class1/dsk/Volume2
```

6b-2) Mounting the snapshot

Mount the ext3 file system on proxy volume Volume2 on temporary mount point /DATA_backup.

```
# mkdir /DATA_backup  
# mount -t ext3 /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
```



Information

When using GDS Management View, select a proxy volume and execute [Operation]:[Stop Volume] in the Main Screen.

7-2) Rejoin proxy volume Volume2 to master volume Volume1.

```
# sdxproxy Rejoin -c Class1 -p Volume2
```

After returning from the command, synchronization copying from Volume1 to Volume2 is performed.



See

When using GDS Management View, see "Rejoin" of "5.3.2.2 Backup (by Synchronization)."

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
```



See

When using GDS Management View, see "5.5.6 Breaking a Proxy."

6.4.2 Online Backup (Snapshot by OPC)

If GDS Snapshot is installed in a system where a disk array unit with the OPC function is used, backup can be performed through instant snapshot by OPC using proxy volumes without stopping services.


However, to secure consistency of snapshot data, the services must be stopped temporarily when creating a snapshot (when updating proxy volumes).



See

For securing consistency of snapshot data, see "A.2.20 Ensuring Consistency of Snapshot Data."

Figure: 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 "A.2.17 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.2.16 Using the Advanced Copy Function in a Proxy Configuration."

Procedure

1) Relating a proxy volume

Before creating snapshots, relate proxy volume Volume2 as a copy destination to master volume Volume1. In this example, Volume1 and Volume2 belong to class Class1. Execute the following command on a node that belongs to the scope of Class1.

```
# sdxproxy Relate -c Class1 -m Volume1 -p Volume2
```



See

When using GDS Management View, see "5.2.4.2 Relate."

2) Stopping the proxy volume

Stop proxy volume Volume2. If Class1 is a shared class, specify the "-e allnodes" option to stop Volume2 on all nodes.

```
# sdxvolume -F -c Class1 -v Volume2
```



Information

When using GDS Management View, select a proxy volume and execute [Operation]:[Stop Volume] in the Main Screen.

3) Stopping the services

To secure consistency of snapshot data, stop the services before creating a snapshot and restrain the master volume from being written in.

3a) When the master volume is being used for a cluster application

Inactivate the cluster application.

3b) When the master volume is not being used for a cluster application

3b-1) Stop the services for which the master volume is being used.

3b-2) When the master volume is 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.

```
# sdxproxy Update -c Class1 -p Volume2 -e instant
```

When returning from the command, the update of Volume2 will be complete. Subsequently, background OPC physical copying is performed, but you may go on to step 5) without waiting until the copying is complete.



See

When using GDS Management View, see "Update" in "5.3.2.3 Backup (by OPC)." For setting of the update environment, select "Yes" to "Instant Snapshot."

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 ext3 file system on master volume Volume1 is mounted on mount point /DATA.

```
# mount -t ext3 /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
```



Information

When using GDS Management View, select a proxy volume and execute [Operation]:[Start Volume] in the Main Screen.

7) Confirming the completion of copying

Confirm that the copying is complete.

```
# sdxinfo -S -c Class1 -o Volume2
OBJ   CLASS  GROUP  DISK   VOLUME  STATUS
-----
slice Class1  Group2 Disk3  Volume2 ACTIVE
slice Class1  Group2 Disk4  Volume2 ACTIVE
```

If all the displayed slices' STATUS fields are "ACTIVE", copying is complete.
If the copying is still in progress, "COPY" will be displayed in the STATUS field.

8) Backing up to tape

Back up the snapshot data on the proxy volume to tape. Execute the following commands on a node that belongs to the scope of class Class1.



See

For details on the backup methods, see the manuals of file systems to be backed up and used commands.

8a) When backing up data with the dd(1) command

```
# dd if=/dev/sfdsk/Class1/dsk/Volume2 of=/dev/st0 bs=32768
```

8b) When backing up the ext3 file system with the tar(1) command

8b-1) Before mounting

Check and repair consistency of the ext3 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 ext3 /dev/sfdsk/Class1/dsk/Volume2
```

8b-2) Mounting the snapshot

Mount the ext3 file system on proxy volume Volume2 on temporary mount point /DATA_backup.

```
# mkdir /DATA_backup
# mount -t ext3 /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
```



Information

When using GDS Management View, select a proxy volume and execute [Operation]:[Stop Volume] in the Main Screen.

10) Breaking the proxy volume

When no more online backup is to be executed, cancel the relationship between master volume Volume1 and proxy volume Volume2.

```
# sdxproxy Break -c Class1 -p Volume2
```



See

When using GDS Management View, see "5.5.6 Breaking a Proxy."

11) Re-executing online backup

To perform online backup again, go back to step 3).

6.4.3 Instant Restore

By keeping proxy volumes parted from master volumes used for services, in the event of a data error in a master volume, data can be restored back from the proxy volume to the master volume (unless background copying is in process after instant snapshot creation).


The master volume must be stopped temporarily for restoration, but you may start the master volume and make it accessible immediately after the restore is initiated without waiting until the copying is complete.



Information Restore from Tape

If an error occurs in master volume data while the master volume used for the services operation and a proxy volume are joined, the same error occurs in data of the proxy volume and instant restore becomes unavailable. If this happens, restore data back to the master volume following the procedure described in "6.3.3 Restoring."

Figure: Schedule for Instant Restore



Note Instant Restore with Disk Unit's Copy Functions

For restore, only (R)OPC is available, but (R)EC, EMC TimeFinder or EMC SRDF is unavailable. For details, see "A.2.16 Using the Advanced Copy Function in a Proxy Configuration."

Procedure

- 1) Viewing the state of the proxy volume

Confirm that proxy volume Volume2 has been parted from master volume Volume1.

```
# sdxinfo -V -c Class1 -o Volume2 -e long
```

OBJ	NAME	TYPE	CLASS	GROUP	DISK	MASTER	PROXY ...
-----	-----	-----	-----	-----	-----	-----	-----
volume	*	mirror	Class1	Group2	*	*	*
volume	Volume2	mirror	Class1	Group2	*	Volume1	Part
volume	*	mirror	Class1	Group2	*	*	*

If Part is displayed in the PROXY field, the proxy volume has been parted.

If Join is displayed in the PROXY field, the proxy volume has been joined to the master volume and instant restore is unavailable. In this situation, data must be restored from tape. For more details, see "Restore from Tape" described above.

2) Stopping the services

2a) When the master volume is being used for a cluster application
Inactivate the cluster application.

2b) When the master volume is not being used for a cluster application

2b-1) Stop the services for which the master volume is being used.

2b-2) When the master volume is being used as a file system, unmount the file system. This example shows the procedure when the mount point is /DATA.

```
# cd /  
# umount /DATA
```

2b-3) Stopping the master volume

Stop master volume Volume1. If Class1 is a shared class, specify the "-e allnodes" option to stop Volume1 on all nodes.

```
# sdxvolume -F -c Class1 -v Volume1
```



Information

When using GDS Management View, select a master volume and execute [Operation]:[Stop Volume] in the Main Screen.

3) Restoring data from the proxy volume

Execute the following commands on a node that belongs to the scope of class Class1.

3a) When the OPC function is unavailable

3a-1) Stop proxy volume Volume2. If Class1 is a shared class, specify the "-e allnodes" option to stop Volume2 on all nodes.

```
# sdxvolume -F -c Class1 -v Volume2
```



Information

When using GDS Management View, select a proxy volume and execute [Operation]:[Stop Volume] in the Main Screen.

3a-2) Restore data from proxy volume Volume2 back to master volume Volume1.

```
# sdxproxy RejoinRestore -c Class1 -p Volume2 -e instant
```

When returning from the command, the restore of Volume1 will be complete. Subsequently, synchronization copying from Volume2 to Volume1 is performed, and you may go on to step 4) without waiting until the copying is complete.



See

When using GDS Management View, see "Restore" in "5.3.3 Restore." For settings of the restore environment, select "Yes" to "Rejoin" and "Yes" to "Instant Restore."

3b) When the OPC function is available
Restore data from proxy volume Volume2 back to master volume Volume1.

```
# sdxproxy Restore -c Class1 -p Volume2 -e instant
```

When returning from the command, the restore of Volume1 will be complete. Subsequently, background OPC physical copying from Volume2 to Volume1 is performed, and you may go on to step 4) without waiting until the copying is complete.



See

When using GDS Management View, see "5.3.3 Restore." For settings of the restore environment, select "No" to "Rejoin" and "yes" to "Instant Restore."



Note Master Volumes with the Mirroring Multiplicity of Two and Higher

By executing the sdxproxy Restore command in step 3b), OPC starts, copying from one of proxy volume's slices to one of master volume's slices. Among slices of the master volume, slices other than the OPC copy destination are excluded from mirroring and thus the data statuses become invalid (INVALID). To recover the master volume mirroring status, perform master volume resynchronization copying by using the sdxcopy -B command. If not executing the sdxcopy -B command, master volume resynchronization copying automatically starts when starting the master volume in step 4) and data will be copied from the OPC copy destination slice to the other slices with the soft copy function.

4) Resuming the services

Without waiting until the copying is complete, you may resume the services.



Note Reusing Proxy Volume Data

By executing the sdxproxy RejoinRestore command in step 3a), Volume1 and Volume2 are joined and Volume2 will also be updated with data written into Volume1. To reuse data in Volume2 for restore without updating, after the synchronization copying from Volume2 to Volume1 is complete, part Volume2 from Volume1 and then resume the services. When the sdxproxy Restore command was executed in step 3b), Volume1 and Volume2 are left parted, and data in Volume2 remains unchanged even if the services are resumed before the copying is complete.

4a) When the master volume is used for a cluster application
Activate the cluster application.

4b) When the master volume is not used for a cluster application

4b-1) Activate master volume Volume1 on the node running the services.

```
# sdxvolume -N -c Class1 -v Volume1
```



Information

When using GDS Management View, select the master volume and execute [Operation]:[Start Volume] in the Main Screen.



Information When the OPC Function Is Available

If the master volume mirroring multiplicity is two and higher and if restore is performed with the OPC function in step 3b), master volume resynchronization copying automatically starts after 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 ext3 file system

```
# mount -t ext3 /dev/sfdsk/Class1/dsk/Volume1 /DATA
```

4b-3) Resume the servicess using the master volume.

5) Viewing the copy status

The status of the copying from proxy volume Volume2 to master volume Volume1 started in step 3) can be viewed by using the `sdxinfo -S` command. The copy destination slice is in the COPY status if copying is in progress and it will be in the ACTIVE status after the copy process ends normally.

```
# sdxinfo -S -c Class1 -o Volume1
OBJ   CLASS  GROUP  DISK   VOLUME  STATUS
----   -
slice Class1  Group1 Disk1   Volume1 ACTIVE
slice Class1  Group1 Disk2   Volume1 ACTIVE
```



Information

By executing the `sdxproxy RejoinRestore` command in step 3a), master volume Volume1 and proxy volume Volume2 are joined. In the even of a data error occurs in Volume1, the same data error occurs in Volume2, and restoring data back from tape will be required. Therefore, after resynchronization copying from Volume2 to Volume1 is complete, it is recommended to part Volume2 from Volume1. For the procedures for parting proxy volumes, see the steps 3) through 5) in "6.4.1 Online Backup (by Synchronization)."



Note When Using (R)EC, EMC TimeFinder, EMC SRDF, for Backup

Restoration of the master by use of the proxy stops sessions of these copying functions. To perform backup using these copying functions, cancel the relation between the master and the proxy with the `sdxproxy Break` command.

6.5 Backing Up and Restoring through Disk Unit's Copy Functions [Linux2.6]

Some sophisticated disk devices contain hardware functions to copy disk data within the disk units or to other disk units. For example, Fujitsu Storage Systems ETERNUS provide the Advanced Copy function and EMC's Symmetrix storage systems provide copy functions such as TimeFinder and SRDF.

This section describes the procedures for backing up and restoring object configurations and data of local disks and shared disks through use of these disk units' 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.

6.5.1 Configuring an 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 "Appendix H Shared Disk Unit 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.

6.5.2 Backing Up

3) Backing up the object configuration of the class

On Node1 or Node2 in the primary domain, back up the object configuration of Class1 to be backed up.

3-1) Saving configuration information

Save outputs of the `sdxinfo` command to a file. In this example, the path to a file is `/var/tmp/Class1.info`.

```
# sdxinfo -c Class1 -e long > /var/tmp/Class1.info
```

3-2) Creating a configuration file

Output the object configuration within Class1 to a file in configuration table format. In this example, the path to a file is `/var/tmp/Class1.conf`.

```
# sdxconfig Backup -c Class1 -o /var/tmp/Class1.conf
```

3-3) Backing up the configuration information and configuration file

Save the files created in steps 3-1) and 3-2) to tape and so on.

4) Detaching the backup disks (suspending synchronization)



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 "A.2.20 Ensuring Consistency of Snapshot Data."

4-1) Viewing the application volume status

Confirm that data of slices comprising application volume Volume1 is in valid status (ACTIVE or STOP).

```
# sdxinfo -S -c Class1 -o Volume1
```

OBJ	CLASS	GROUP	DISK	VOLUME	STATUS
slice	Class1	Group1	Disk1	Volume1	ACTIVE
slice	Class1	Group1	Disk2	Volume1	ACTIVE

If it is not in valid status (ACTIVE or STOP), recover the slice status, referring to "F.1.1 Slice Status Abnormality."

4-2) Checking the statuses of the disk unit's copy functions

Confirm that the application disks and the backup disks are synchronized.



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
```



Note 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.



See

For details on the backup methods, see the manuals of file systems to be backed up and used commands.

6a) When backing up data with the dd(1) command

```
# dd if=/dev/sfdsk/Class2/dsk/Volume1 of=/dev/st0 bs=32768
```

6b) When backing up the ext3 file system with the tar(1) command

6b-1) Check and repair consistency of the ext3 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 ext3 /dev/sfdsk/Class2/dsk/Volume1
```

6b-2) Mount the ext3 file system on backup volume Volume1 on /mnt1, a temporary mount point, in read only mode.

```
# mkdir /mnt1
# mount -t ext3 -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.

6.5.3 Restoring from Backup Disks

9) Stopping services

Stop applications using application volume Volume1 on nodes Node1 and Node2 in the primary domain. When using Volume1 as a file system, unmount the file system.

10) Stopping the application volume

On Node1 and Node2 in the primary domain, stop application volume Volume1. Execute the following command on Node1 or Node2

```
# sdxvolume -F -c Class1 -v Volume1 -e allnodes
```

11) Deleting the class

In the primary domain, delete class Class1 to which application volume Volume1 belongs. Execute the following command on Node1 or Node2 in the primary domain.

```
# sdxvolume -R -c Class1 -v Volume1
# sdxgroup -R -c Class1 -g Group1
# sdxdisk R -c Class1 -d Disk1
# sdxdisk R -c Class1 -d Disk2
```

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 "6.5.2 Backing Up," restore the object configuration of Class1.

After restoring the object configuration, reboot the restored node.

```
# sdxconfig Restore -c Class1 -i /var/tmp/Class1.conf
# shutdown -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 "6.5.2 Backing Up." In this example, the scope of backed up class Class1 is Node1:Node2.

15-1) Stop the volume in the class.

```
# sdxvolume -F -c Class1
```

15-2) Change the class type and expand the class scope.

```
# sdxattr -C -c Class1 -a type=shared,scope=Node1:Node2
```

16) Starting the application volume

On Node1 or Node2 in the primary domain, start application volume Volume1. Execute the following command on Node1 or Node2.

```
# sdxvolume -N -c Class1 -v Volume1
```

17) Resuming the services

When the file system on application volume Volume1 was unmounted in step 8) of "6.5.2 Backing up," unmount it again.

Start the applications using Volume1.

6.6 Backing Up and Restoring through an External Server

This section discusses the method of backing up data from and restoring data back to logical volumes (called application volumes in this manual) in a local or shared class in the primary domain through a server in a domain different from the primary domain (called an external backup server in this manual).

The backup and restore operations through an external backup server can be categorized into 4 patterns.

1. Backing up and restoring a logical volume with no replication



See

For details, see "6.6.1."

2. Backing up and restoring through snapshot by slice detachment



See

For details, see "6.6.2."

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



See

For details, see "6.6.3."

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



See

For details, see "6.6.4."

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
			Advanced Copy function	One of (*5): single mirror	
			EMC TimeFinder or EMC SRDF (*6)	One of (*6): mirror	
4	Available	Non-SDX disks as the disk unit's copy function destinations	EMC TimeFinder or EMC SRDF (*7)	Any	GDS

(*1):

The table above describes the volume types according to the following classification.

Type	Description
single	Single volume created in a single disk.
mirror	Mirror volume connected to a mirror group to which one or more disks are connected. This is excluded when a lower level group is connected to the mirror group. Making disk data redundant can improve the services continuity.
concat	Volume created in a concatenation group. When no large volume exists, it can be created by concatenating multiple disks.
stripe	Stripe volume created in a stripe group. I/O loads of the services can be balanced up on multiple disks.
concat + mirror	Mirror volume created in a mirror group to which one or more concatenation groups are connected. This is excluded when a disk is connected to the mirror group. The effects of both concat and mirror can be gained.
stripe + mirror	Mirror volume created in a mirror group to which one or more stripe group are connected. This is excluded when a disk is connected to the mirror group. The effects of both stripe and mirror can be gained.
(concat + mirror)	Mirror volume created in a mirror group to which one disk and one or more concatenation groups are connected. Since the size of this volume is limited by the size of the disk connected to the mirror group, the effects of concat cannot be gained.
(stripe + mirror)	Mirror volume created in a mirror group to which one disk and one or more stripe group are connected. Since I/O loads to the disk connected to the mirror group are not balanced up, the effects of stripe cannot be gained. However, while the slice is temporarily detached from the disk, the effects of stripe can be gained since I/O loads to the volume are balanced up on multiple disks connected to the stripe group. Note, however, that a slice cannot be temporarily detached from a master volume related to a proxy volume.

(*2):

For an external backup server, GDS and GDS Snapshot must be installed for creating shadow volumes.

(*3):

See "A.2.10 Creating a Snapshot by Slice Detachment."

(*4):

See "A.1.8 Exception to Proxy Configuration."

(*5):

See "A.1.8 Exception to Proxy Configuration" and "A.2.16 Using the Advanced Copy Function in a Proxy Configuration."

(*6):

See "A.1.8 Exception to Proxy Configuration" and "A.2.19 Using EMC TimeFinder or EMC SRDF in a Proxy Configuration [Linux2.6]."

(*7):

There are two operation patterns that use EMC TimeFinder or EMC SRDF. The features of respective types are as follows.

- Backup and restore using snapshots from proxy volumes
Snapshot operations are available through only GDS and GDS Snapshot commands without using a SYMCLI command. When striping or concatenation has been applied to the application volume, copying to proxy volumes is conducted by the soft copy function instead of TimeFinder and SRDF.
- Backup and restore using the hard copy function
Disk areas that are copy destinations of TimeFinder or SRDF can be backed up, regardless of the type of the application volume.

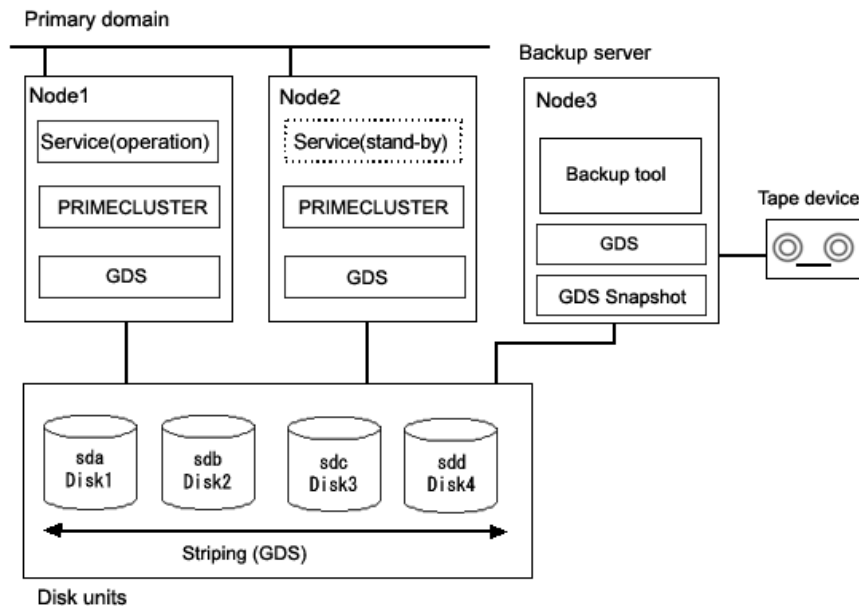
6.6.1 Backing Up and Restoring a Logical Volume with No Replication

This sub-section describes the method of backing up data from and restoring data back to logical volumes in the primary domain through a backup server in another domain.

The following is an example of backing up and restoring a stripe volume. Mirror volumes, single volumes, and volumes in concatenation groups can also be backed up and restored in a similar manner. However, for backing up mirror volumes, the method described in “6.6.2 Backing Up and Restoring through Snapshot by Slice Detachment” is recommended.

6.6.1.1 System Configuration

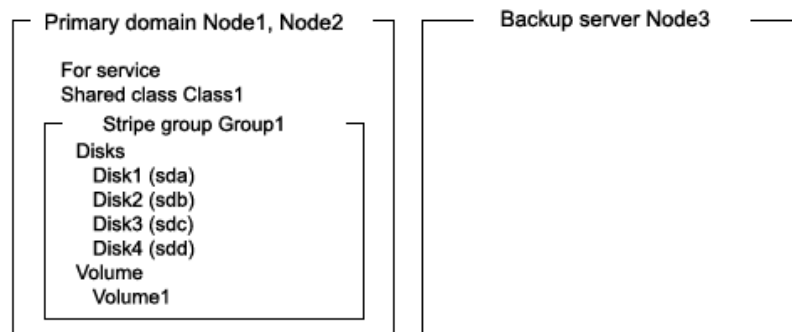
Figure: System Configuration



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: Object Configuration in Normal Operation



6.6.1.2 Summary of Backup

Backups can be created while the services are stopped and the application volume is not in use.

Figure: Backup

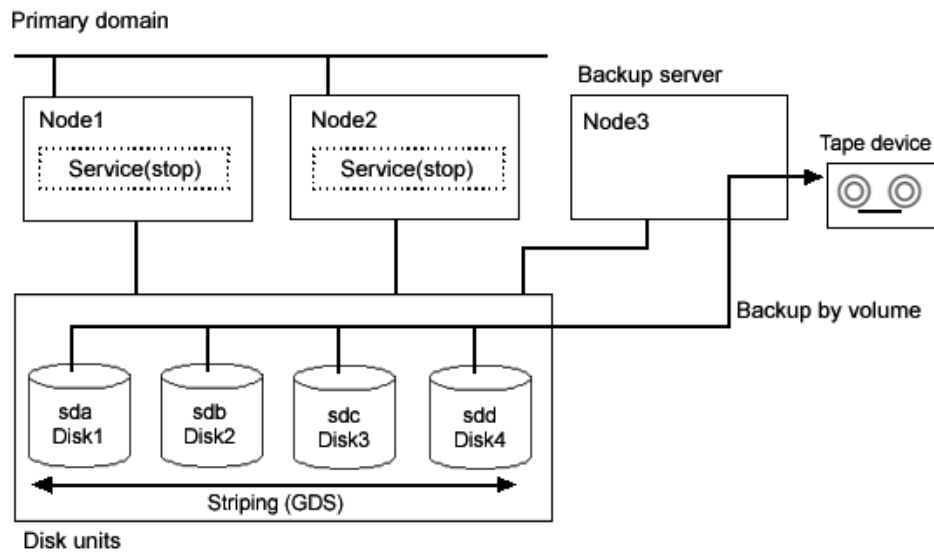


Figure: Object Configuration When Backing Up

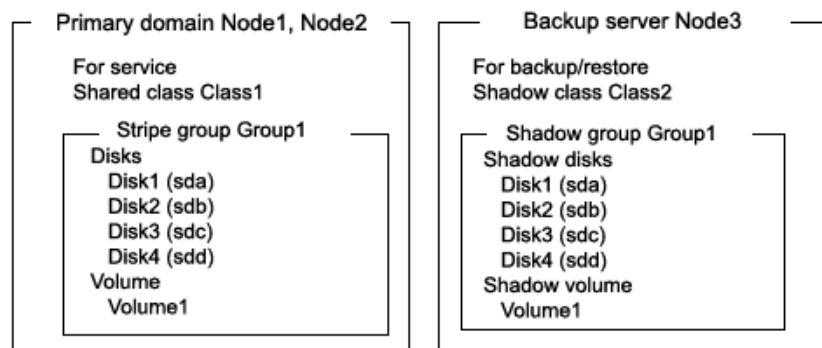



Figure: Backup Schedule



6.6.1.3 Summary of Restore

If volume data is damaged, it can be restored from tape. Data can be restored while the services are stopped and the application volume is not in use.

Figure: Restore

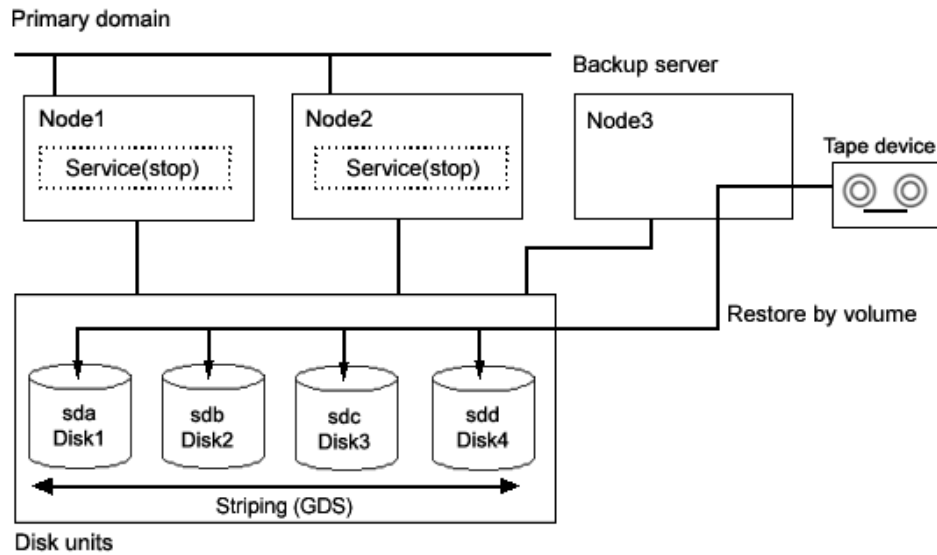


Figure: Object Configuration When Restoring

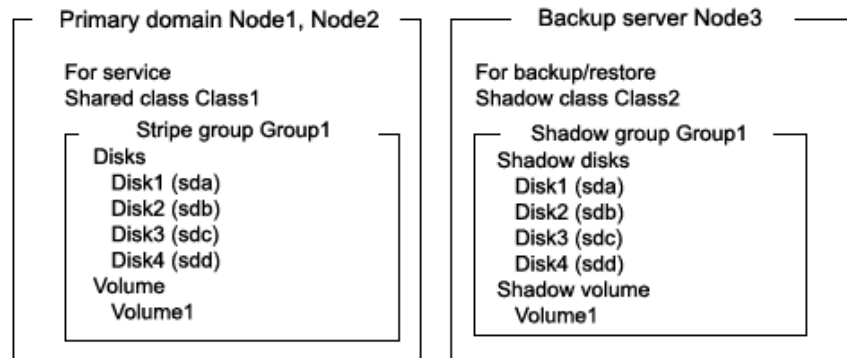
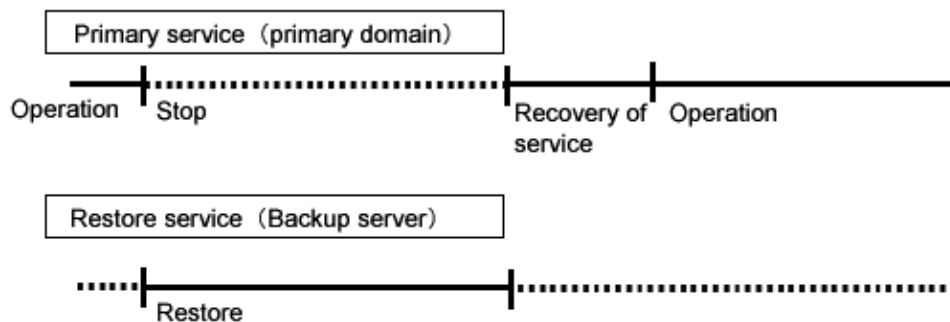


Figure: Restore Schedule



6.6.1.4 Summary of Procedure

Figure: Outline of the Configuration Procedure



Figure: Outline of the Backup Procedure

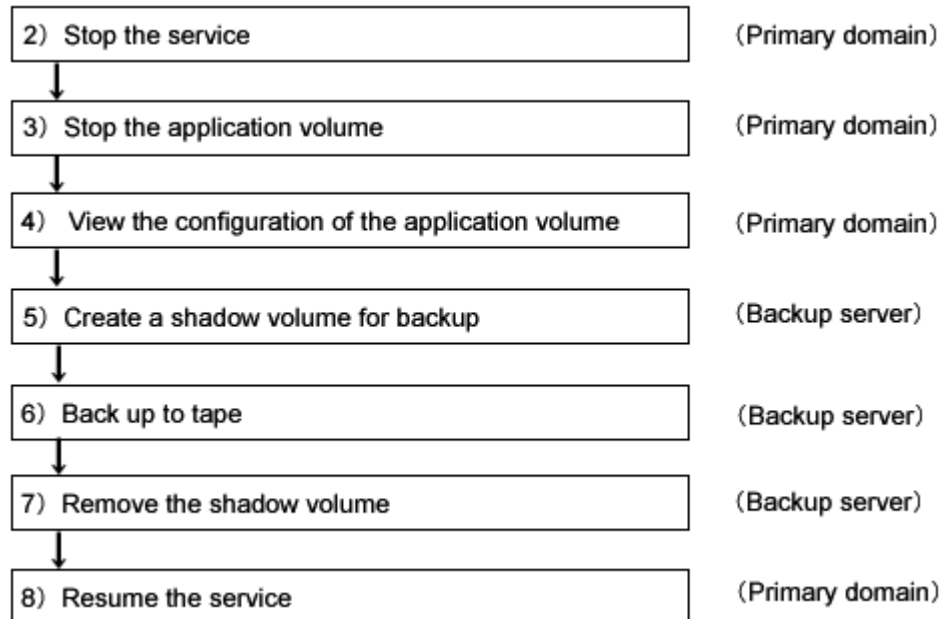



Figure: Outline of the Restore Procedure



6.6.1.5 Configuring an 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 “Appendix H Shared Disk Unit 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,svc=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
```

6.6.1.6 Backing Up

2) Stopping the services

Exit all applications accessing the application volume Volume1 in the primary domain on Node1 and Node2.

When Volume1 is used as a file system, it should be unmounted.

3) Stopping the application volume

To write-lock volume Volume1, inactivate Volume1 on Node1 and Node2 in the primary domain. Execute the following command on Node1 or Node2.

```
# sdxvolume -F -c Class1 -v Volume1 -e allnodes
```

4) Viewing the configuration of the application volume

On Node1 or Node2 in the primary domain, view the configuration of application volume Volume1 that is the backup target.

```
# sdxinfo -c Class1
```

OBJ	NAME	TYPE	SCOPE	SPARE					
class	Class1	shared	Node1:Node2	0					
OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	DEVCONNECT	STATUS	
disk	<u>Disk1</u>	stripe	Class1	Group1	sda	8380800	Node1:Node2	ENABLE	
disk	<u>Disk2</u>	stripe	Class1	Group1	sdb	8380800	Node1:Node2	ENABLE	
disk	<u>Disk3</u>	stripe	Class1	Group1	sdc	8380800	Node1:Node2	ENABLE	
disk	<u>Disk4</u>	stripe	Class1	Group1	sdd	8380800	Node1:Node2	ENABLE	
OBJ	NAME	CLASS	DISKS	BLKS	FREEBLKS	SPARE			
group	Group1	Class1	<u>Disk1:Disk2:Disk3:Disk4</u>	32964608	31850496	*			
OBJ	NAME	CLASS	GROUP	SKIP	JRM	1STBLK	LASTBLK	BLOCKS	STATUS
volume	*	Class1	Group1	*	*	0	65535	65536	PRIVATE
volume	Volume1	Class1	Group1	*	*	<u>65536</u>	1114111	<u>1048576</u>	STOP
volume	*	Class1	Group1	*	*	1114112	32964607	31850496	FREE
OBJ	CLASS	GROUP	DISK	VOLUME	STATUS				
slice	Class1	Group1	*	Volume1	STOP				

If application volume Volume1 belongs to a stripe group, also pay attention to the stripe width.

```
# sdxinfo -G -c Class1 -o Group1 -e long
```

OBJ	NAME	CLASS	DISKS	BLKS	FREEBLKS	SPARE	MASTER	TYPE	WIDTH
group	Group1	Class1	Disk1:Disk2:Disk3:Disk4	32964608	31850496	*	*	stripe	<u>256</u>

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.

```
# sdxinfo -c Class2
```

OBJ	NAME	TYPE	SCOPE	SPARE					
class	Class2	local	Node3	0					
OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	DEVCONNECT	STATUS	
disk	<u>Disk1</u>	stripe	Class2	Group1	sda	8380800	Node3	ENABLE	
disk	<u>Disk2</u>	stripe	Class2	Group1	sdb	8380800	Node3	ENABLE	
disk	<u>Disk3</u>	stripe	Class2	Group1	sdc	8380800	Node3	ENABLE	
disk	<u>Disk4</u>	stripe	Class2	Group1	sdd	8380800	Node3	ENABLE	
OBJ	NAME	CLASS	DISKS	BLKS	FREEBLKS	SPARE			
group	Group1	Class2	<u>Disk1:Disk2:Disk3:Disk4</u>	32964608	31850496	*			
OBJ	NAME	CLASS	GROUP	SKIP	JRM	1STBLK	LASTBLK	BLOCKS	STATUS
volume	*	Class2	Group1	*	*	0	65535	65536	PRIVATE
volume	Volume1	Class2	Group1	*	*	<u>65536</u>	1114111	<u>1048576</u>	ACTIVE
volume	*	Class2	Group1	*	*	1114112	32964607	31850496	FREE
OBJ	CLASS	GROUP	DISK	VOLUME	STATUS				
slice	Class2	Group1	*	Volume1	ACTIVE				

For a stripe volume, see also the stripe width.

```
# sdxinfo -G -c Class2 -o Group1 -e long
```

OBJ	NAME	CLASS	DISKS	BLKS	FREEBLKS	SPARE	MASTER	TYPE	WIDTH
group	Group1	Class2	Disk1:Disk2:Disk3:Disk4	32964608	31850496	*	*	stripe	<u>256</u>

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`.



See

For details on the backup method, see the manuals of file systems to be backed up and used commands.



Information In a GFS Shared File System

Back up through the method as described in step 6a).

6a) When backing up data with the `dd(1)` command

```
# dd if=/dev/sfdsk/Class2/dsk/Volume1 of=/dev/st0 bs=32768
```

6b) When backing up the ext3 file system with the tar(1) command

6b-1) Mount the ext3 file system on shadow volume Volume1 on /mnt1, a temporary mount point, in read only mode.

```
# mkdir /mnt1
# mount -t ext3 -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.

6.6.1.7 Restoring

9) Stopping the services

Exit all applications accessing application volume Volume1 on Node1 and Node2 in the primary domain. When Volume1 is used as a file system, it should be unmounted.

10) Stopping the application volume

To write-lock application volume Volume1, inactivate Volume1 on Node1 and Node2 in the primary domain. Execute the following command on Node1 or Node2.

```
# sdxvolume -F -c Class1 -v Volume1 -e allnodes
```

11) Viewing the configuration and status of the application volume

On Node1 or Node2 in the primary domain, view the configuration and status of application volume Volume1 that is the restore target. Confirm that Volume1 is in the STOP status. If the volume status is invalid, repair it referencing to "F.1.3 Volume Status Abnormality"

```
# sdxinfo -c Class1
OBJ  NAME  TYPE  SCOPE  SPARE
----  ----  ----  -----  ----
class Class1  shared  Node1:Node2  0

OBJ  NAME  TYPE  CLASS  GROUP  DEVNAM  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.

```
# sdxinfo -G -c Class1 -o Group1 -e long
OBJ  NAME  CLASS  DISKS  BLKS  FREEBLKS  SPARE  MASTER  TYPE  WIDTH
----  ----  ----  -----  ----  ----  ----  ----  ----  ----
group Group1  Class1  Disk1:Disk2:Disk3:Disk4  32964608  31850496  *  *  stripe  256
```

12) Creating a shadow volume for restore

On backup server Node3, create a volume for restore (shadow volume) in the same configuration as the application volume found in step 11). The following settings are necessary on backup server Node3. A shadow volume for restore and a shadow volume for backup are common. When it has already been created, simply change the access mode as described in step 12-4).



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.

```
# 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
----  ----  ----  -----  ----  ----  ----  ----  ----  ----
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 ext3 file system with the tar(1) command

13b-1) Create the ext3 file system on Volume1.

```
# mkfs -t ext3 /dev/sfdsk/Class2/dsk/Volume1
```

13b-2) Mount the ext3 file system on shadow volume Volume1 on /mnt1, a temporary mount point.

```
# mkdir /mnt1
# mount -t ext3 /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

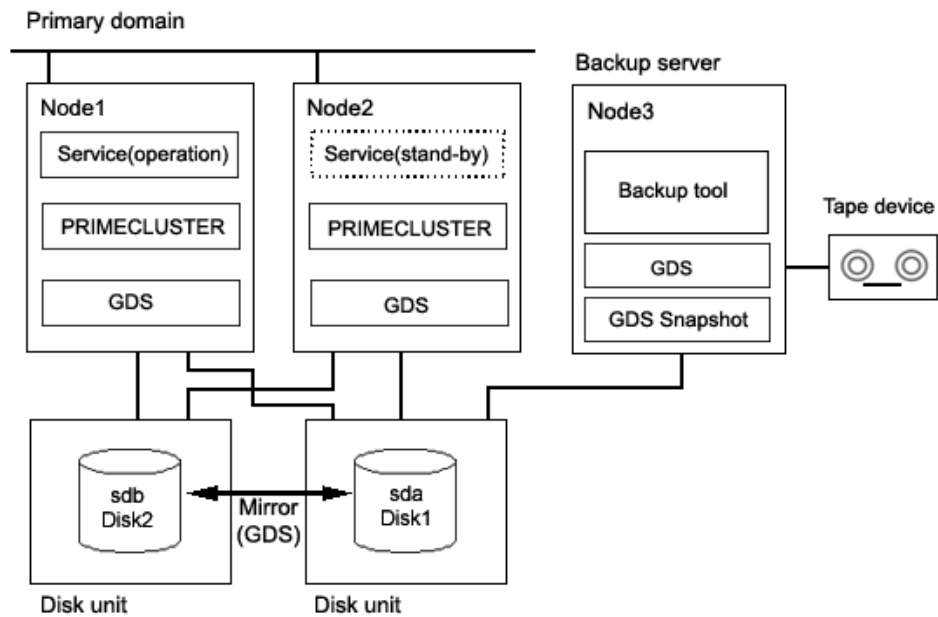
When the file system on application volume Volume1 was unmounted in step 9), mount it again. Start the applications using Volume1.

6.6.2 Backing Up and Restoring through Snapshot by Slice Detachment

This sub-section describes the method of backing up data from and restoring data back to mirror volumes in the primary domain through a backup server in another domain.

6.6.2.1 System Configuration

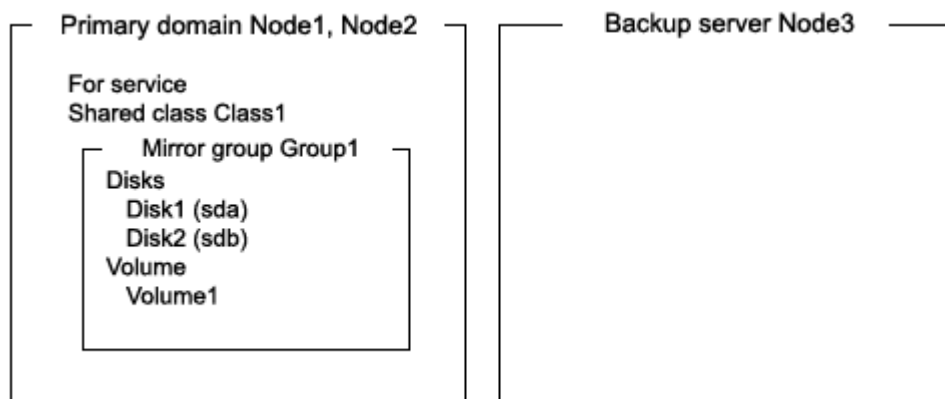
Figure: System Configuration



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: Object Configuration in Normal Operation



6.6.2.2 Summary of Backup

Data in a slice temporarily detached from a volume can be backed up to tape during the service operation.

To secure consistency of data in a detached slice, the services must be stopped temporarily when detaching the slice.



Information Consistency of Snapshot Data

When detaching a slice while the services are operating, data consistency must be secured through the method specific to that software, such as a file system and a database system, which manages volume data. For details, see "A.2.20 Ensuring Consistency of Snapshot Data."

Figure: Backup




Figure: Object Configuration When Backing Up

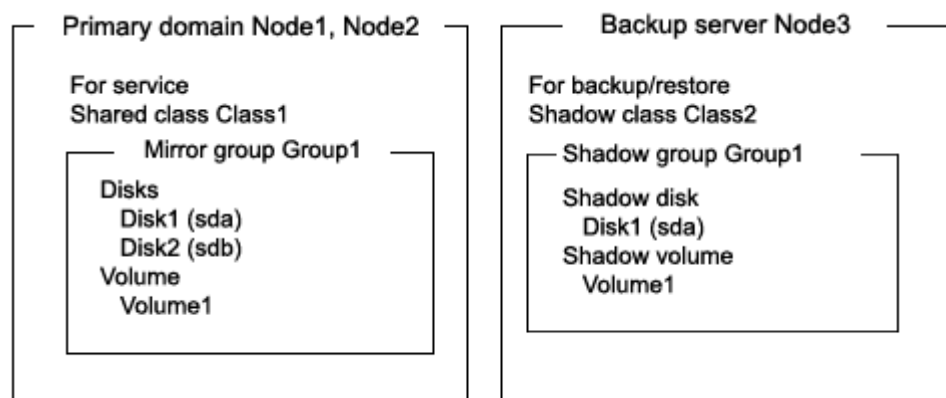
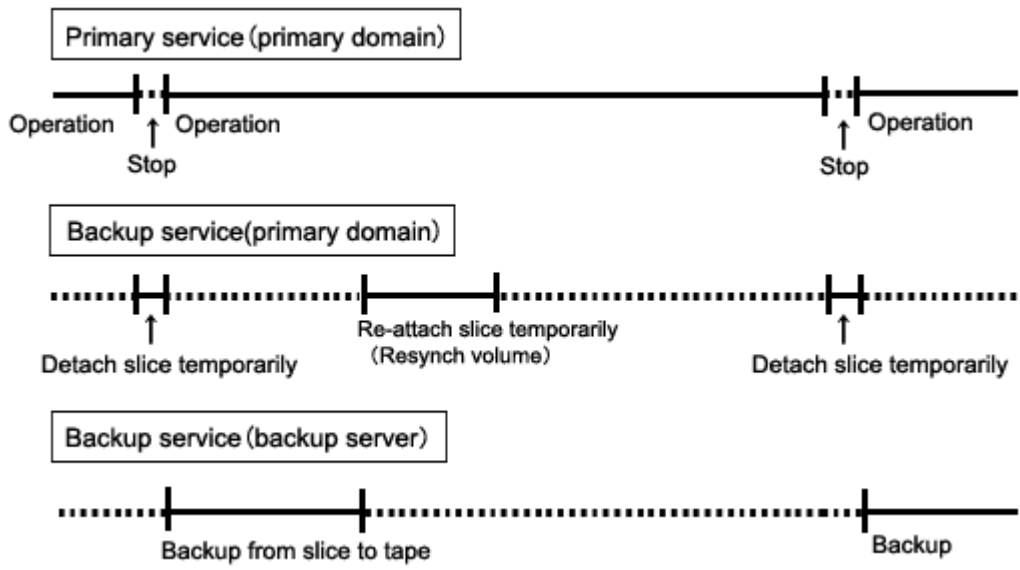


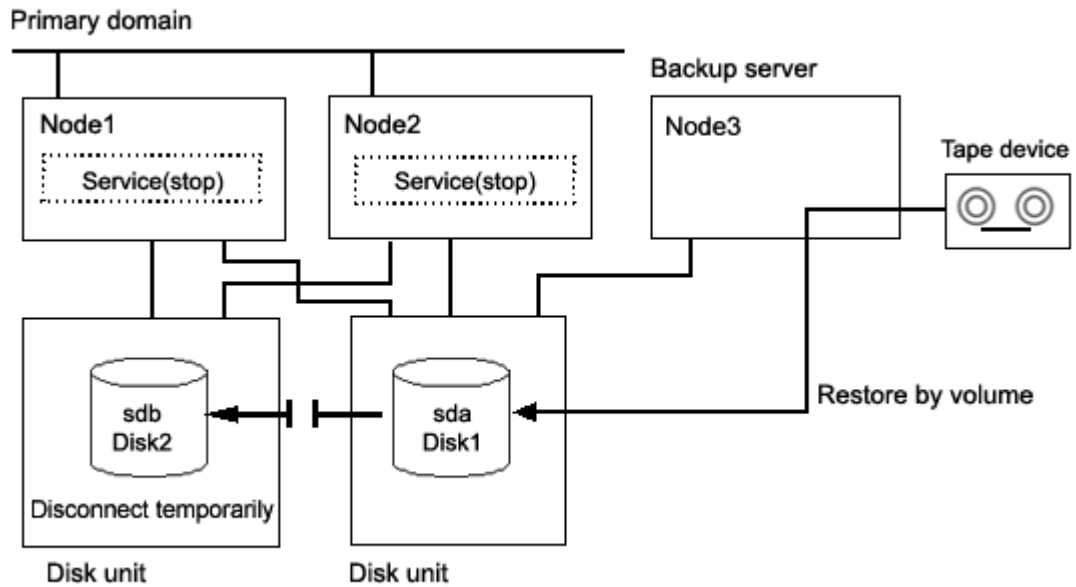
Figure: Backup Schedule



6.6.2.3 Summary of Restore

If volume data is damaged, it can be restored from tape.
Data can be restored while service is stopped and the application volume is not in use.

Figure: 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 "6.6.1 Backing Up and Restoring a Logical Volume with No Replication."

Figure: Object Configuration When Restoring

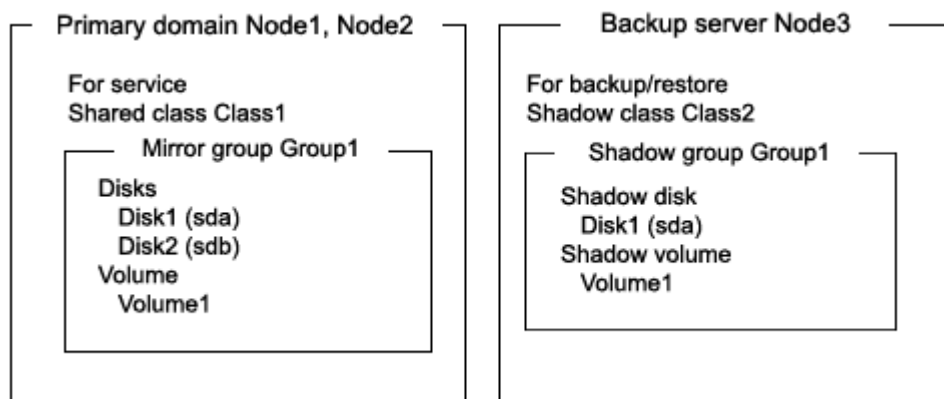
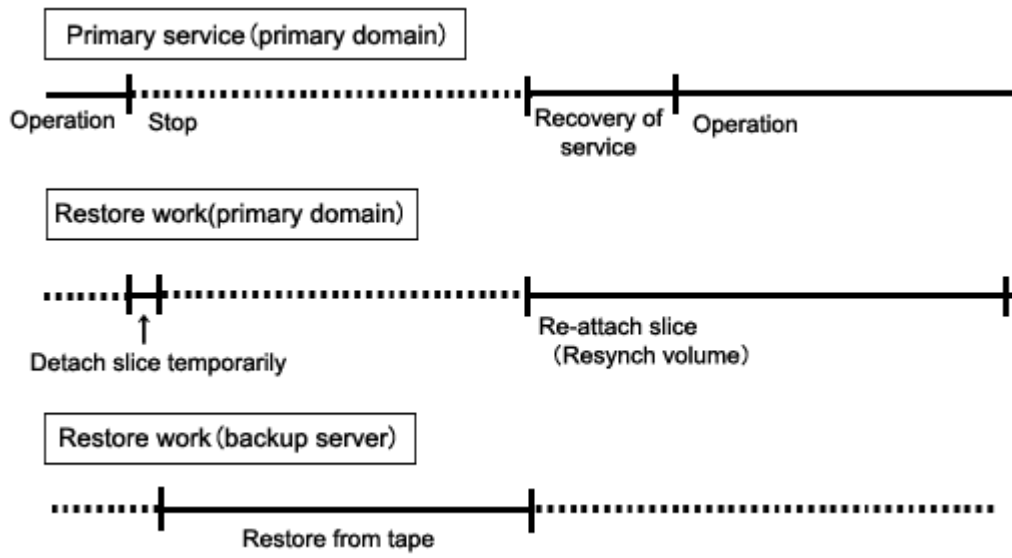


Figure: Restore Schedule



6.6.2.4 Summary of Procedure

Figure: Outline of the Configuration Procedure



Figure: Outline of the Backup Procedure

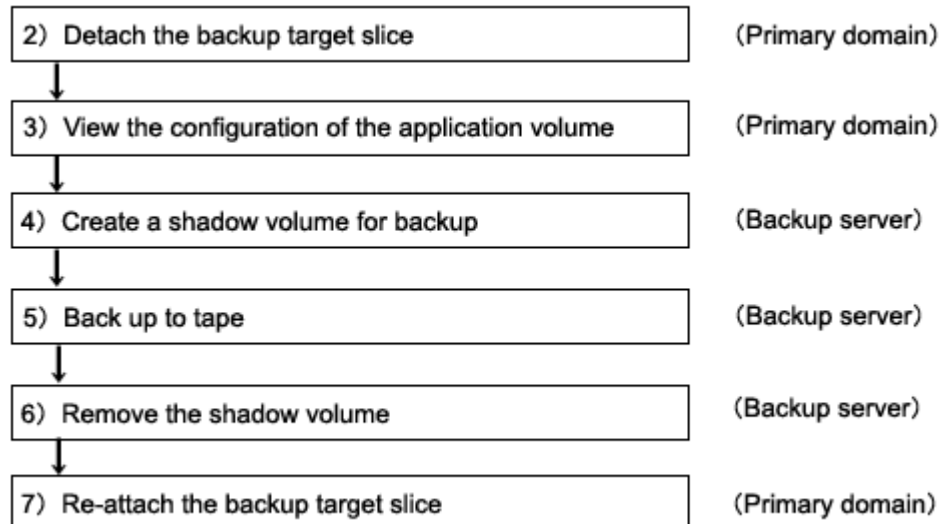
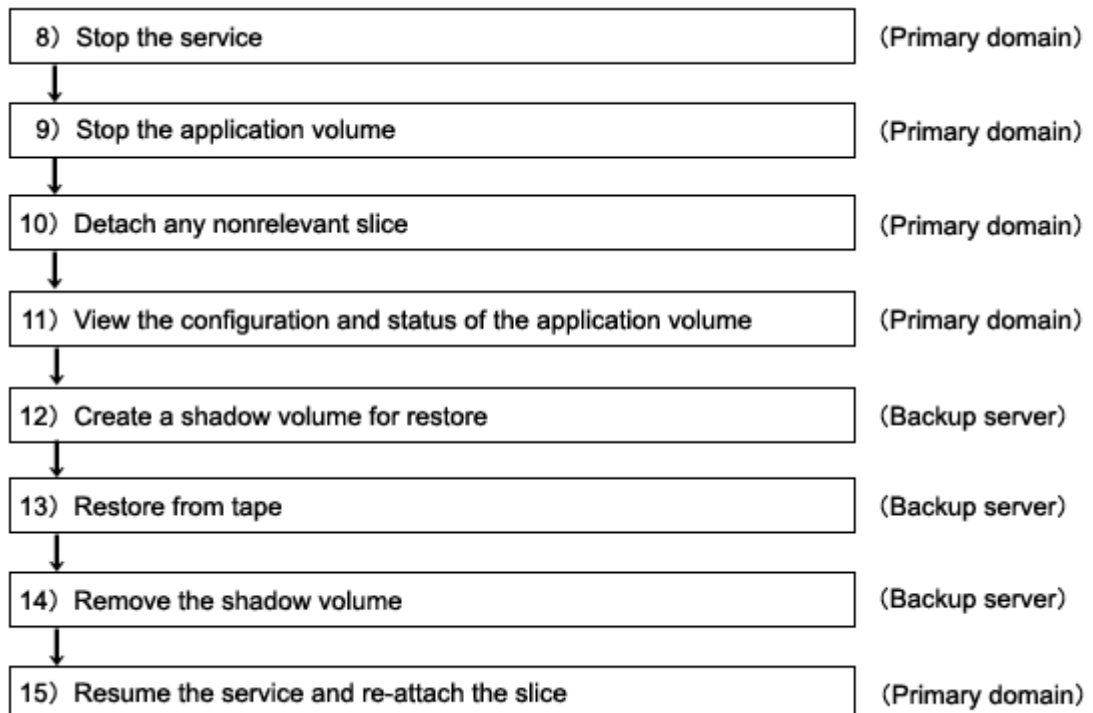


Figure: Outline of the Restore Procedure



6.6.2.5 Configuring an 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 “Appendix H Shared Disk Unit 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
```


6.6.2.6 Backing Up

2) Detaching the backup target slice

Temporarily detach the slice on Disk1 that is the backup target, among slices in application volume Volume1. The following procedure must be performed on Node1 or Node2 in the primary domain.



Information

The following example secures data consistency by stopping the services when a slice is detached. Steps 2-1) and 2-3) are not required if your software, such as a file system and a database system, that manages volume data provides functionality ensuring data consistency or repairing consistency for a detached slice. Alternatively, data consistency must be secured with the method specific to that software. For details, see "A.2.20 Ensuring Consistency of Snapshot Data."

2-1) Stopping the services

To secure consistency of data in a detached slice, exit all applications accessing application volume Volume1 on Node1 and Node2.

When Volume1 is used as a file system, it should be unmounted.

2-2) Detaching the slice

Temporarily detach the slice on disk Disk1 from Volume1. To write-lock the detached slice, set the access mode of the slice to ro (read only).

```
# sdxslice -M -c Class1 -d Disk1 -v Volume1 -a jrm=off,mode=ro
```



Note **Just Resynchronization Mode for Silice**

On backup server Node3, data may be written from Node3 into Disk1 when data in Disk1 is backed up to tape. GDS in the primary domain cannot recognize the write occurrence from Node3. Consequently, if the JRM mode of the detached slice is "on", the portions updated from Node3 may not be involved in resynchronization copying performed when the slice is reattached. If this happens, synchronization of Volume1 is no longer ensured. For this reason, the JRM mode of a detached slice must be set to off in advance.

2-3) Resuming the services

When the file system was unmounted in step 2-1), mount it again.

Resume the application stopped in step 2-1).

3) Viewing the configuration of the application volume

On Node1 or Node2 in the primary domain, view the configuration of application volume Volume1 that is the backup target.

```
# sdxinfo -c Class1
```

OBJ	NAME	TYPE	SCOPE	SPARE					
class	Class1	shared	Node1:Node2	0					
OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	DEVCONNECT	STATUS	
disk	Disk1	mirror	Class1	Group1	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.



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 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.

```
# 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				

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.



See

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 ext3 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 ext3 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 ext3 -y /dev/sfdsk/Class2/dsk/Volume1
```

5b-3) Mount the ext3 file system on shadow volume Volume1 on /mnt1, a temporary mount point, in the read only mode.

```
# mkdir /mnt1
# mount -t ext3 -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).

# <code>sdxinfo -S -c Class1 -o Volume1</code>					
OBJ	CLASS	GROUP	DISK	VOLUME	STATUS
-----	-----	-----	-----	-----	-----
slice	Class1	Group1	Disk1	Volume1	ACTIVE
slice	Class1	Group1	Disk2	Volume1	COPY

6.6.2.7 Restoring



Information

If access can be gained from backup server Node3 to all of the disks constituting Volume1 (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 “F.1.3 Volume Status Abnormality” and “F.1.1 Slice Status Abnormality.”

```
# sdxinfo -c Class1
```

OBJ	NAME	TYPE	SCOPE	SPARE					
class	Class1	shared	Node1:Node2	0					
OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	DEVCONNECT	STATUS	
disk	Disk1	mirror	Class1	Group1	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 restore

On backup server Node3, create a volume for restore (shadow volume) on disk sda. The following settings are necessary on backup server Node3. A shadow volume for restore and a shadow volume for backup are common. If one already exists, this procedure is not required.



Note

Application volume data may be damaged if data is written into a shadow volume in incorrect configuration. Be sure to confirm that the shadow volume configuration is correct in step 12-5).



Information

If access can be gained from backup server Node3 to all of the disks constituting Volume1 (sda and sdb) a shadow volume for restore must be created in the same configuration as Volume1. Under these circumstances, those shadow volumes for restore and backup are not common.

12-1) Registering a shadow disk

Register disk 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 3).
- 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.

```
# 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 ext3 file system with the `tar(1)` command

13b-1) Create the ext3 file system to shadow volume Volume1.

```
# mkfs -t ext3 /dev/sfdsk/Class2/dsk/Volume1
```

13b-2) Mount the ext3 file system on shadow volume Volume1 on /mnt1, a temporary mount point.

```
# mkdir /mnt1
# mount -t ext3 /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

6.6.3 Backing Up and Restoring Using Snapshots from a Proxy Volume

This sub-section describes the method of backing up data from and restoring data back to logical volumes in the primary domain through a backup server in another domain by use of snapshots from the proxy volume.

6.6.3.1 System Configuration

Figure: System Configuration

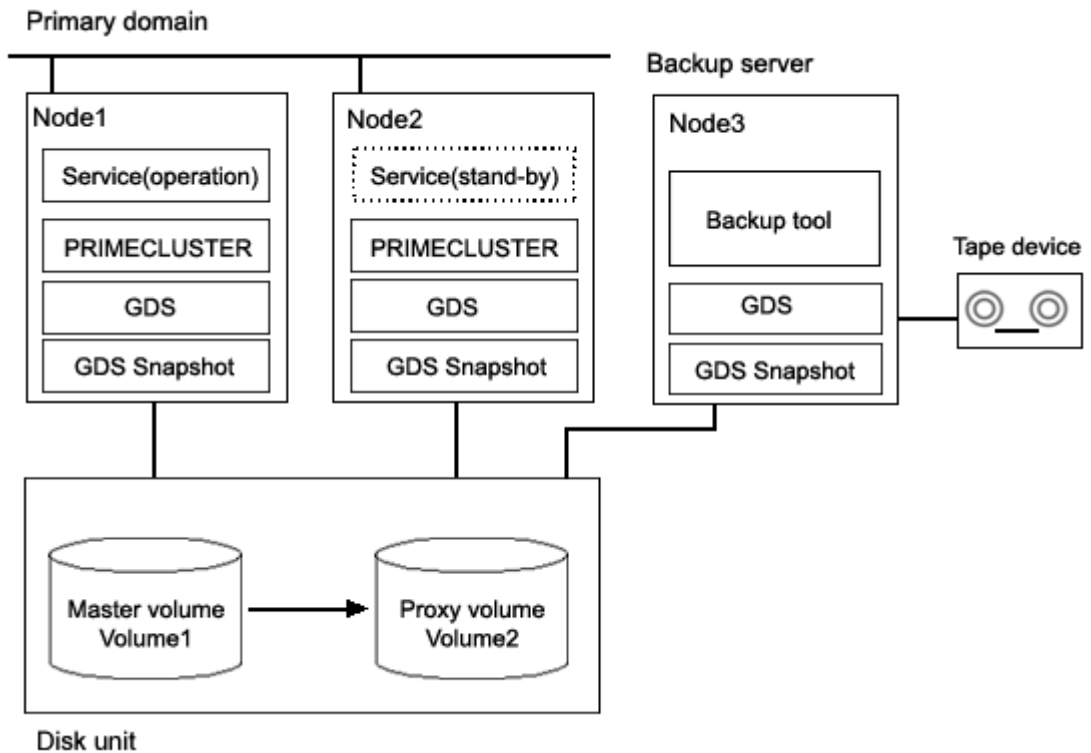



Figure: Object Configuration in Normal Operation



6.6.3.2 Summary of Backup

Data in a proxy volume parted from the master volume during the services operation can be backed up to tape.

To secure consistency of data in the proxy volume to be parted, the services should be stopped temporarily when it is parted.



Information Consistency of Snapshot Data

When detaching a proxy volume while the service is operating, data consistency must be secured through the method specific to that software, such as a file system and a database system, which manages volume data. For details, see "A.2.20 Ensuring Consistency of Snapshot Data."

Figure: Backup

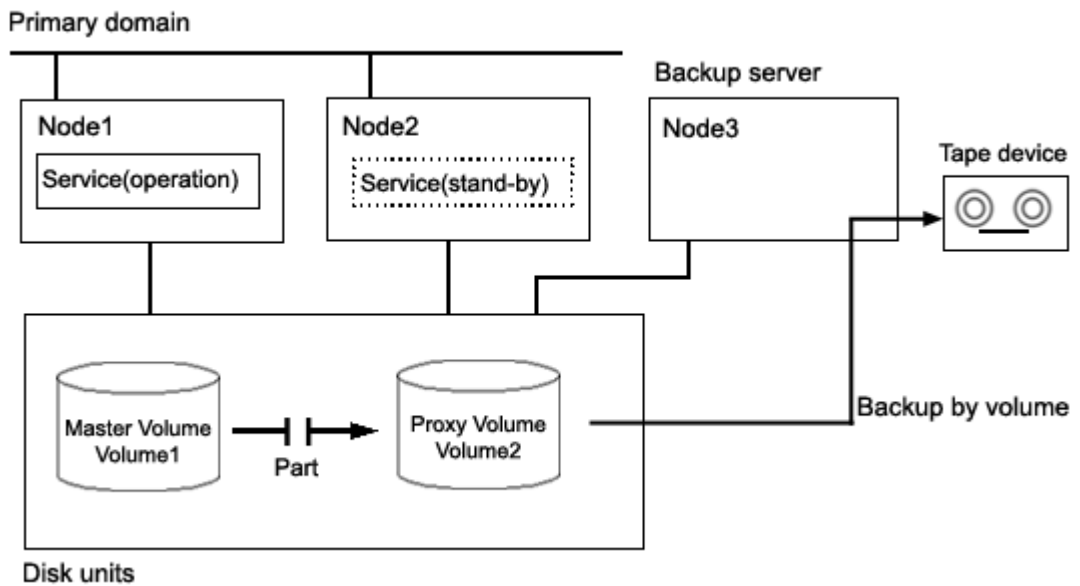


Figure: Object Configuration When Backing Up

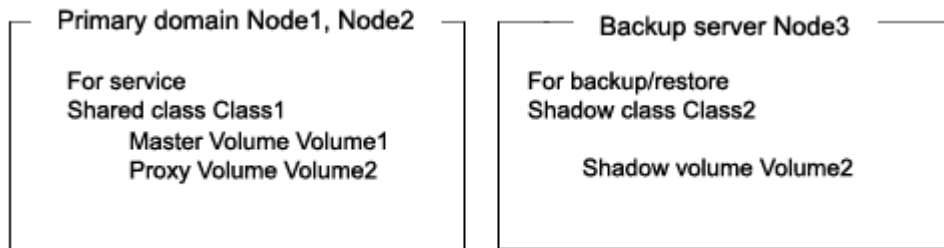
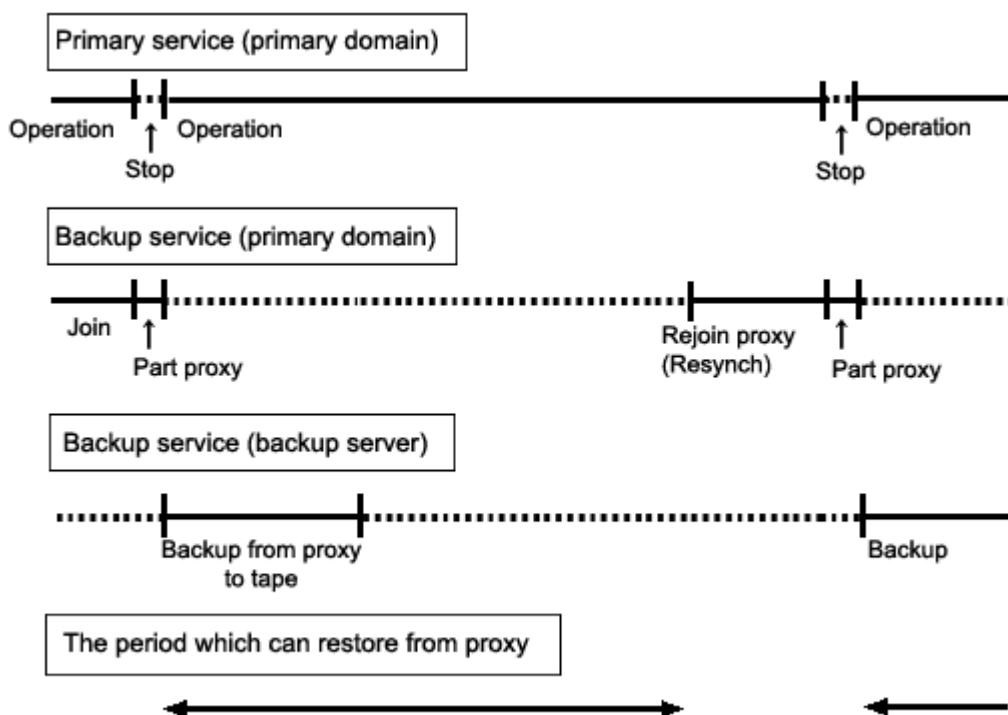


Figure: Backup Schedule



6.6.3.3 Summary of Restore from a Proxy Volume

If master volume data is damaged while a proxy volume is parted from the master volume used for the services operation, data can be restored from the proxy volume back to the master volume.

When restoring, access to the volume must be suspended temporarily.

Figure: Restore from a Proxy Volume




Figure: Object Configuration When Restoring from a Proxy Volume


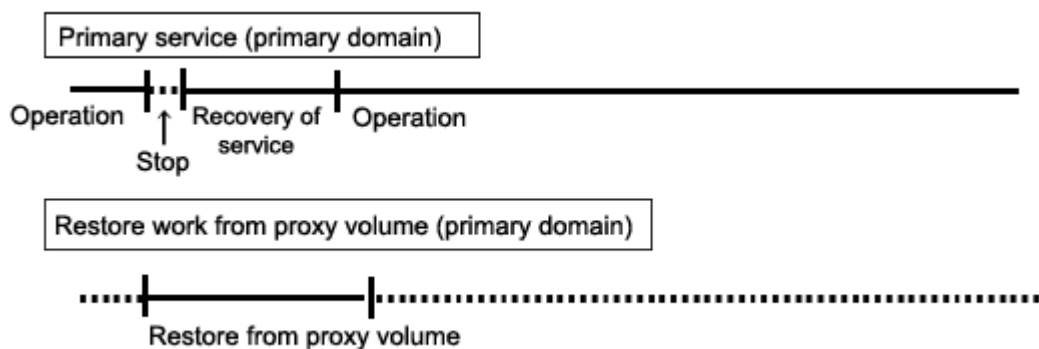


Figure: Schedule for Restore a Proxy Volume

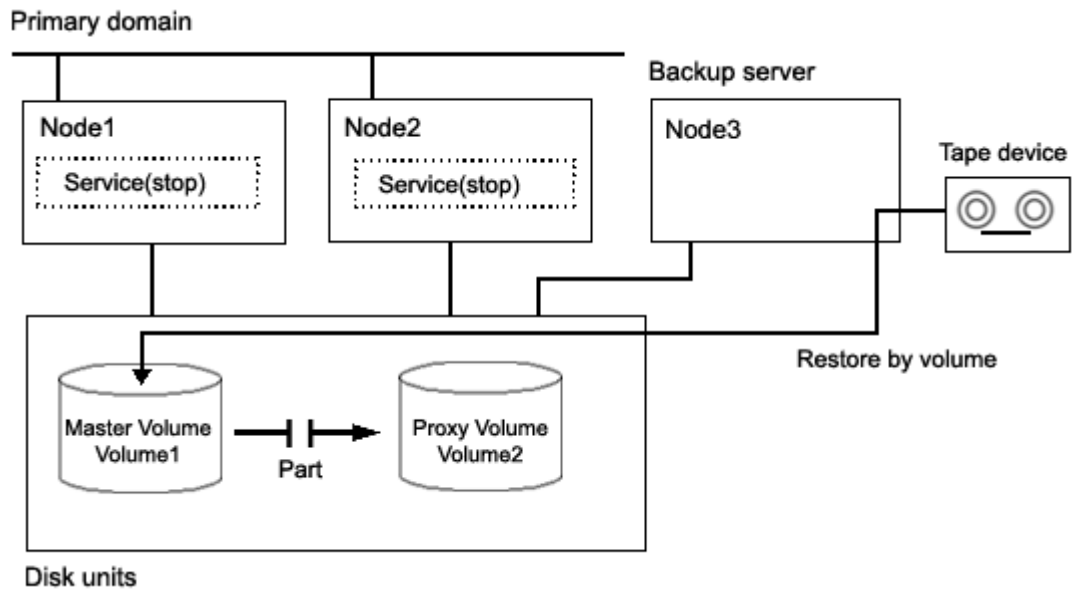


6.6.3.4 Summary of Restore from Tape

If master volume data is damaged while the master volume for the services operation and the proxy volume are in the joined state, the proxy data is also damaged. In this case data can be restored from tape back to the master volume.

Data can be restored while the service is stopped and the master volume is not in use.

Figure: 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: Object Configuration When Restoring from Tape


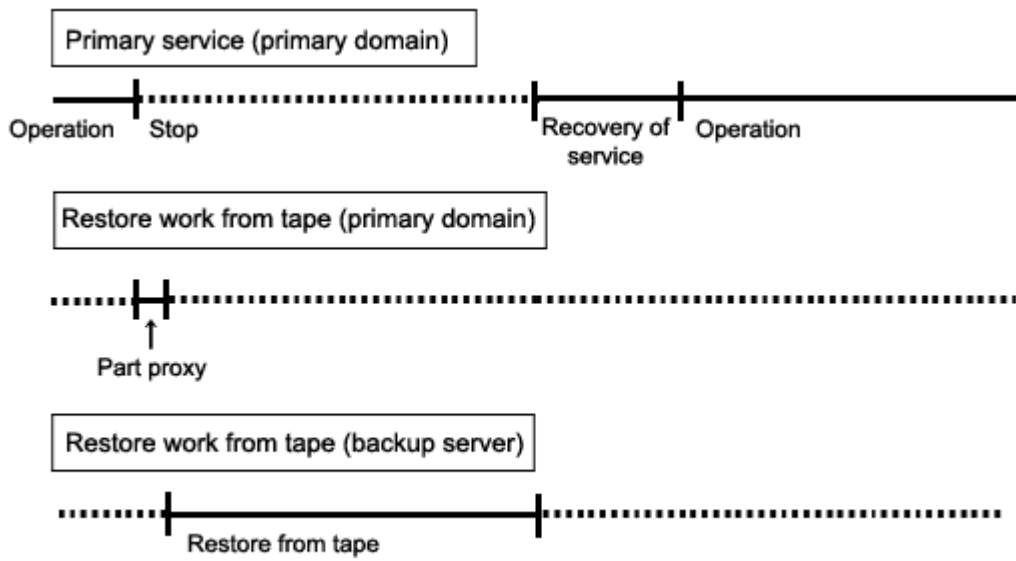


Figure: Schedule for Restore from Tape



6.6.3.5 Summary of Procedure

Figure: Outline of the Configuration Procedure




Figure: Outline of the Backup Procedure




Figure: Outline of the Procedure for Restoring from a Proxy Volume


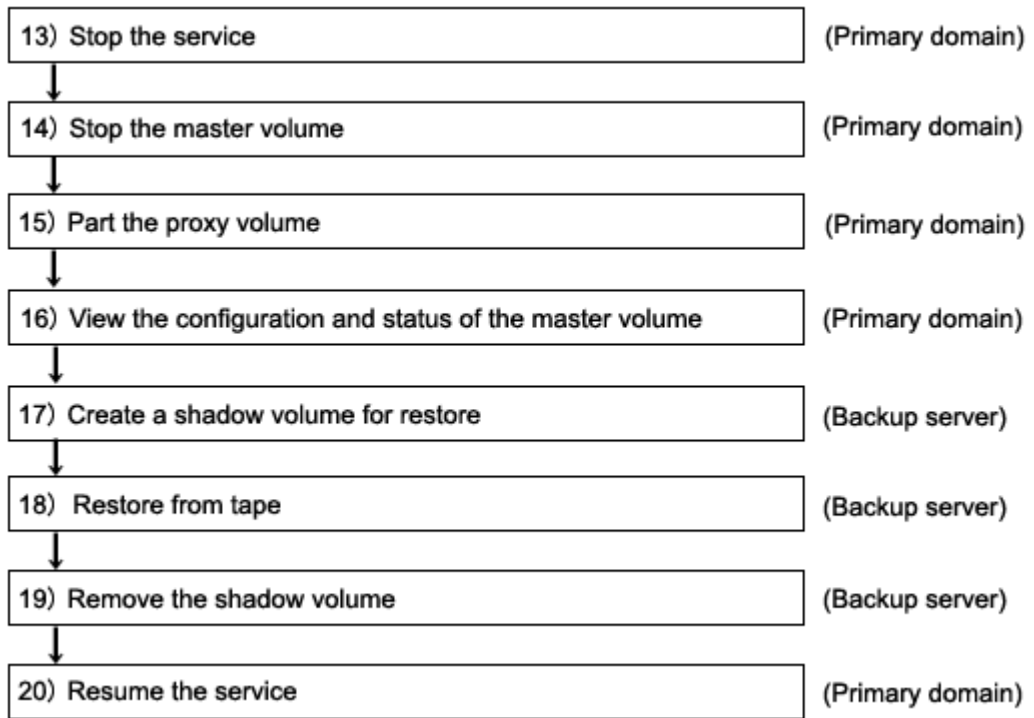


Figure: Outline of the Procedure for Restoring from Tape



6.6.3.6 Configuring an 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 “Appendix H Shared Disk Unit 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.

6.6.3.7 Backing Up

3) Parting the proxy volume

Part the proxy volume from the master volume. The following procedure must be performed on Node1 or Node2 in the primary domain.



Information

The following example secures data consistency by stopping the services when the proxy volume is parted. Steps 3-2) and 3-4) are not required if software, such as a file system and a database system, that manages volume data provides functionality ensuring data consistency or repairing consistency for a parted volume is present. Alternatively, data consistency must be secured with the method specific to that software. For details, see “A.2.20 Ensuring Consistency of Snapshot Data.”

3-1) Viewing the status of the proxy volume

Confirm that master volume Volume1 and proxy volume Volume2 are in sync with each other.

Confirm that proxy volume Volume2 is in the joined state. If Join is displayed in the PROXY field, the proxy volume is in the joined state.

```
# sdxinfo -V -c Class1 -o Volume2 -e long
```

OBJ	NAME	TYPE	CLASS	GROUP	DISK	MASTER	PROXY	...
volume	*	mirror	Class1	Group2	*	*	*	...
volume	Volume2	mirror	Class1	Group2	*	Volume1	Join	...
volume	*	mirror	Class1	Group2	*	*	*	...

Confirm that data in all the slices of proxy volume Volume2 is valid (STOP).

```
# sdxinfo -S -c Class1 -o Volume2
```

OBJ	CLASS	GROUP	DISK	VOLUME	STATUS
slice	Class1	Group2	Disk3	Volume2	STOP

If data is not valid (STOP), repair the slice status referring to “F.1.1 Slice Status Abnormality.”

3-2) Stopping the services

To secure consistency of data in proxy volume Volume2 parted, exit all applications using master volume Volume1 on Node1 and Node2.

When Volume1 is used as a file system, it should be unmounted.

3-3) Parting the proxy volume

Part proxy volume Volume2 from master volume Volume1.

```
# sdxproxy Part -c Class1 -p Volume2 -a pjrm=off
```



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 ext3 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 ext3 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 ext3 -y /dev/sfdsk/Class2/dsk/Volume2
```

6b-3) Mount the ext3 file system on shadow volume Volume2 on /mnt1, a temporary mount point, in the read only mode.

```
# mkdir /mnt1
# mount -t ext3 -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
```


6.6.3.8 Restoring from a Proxy Volume

9) Stopping the services

Exit all applications using master volume Volume1 on Node1 and Node2 in the primary domain.

When Volume1 is used as a file system, it should be unmounted.

10) Stopping the master volume

Stop master volume Volume1 on Node1 and Node2 in the primary domain. Execute the following command on Node1 or Node2.

```
# sdxvolume -F -c Class1 -v Volume1 -e allnodes
```

11) Restoring from the proxy volume

In the primary domain, restore data from proxy volume Volume2 back to master volume Volume1. Execute the following command on Node1 or Node2.

```
# sdxproxy RejoinRestore -c Class1 -p Volume2
```

After returning from the command, synchronization copying from Volume2 to Volume1 is executed.

12) Resuming the services

After synchronization copying is started from proxy volume Volume2 to master volume Volume1 in step 11), the services can be resumed before the copy process is completed. The following procedure must be performed on the node that runs the services.

12-1) Activating the master volume

Activate master volume Volume1.

```
# sdxvolume -N -c Class1 -v Volume1
```

12-2) 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 "6.6.3.7 Backing Up."

6.6.3.9 Restoring from Tape

This sub-section shows an example that access can be gained from backup server Node3 to all of the disks constituting master volume Volume1.



Information When Using a Disk Unit's Copy Function

When access can be gained from backup server Node3 to all of the disks constituting master volume Volume1 and the disk unit's copy function is used to synchronize a master and a proxy, parting the proxy volume in step 15) is not required.



Information When access cannot be gained from the backup server to the disks constituting the master volume

While proxy volume Volume2 is parted, copy data from tape to the proxy volume, and then restore master volume data using the proxy volume.

13) Stopping the services

Exit all applications using master volume Volume1 on Node1 and Node2 in the primary domain.

When Volume1 is used as a file system, it should be unmounted.

14) Stopping the master volume

On Node1 and Node2 in the primary domain, stop master volume Volume1 to prevent improper access to it. Execute the following command on Node1 or Node2.

```
# sdxvolume -F -c Class1 -v Volume1 -e allnodes
```

15) Parting the proxy volume

In the primary domain, part proxy volume Volume2 from master volume Volume1. Execute the following command on Node1 or Node2.

```
# sdxproxy Part -c Class1 -p Volume2 -a pjrm=off
```



Note Just Resynchronization Mode for Proxy

After data held in master volume Volume1 is restored from tape on backup server Node3, proxy volume Volume2 is supposed to be rejoined to master volume Volume1. At this point the entire Volume1 data must be copied to Volume2. For this reason, the JRM mode of a parted proxy volume must be set to off in advance.



Information When Using a Disk Unit's Copy Function

When the disk unit's copy function is used to synchronize a master and a proxy, this procedure (parting a proxy) is not required.



Information When access cannot be gained from the backup server to the disks constituting the master volume

After proxy volume Volume2 is parted, inactivate Volume2 on Node1 and Node2 to prevent Volume2 from being written in improperly.

16) Viewing the status and configuration of the master volume

On Node1 and Node2 in the primary domain, view the configuration and status of master volume Volume1 that is the restore target. Confirm that all of the slices constituting Volume1 are in STOP status. If the status of a slice is invalid, repair it referring to "F.1.1 Slice Status Abnormality."

```
# sdxinfo -S -c Class1 -o Volume1
OBJ      CLASS  GROUP  DISK    VOLUME  STATUS
-----  -----  -----  -----  -----  -----
slice    Class1  Group1  Disk1   Volume1  STOP
slice    Class1  Group1  Disk2   Volume1  STOP
```



Information When Using a Disk Unit's Copy Function

When the proxy volume was not parted in step 15) because the disk unit's copy function is used for synchronizing a master and a proxy, you must also confirm that all of the slices constituting proxy volume Volume2 are in STOP status.



Information When access cannot be gained from the backup server to the disks constituting the master volume

View the configuration and the status of proxy volume Volume2 that is the restore target.

17) Creating a shadow volume for restore

On backup server Node3, create a volume for restore (shadow volume).

17-1) Creating a shadow volume

```
# sdxshadowdisk -M -c Class2 -d 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 restore in a similar procedure that created a proxy volume in step 2).

17-2) Setting the access mode of the shadow volume

Activate shadow volume Volume1 in the read and write access mode (rw).

```
# sdxshadowvolume -F -c Class2 -v Volume1
# sdxshadowvolume -N -c Class2 -v Volume1 -e mode=rw
```

17-3) Viewing the configuration of the shadow volume

Using the sdxinfo command, confirm that the group configuration and the volume configuration are correct based on group information in the DISKS field, volume information in the 1STBLK field and in the BLOCKS field and so on.

```
# sdxinfo -c Class2
```

18) Restoring from tape

On backup server Node3, restore shadow volume data from tape to which it was backed up in step 6). In the following examples, restore data held in shadow volume Volume1 from a tape medium of tape device /dev/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 ext3 file system with the tar(1) command

18b-1) Create the ext3 file system to shadow volume Volume1.

```
# mkfs -t ext3 /dev/sfdsk/Class2/dsk/Volume1
```

18b-2) Mount the ext3 file system on shadow volume Volume1 on /mnt1, a temporary mount point.

```
# mkdir /mnt1
# mount -t ext3 /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 "6.6.3.8 Restoring from a Proxy Volume."

20-1) Activating the master volume

Activate master volume Volume1.

```
# sdxvolume -N -c Class1 -v Volume1
```

20-2) Resuming the services

When the file system on Volume1 was unmounted in step 13), mount it again. Start the applications using Volume1.

6.6.4 Backing Up and Restoring by the Disk Unit's Copy Function

This sub-section describes the method of backing up data from and restoring data back to mirror volumes in the primary domain through a backup server in another domain using the disk unit's copy function provided by a disk array unit.

If volume data is updated with copy functions of disk units, GDS does not recognize the update. If mirrored disk data is updated with copy functions of disk units, synchronization of the mirrors are no longer ensured. Therefore, when restoring mirror volume data using the disk unit's copy function, it is necessary to disconnect the other disk from mirroring once and reconnect it to mirroring after restoration.

The following example illustrates using an EMC's Symmetrix storage unit as a disk array unit and EMC TimeFinder as a copy function.

When restoration is performed with TimeFinder, configuration information within the private slice is also restored. For this reason, simultaneous backup and restore for all disks within the class is required. Additionally, the object configuration and status must match at backup and at restore, and so as in restore, it is necessary to disconnect the disks from mirroring before backup and reconnect after.

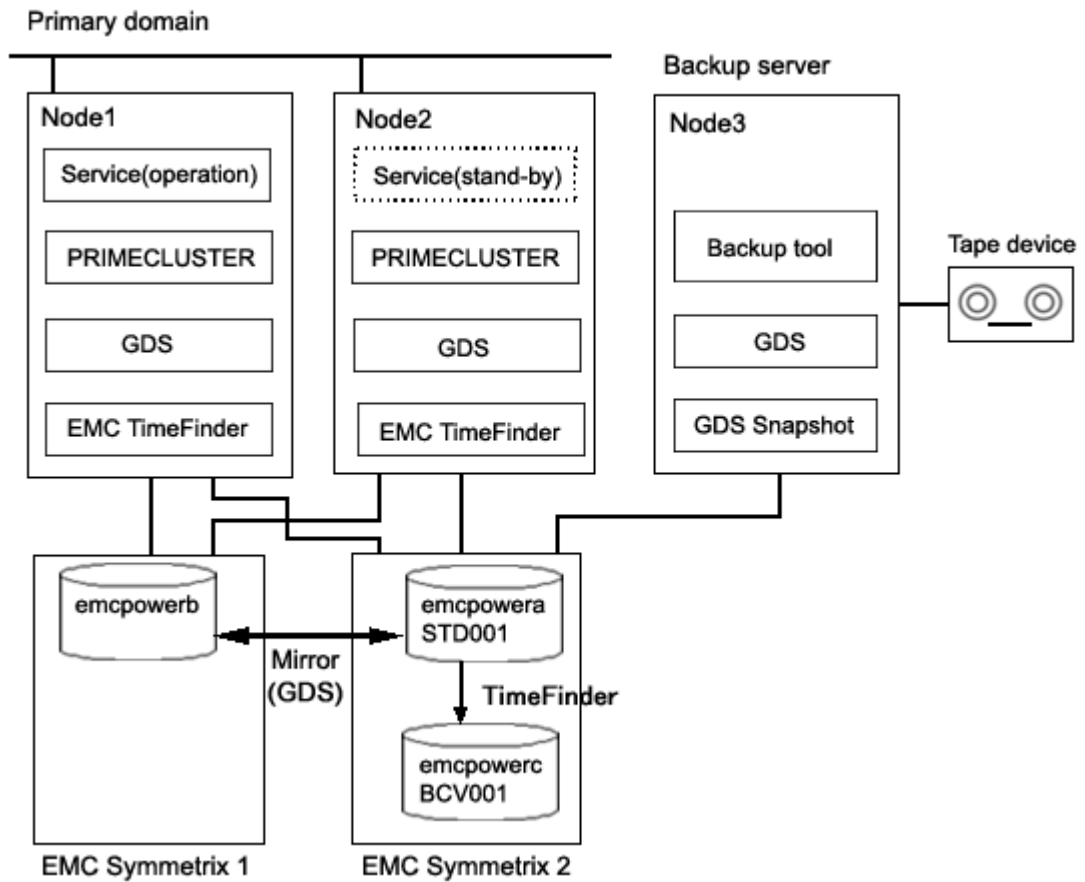


See

- For using EMC's Symmetrix storage units, refer to notes described in "A.2.18 To Use EMC Symmetrix."

6.6.4.1 System Configuration

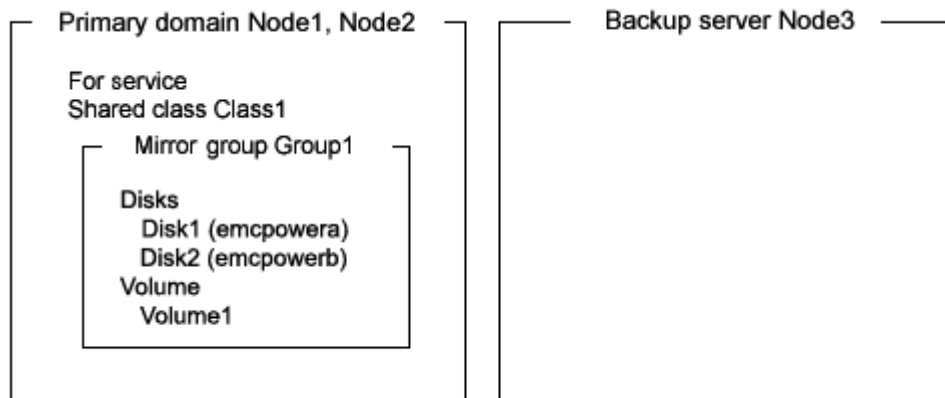
Figure: System Configuration



Information

- A configuration that makes a node in the primary domain (e.g. Node2) work as a backup server is also available.

Figure: Object Configuration in Normal Operation



6.6.4.2 Summary of Backup

Data in a split BCV can be backed up to tape during the services. To secure consistency of BCV data, the services should be stopped temporarily when the BCV is split.



Information Consistency of Snapshot Data

When detaching BCV while the services are operating, data consistency must be secured through the method specific to that software, such as a file system and a database system, which manages volume data. For details, see "A.2.20 Ensuring Consistency of Snapshot Data."

Figure: Backup

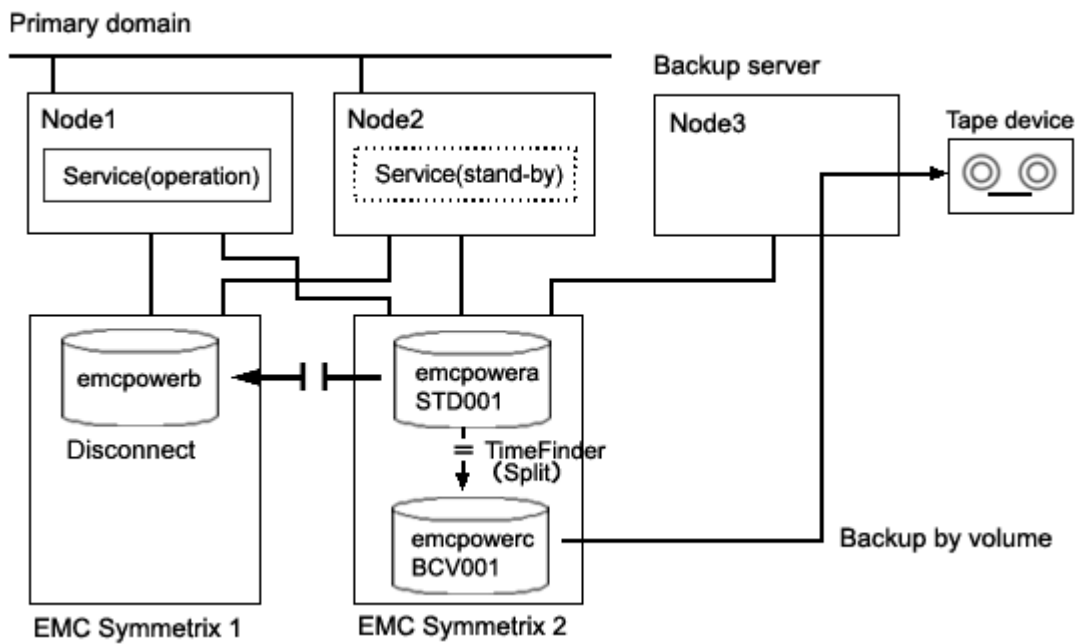


Figure: Object Configuration When Backing Up

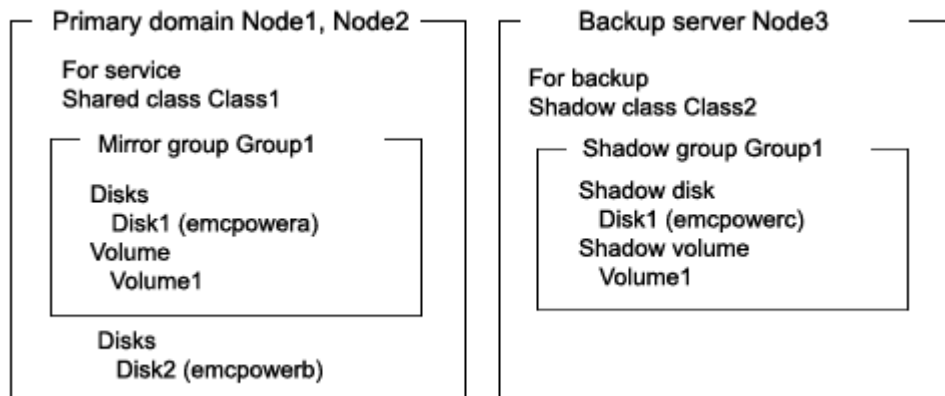
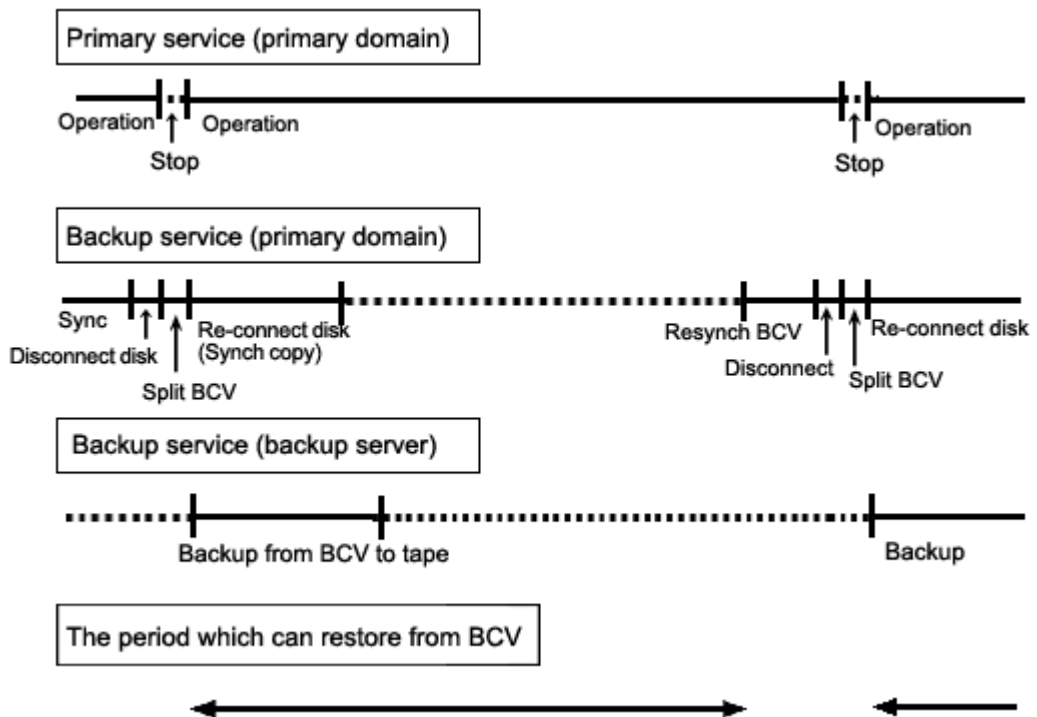


Figure: Backup Schedule



Note When Restoring from the BCV

The following conditions must be met when performing a backup to the BCV (disconnecting the BCV).

- Back up all disks registered with the backup target class (except for disks that are disconnected and in SWAP status) to the BCV. For disks not backed up to the BCV, it is necessary to disconnect from the class before the BCV is detached.
- Before backup to the BCV is completed for all disks within the class (except for disks that are disconnected and in SWAP status), do not change the object configuration or status in the class.



Information When Not Restoring from the BCV

When restore is always from tape instead of BCV, it is not necessary to perform disk disconnection and reconnection in the backup process.

6.6.4.3 Summary of Restore from a BCV

If data in a standard device is damaged while a BCV is split from the disk (standard disk) used for the services operation, the standard device data can be restored from the BCV.

Data can be restored while the service is stopped and the application volume is not in use.



Note Conditions for Restore from the BCV

The following conditions must be met when performing a restore from the BCV.

- To all disks within the restore target class (except for disks that are disconnected and in SWAP status), restore data from the BCV. For disks not backed up to the BCV, it is necessary to disconnect from the class before restore from the BCV is performed.
- The configuration of objects within the restore target class must be the same as it was when the backup to the BCV (disconnection of the BCV) was performed.
- Before restore to the BCV is completed for all disks within the class (except for disks that are disconnected and in SWAP status), do not reboot any node in the primary domain.

Figure: Restore from a BCV

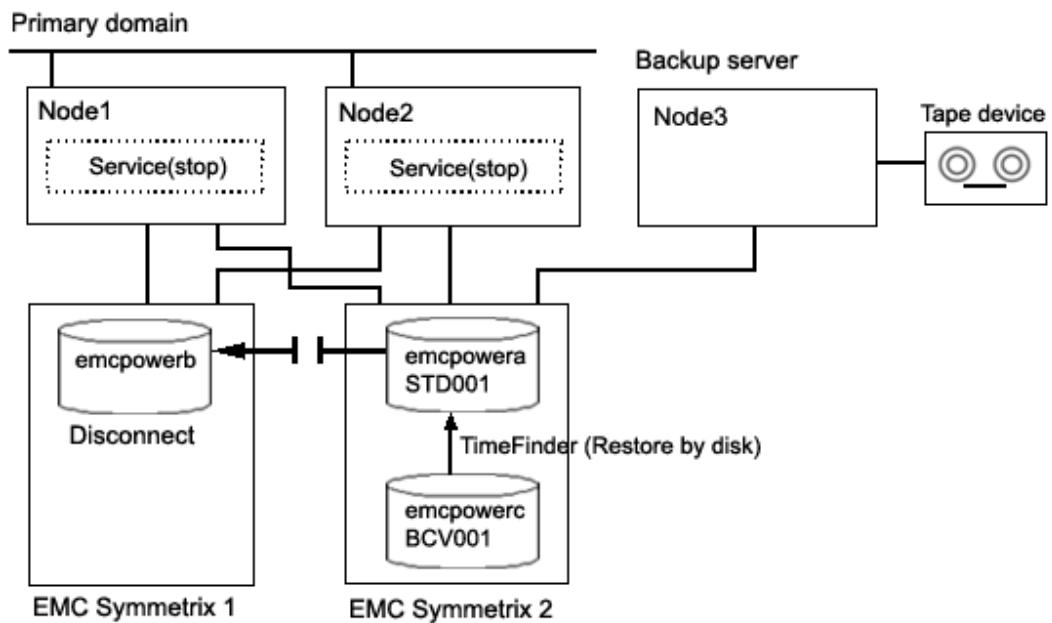


Figure: Object Configuration When Backing Up from a BCV


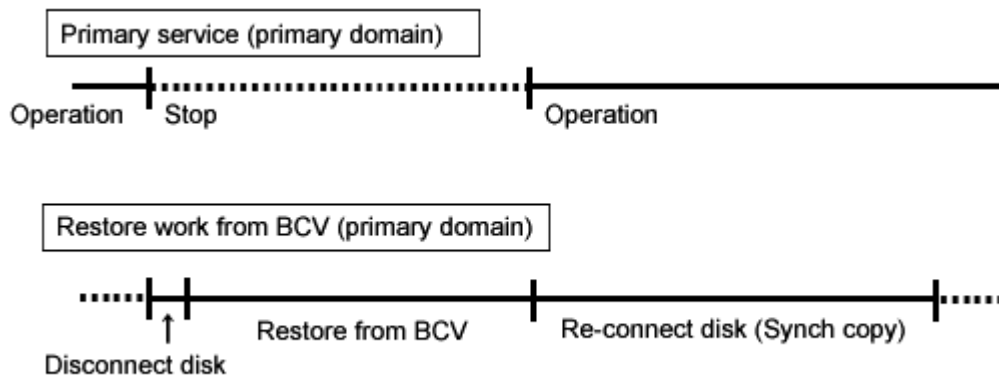


Figure: Schedule for Restore from a BCV




6.6.4.4 Summary of Restore from Tape

If data in a standard device used for the service is damaged while it is in sync with a BCV, data in the BCV is also damaged. In this case data can be restored from tape back to the standard device.

Data can be restored while the services are stopped and the application volume is not in use.

Figure: 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: Object Configuration When Restoring from Tape

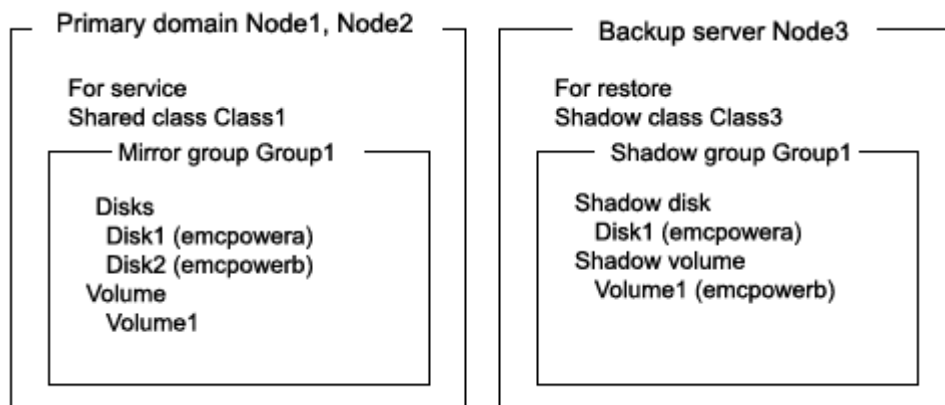
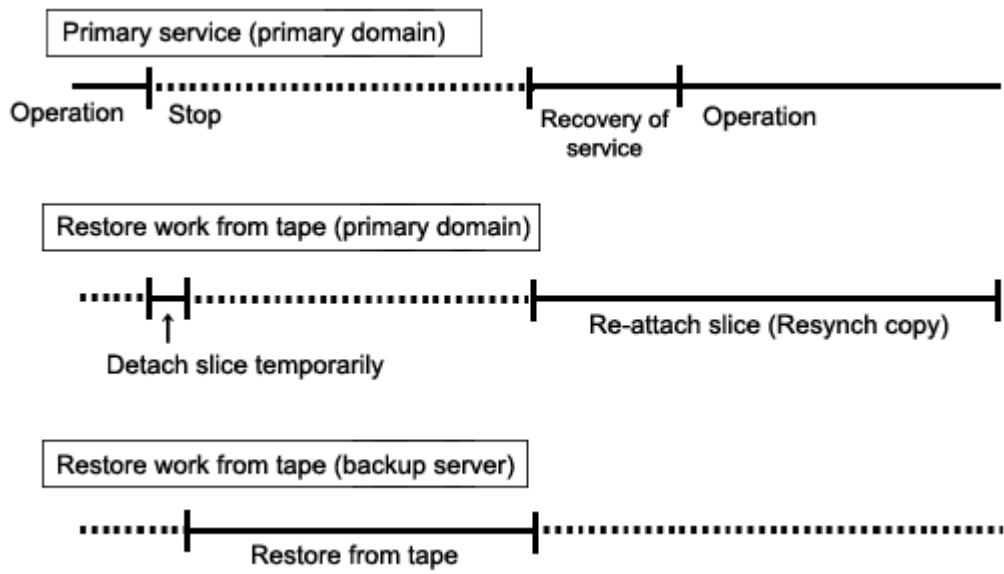


Figure: Schedule for Restore from Tape



6.6.4.5 Summary of Procedure

Figure: Outline of the Configuration Procedure




Figure: Outline of the Backup Procedure

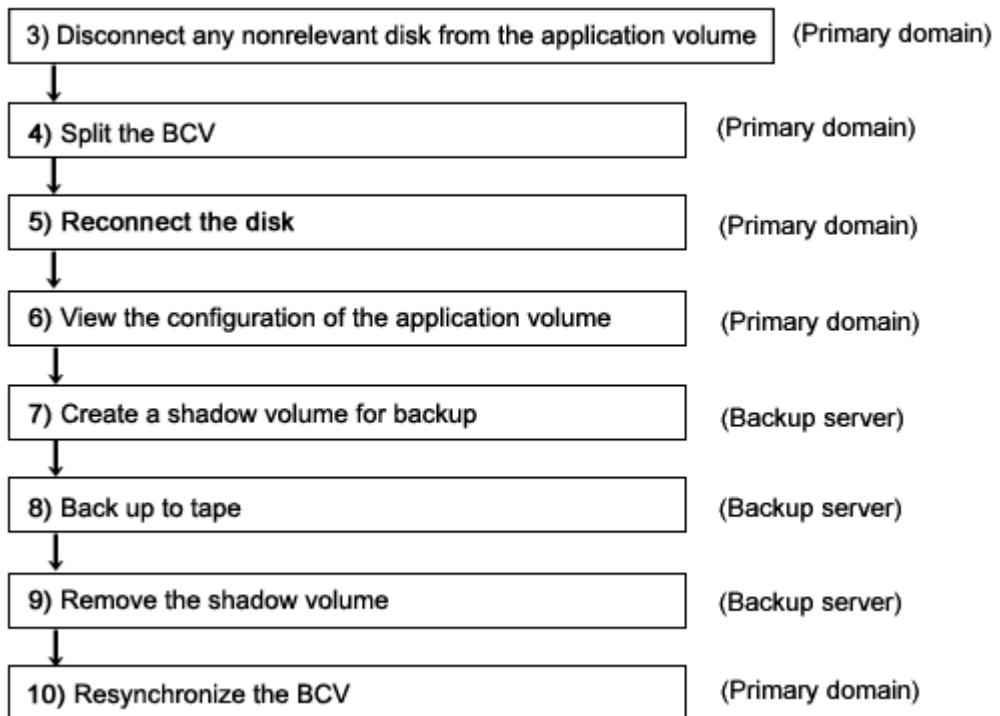


Figure: Outline of the Procedure for Restoring from a BCV


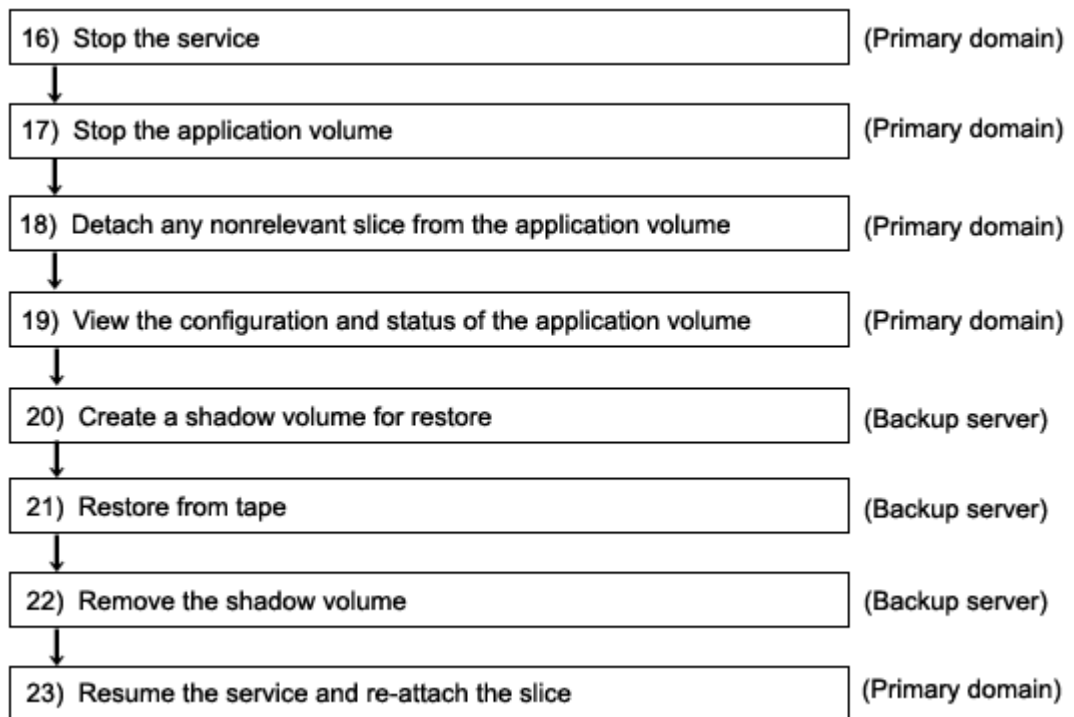


Figure: Outline of the Procedure for Restoring from Tape



6.6.4.6 Configuring an 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 “Appendix H Shared Disk Unit 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
```


6.6.4.7 Backing Up

3) Disconnecting a disk of the application volume

In the primary domain, among disks registered with class Class1 to which application volume Volume1 belongs, disconnect a disk (Disk2) other than the backup target disk Disk1 from Class1. Execute the following command on node Node1 or Node2 in the primary domain.

```
# sdxswap -O -c Class1 -d Disk2
```

4) Splitting the BCV

Split BCV device 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 "A.2.20 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 "F.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 -l -c Class1 -d Disk2 -e nowaitsync
```

After returning from the command, synchronization copying from the slice on Disk1 of volume Volume1 to the slice on Disk2 is executed.

6) Viewing the configuration of the application volume

On Node1 or Node2 in the primary domain, view the configuration of services volume Volume1 that is the backup target.

```
# sdxinfo -c Class1
OBJ  NAME  TYPE  SCOPE  SPARE
----  ----  ----  -----  ----
class Class1  shared  Node1:Node2  0

OBJ  NAME  TYPE  CLASS  GROUP  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  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.



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 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 emcpowera 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.

```
# 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 ext3 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 ext3 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 ext3 /dev/sfdsk/Class2/dsk/Volume1
```

8b-3) Mount the ext3 file system on shadow volume Volume1 on /mnt1, a temporary mount point, in the read only mode.

```
# mkdir /mnt1  
# mount -t ext3 -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 Id BCV001
```

To back up again, follow the procedure from step 4).

6.6.4.8 Restoring form a BCV

11) Stopping the services

Exit all applications using application volume Volume1 on Node1 and Node2 in the primary domain. When Volume1 is used as a file system, unmount it.

12) Stopping the application volume

Stop application volume Volume1 on Node1 and Node2 in the primary domain. Execute the following command on Node1 or Node2.

```
# sdxvolume -F -c Class1 -v Volume1 -e allnodes
```

13) Disconnecting any nonrelevant disk from the application volume

In the primary domain, disconnect any disk (Disk2) other than Disk1 that is the restore target from Class1, among disks connected to Class1 to which application volume Volume1 belongs. Execute the following command on Node1 or Node2 in the primary domain.

```
# sdxswap -O -c Class1 -d Disk2
```

14) Restoring from the BCV

Restore data held in standard device STD001 from BCV device BCV001 in the primary domain. The following procedure must be performed on Node1 or Node2 in the primary domain.

14-1) Restoring from the BCV

Restore data held in standard device STD001 from BCV device BCV001.

```
# symmir -g DevGroup restore STD001 BCV Id 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
-----			-----			-----
Logical	Inv. Sym	Tracks	Logical	Sym	Inv. 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 -l -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
```

6.6.4.9 Restoring from Tape

16) Stopping the services

Exit all applications using application volume Volume1 on Node1 and Node2 in the primary domain. When Volume1 is used as a file system, unmount it.

17) Stopping the services application

To write-lock volume Volume1, inactivate Volume1 on Node1 and Node2 in the primary domain. Execute the following command on Node1 or Node2.

```
# sdxvolume -F -c Class1 -v Volume1 -e allnodes
```

18) Detaching any nonrelevant slice from the application volume

Temporarily detach the slice on any disk (Disk2) other than Disk1 that is the restore target from Volume1, among slices in application volume Volume1. Execute the following command on Node1 or Node2 in the primary domain.

```
# sdxslice -M -c Class1 -d Disk2 -v Volume1 -a jrm=off
```



Point

On backup server Node3, after Disk1 data is restored from tape, the slice on Disk2 is supposed to be reattached to application volume Volume1 in the primary domain. At this point the entire volume data must be copied to the attached slice. For this reason, the JRM mode of a detached slice must be set to off in advance.



Information

If access can be gained from backup server Node3 to Disk2, data can be restored from tape back to both Disk1 and Disk2 on Node3. Under these circumstances, this procedure (detaching a slice) should not be performed.

19) Viewing the configuration and status of the application volume

On Node1 and Node2 in the primary domain, see the configuration and status of application volume Volume1 that is restore target. Confirm that Volume1 is in STOP status and that only restore target slice Volume1.Disk1 is in STOP status among the slices constituting the volume and the other slices are in TEMP or TEMP-STOP status. If the volume or slice status is invalid, repair it referring to "F.1.3 Volume Status Abnormality" or "F.1.1 Slice Status Abnormality."

```
# sdxinfo -c Class1
```

OBJ	NAME	TYPE	SCOPE	SPARE					
class	Class1	shared	Node1:Node2	0					
OBJ	NAME	TYPE	CLASS	GROUP	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 restore

On backup server Node3, create a volume for restore (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  emcpwera  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 ext3 file system with the tar(1) command

21b-1) Create the ext3 file system to shadow volume Volume1.

```
# mkfs -t ext3 /dev/sfdsk/Class3/dsk/Volume1
```

21b-2) Mount the ext3 file system on shadow volume Volume1 on /mnt2, a temporary mount point.

```
# mkdir /mnt2
# mount -t ext3 /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

6.7 Backing Up and Restoring Object Configurations [Linux2.6]

Due to multiple disk failures, the valid configuration database of a class may be lost, resulting in loss of configuration information of objects within the class. If that happens, after recovering the failed disks, objects such as volumes must be re-created. By backing up object configuration information in advance, efforts to restore object configurations can be reduced.

This section describes the procedures for backing up and restoring object configurations of classes.



Note **Classes Unavailable for Object Configuration Backup**

Object configurations of the following classes cannot be backed up.

- Root class
- Shared class that include a switch group
- Class that include a proxy object
- Shadow class



Note **Systems Available for Object Configuration Restore**

To restore object configuration according to backed up object configuration information, the system for restoration must be connecting disks that are equivalent in size to physical disks registered with the backed up class.

6.7.1 Backing Up

This subsection describes the procedures for backing up configuration information of objects within class Class1.

1) Saving configuration information

Save outputs of the `sdxinfo` command to a file. In this example, the path to a file is `"/var/tmp/Class1.info"`.

```
# sdxinfo -c Class1 -e long > /var/tmp/Class1.info
```

2) Creating a configuration file

Output the object configuration within Class1 to a file in configuration table format. In this example, the path to a file is `"/var/tmp/Class1.conf"`.

```
# sdxconfig Backup -c Class1 -o /var/tmp/Class1.conf
```

3) Backing up the configuration file and configuration information

Save the files created in steps 1) and 2) to tape and so on.

6.7.2 Restoring

This subsection describes the procedures for restoring the object configuration within class Class1 according to the configuration file and configuration information saved in advance as shown in "6.8.1 Backing Up" in the event of loss of the Class1 object configuration caused by a problem of some kind.

1) Checking the class scope

With a cluster system, check names of nodes sharing the class. For node names that belong to the class scope, check the SCOPE field of class information output by the sdxinfo command and saved as shown in step 1) of "6.8.1 Backing Up."

2) Placing the configuration file

On a node where the object configuration of the class is restored (with a cluster system, on a node that belongs to the class scope), place the configuration file created in step 2) of "6.8.1 Backing Up" In this example, the path to a destination file is "/var/tmp/Class1.conf".

3) Restoring the object configuration of the class

Execute the following command on the node where the configuration file was placed in step 2) to restore the object configuration of class Class1 according to descriptions in the configuration file "/var/tmp/Class1.conf". Class1 is restored as a local class of that node.

After restoring the object configuration, reboot the node.

```
# sdxconfig Restore -c Class1 -i /var/tmp/Class1.conf
# shutdown -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
```

4) Changing the class type and expanding the class scope

If the backed up class, Class1, is a shared class, change the type and scope attributes of Class1. In this example, the scope of the backed up class is node1:node2.

4-1) Stop the volume in the class.

```
# sdxvolume -F -c Class1
```

4-2) Change the class type and expand the class scope.

```
# sdxattr -C -c Class1 -a type=shared,scope=node1:node2
```

