

PRIMECLUSTER™

**Global Disk Services
Configuration and Administration Guide 4.1
(Solaris™ Operating Environment)**

February 2003 Second Edition

Contents

Chapter 1	Function	1
1.1	GDS Features.....	2
1.2	Functions for High Availability.....	4
1.2.1	Disk Mirroring.....	4
1.2.2	Hot Spare	8
1.2.3	Hot Swap	9
1.2.4	Just Resynchronization Mechanism (JRM).....	10
1.3	Functions for High Manageability.....	11
1.3.1	Operation Management Interface	11
1.3.2	Centralized Disk Management.....	12
1.3.3	Name Management	13
1.3.4	Single System Image Environment.....	14
1.3.5	Access Control.....	15
1.3.6	Realizing Large Capacity and I/O Load Balancing.....	18
1.3.7	Snapshot by Detaching Slice	21
1.3.8	Proxy volume.....	23
Chapter 2	Objects	35
2.1	SDX Object.....	36
2.2	Disk Class.....	37
2.3	SDX Disk	39
2.4	Disk Group.....	41
2.5	Logical Volume	43
2.6	Logical Slice	46
2.7	Proxy Object	49
Chapter 3	Starting and Exiting GDS Management View	55
3.1	Preparation for Starting GDS Management View	56
3.1.1	Deciding the User Group	56
3.1.2	Setting up the Client Environment	58
3.1.3	Setting up the Web Environment	59
3.2	Starting the GDS Management View	60
3.2.1	Starting Web-Based Admin View Operation Menu	60
3.2.2	Web-Based Admin View Operation Menu Functions	61
3.2.3	Starting GDS Management View	62
3.3	Exiting GDS Management View.....	63
3.4	Changing the Web-Based Admin View Settings	64
Chapter 4	Management View Screen Elements	65
4.1	Screen Configuration	66
4.2	Menu Configuration and Functions.....	69
4.2.1	General.....	69
4.2.2	Settings.....	70
4.2.3	Operation.....	77
4.2.4	View.....	79
4.2.5	Help	80
4.3	Icon Types and Object Status.....	81
Chapter 5	Operation	85
5.1	Operation Outline.....	86
5.1.1	System Disk Settings.....	86
5.1.2	Automatic Configuration	87
5.1.3	Configuration Settings	88
5.1.4	Backup.....	90
5.1.5	In Case of Disk Abnormality	91
5.1.6	Configuration Changes	92
5.1.7	Unmirroring the System Disk	93
5.1.8	Operation using GDS Management View	94
5.2	Settings.....	96
5.2.1	System Disk Settings.....	96
5.2.2	Automatic Configuration Wizard	103
5.2.3	Operating from the Settings Menu	109
5.2.4	File System Configuration.....	121

5.3	Operation in Use	124
5.3.1	Confirming Configuration and Monitoring Status	124
5.3.2	Backup	129
5.3.3	In Case of Disk Abnormality	136
5.3.4	Copying Operation	140
5.4	Changes	144
5.4.1	Class Configuration	144
5.4.2	Group Configuration	150
5.4.3	Volume Configuration	153
5.5	Removals	154
5.5.1	Removing a File System	154
5.5.2	Removing a Volume	155
5.5.3	Removing a Group	157
5.5.4	Removing a Class	158
5.5.5	Unmirroring the System Disk	159
Appendix A General Notes		161
A.1	Rules	161
A.1.1	Object Name	161
A.1.2	Number of Classes	162
A.1.3	Number of Disks	162
A.1.4	Number of Groups	162
A.1.5	Number of Volumes	163
A.1.6	Number of Spare Disks	163
A.1.7	Number of Keep Disks	163
A.1.8	Creating Group Hierarchy	163
A.1.9	Exception to Proxy Configuration	165
A.1.10	Number of Proxy Volumes	165
A.1.11	Proxy Volume Size	165
A.1.12	Proxy Group Size	165
A.2	Important Points	166
A.2.1	Physical Disk Driver	166
A.2.2	Managing System Disks	166
A.2.3	Restraining Access to Physical Special File	166
A.2.4	Bootng from a CD-ROM Device	167
A.2.5	Initializing Disk	167
A.2.6	Disk Size	167
A.2.7	Volume Size	168
A.2.8	Spare Disk	168
A.2.9	System Disk Mirroring	169
A.2.10	Keep Disk	170
A.2.11	GFS Partition on a Keep Disk	170
A.2.12	Creating a Snapshot by Detaching Slice	170
A.2.13	The Difference between a Mirror Slice and a Proxy Volume	170
A.2.14	Just Resynchronization Mechanism (JRM)	171
A.2.15	Object Operation When Using Proxy	171
A.2.16	Snapshot by Cooperation with Advanced Copy of ETERNUS3000 or ETERNUS GR700 series	172
A.2.17	Instant Snapshot Created by OPC	172
A.2.18	To Use EMC Symmetrix	173
A.2.19	Snapshot by Cooperation with EMC TimeFinder or EMC SRDF	175
A.2.20	Ensuring Integrity of Snapshot Data	176
A.2.21	Data Consistency at the time of Simultaneous Access	176
A.2.22	Volume Access Mode	176
A.2.23	Configuration Operation in Cluster System	177
A.3	General Points	178
A.3.1	Spare Disk Size	178
A.3.2	Guidelines for Mirroring	178
A.3.3	Guidelines for Striping	178
A.3.4	Guidelines for Concatenation	179
A.3.5	Guidelines for Combining Striping with Mirroring	179
A.3.6	Object Status	180
Appendix B Logging View		183

Appendix C Environment Setting	185
Appendix D Command Reference	187
D.1 sdxclass - Class operations	188
D.2 sdxdisk - Disk operations	189
D.3 sdxgroup - Group operations	195
D.4 sdxvolume - Volume operations	198
D.5 sdxslice - Slice operations	202
D.6 sdxinfo - Display object configuration and status information	205
D.7 sdxattr - Set objects attributes	214
D.8 sdxswap - Swap disk	219
D.9 sdxfix - Restore a failed object	221
D.10 sdxcopy - Synchronization copying operation	223
D.11 sdxroot - Define or cancel mirroring of root file system	225
D.12 sdxparam - Configuration parameter operations	228
D.13 sdxproxy – Proxy object operations	230
D.14 Volume Creation Using Command	235
D.15 Snapshot Creation Using Command	239
D.16 Online Backup Making Using Command	241
Appendix E GDS Messages	247
E.1 Web-Based Admin View Messages	247
E.2 Driver Messages	247
E.2.1 Panic Messages	248
E.2.2 Warning Messages	248
E.2.3 Information Messages	252
E.2.4 Internal Error Messages	253
E.3 Daemon Messages	254
E.3.1 Halt Messages	255
E.3.2 Error Messages	256
E.3.3 Warning Messages	266
E.3.4 Information Messages	273
E.3.5 Internal Error Messages	289
E.4 Command Messages	290
E.4.1 Error Messages	291
E.4.2 Warning Messages	328
E.4.3 Information Messages	332
E.4.4 Fix Messages	340
E.4.5 Internal Error Messages	340
E.5 Operation Management View Messages	341
E.5.1 Error Messages	341
E.5.2 Warning Messages	344
E.5.3 Information Messages	351
Appendix F Troubleshooting	357
F.1 Resolving Problems	357
F.1.1 Slice Status Abnormality	357
F.1.2 Disk Status Abnormality	364
F.1.3 Volume Status Abnormality	365
F.1.4 Class Status Abnormality	380
F.1.5 System Disk Abnormality	383
F.1.6 GDS Management View Abnormality	389
F.1.7 EMC Symmetrix Abnormality	392
F.2 Collecting Investigation Material	397
F.2.1 Collecting with fjsnap Command	397
F.2.2 Collecting Initial Investigation Material with sdxsnap.sh Script	398
Appendix G Frequently Asked Questions (FAQ)	399
G.1 Operation Design	399
G.2 Environment Configuration	402
G.3 Operation	404
Glossary	407

Preface

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

Purpose

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

Target Reader

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

Organization

This manual consists of 5 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.

	Contents	Operator	Administrator
Chapter 1 Functions	Explains features of GDS.	YES	YES
Chapter 2 Objects	Explains the objects used by GDS.	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 regarding GDS operation such as operation flow, settings, maintenance and management.	YES	YES
Appendix A General Notes	Explains guidelines, general precautions, and configuration tips necessary for using GDS.	-	YES
Appendix B Logging View	For details, see the supplementary "Web-Based Admin View Operation Guide."	-	YES
Appendix C 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	-	YES
Appendix E GDS Messages	Explains the contents, possible causes, and resolutions for GDS messages that appear when setting or operating GDS	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	YES	YES
Glossary	Explains GDS terminology. Please refer to it as necessary.	YES	YES

Related documentation

Please refer to the following documents according to need.

- PRIMECLUSTER Installation Guide
- PRIMECLUSTER Concepts Guide
- PRIMECLUSTER Cluster Foundation Configuration and Administration
- PRIMECLUSTER RMS Configuration and Administration
- PRIMECLUSTER Web-Based Admin View Operation Guide
- PRIMECLUSTER Global File Services Configuration and Administration Guide





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:

-  **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

- Solaris(TM) Operating Environment is abbreviated as Solaris OE.
- Solaris(TM) 8 Operating Environment is abbreviated as Solaris 8 OE.
- Solaris(TM) 9 Operating Environment is abbreviated as Solaris 9 OE.
- Microsoft(R) Windows NT(R) Server operating System Version4.0 and Microsoft(R) Windows NT(R) Server workstation System Version4.0 are abbreviated as Windows NT(R).
- Microsoft(R) Windows(R) 95 operating system is abbreviated as Windows(R) 95.
- Microsoft(R) Windows(R) 98 operating system and Microsoft(R) Windows(R) 98 Second Edition operating system are abbreviated as Windows(R) 98.
- Microsoft(R) Windows(R) 2000 operating system is abbreviated as Windows(R) 2000.
- Microsoft(R) Windows(R) Millennium Edition is abbreviated as Windows(R) Me.
- Microsoft(R) Windows(R) XP operating system is abbreviated as Windows(R) XP.
- Windows(R) 95, Windows(R) 98, Windows NT(R), Windows(R) 2000, Windows(R) Me and Windows(R) XP are abbreviated as Microsoft(R) Windows.

February 2003 Second Edition

Trademarks

Sun, Sun Microsystems, the Sun Logo, Solaris and all Solaris based marks and logos are trademarks of Sun Microsystems, Inc. in the U.S. and other countries, and are used under license.

Netscape and Netscape Navigator are trademarks of Netscape Communications Corporation, USA.

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

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

EMC, PowerPath and Symmetrix are registered trademarks of EMC Corporation. TimeFinder and SRDF are trademarks 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) 2003 Fujitsu LIMITED.

All Rights Reserved, Copyright(C) 2003 Fujitsu Siemens Computers GmbH.

Chapter 1 Function

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

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 following functions, which are closely related:

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

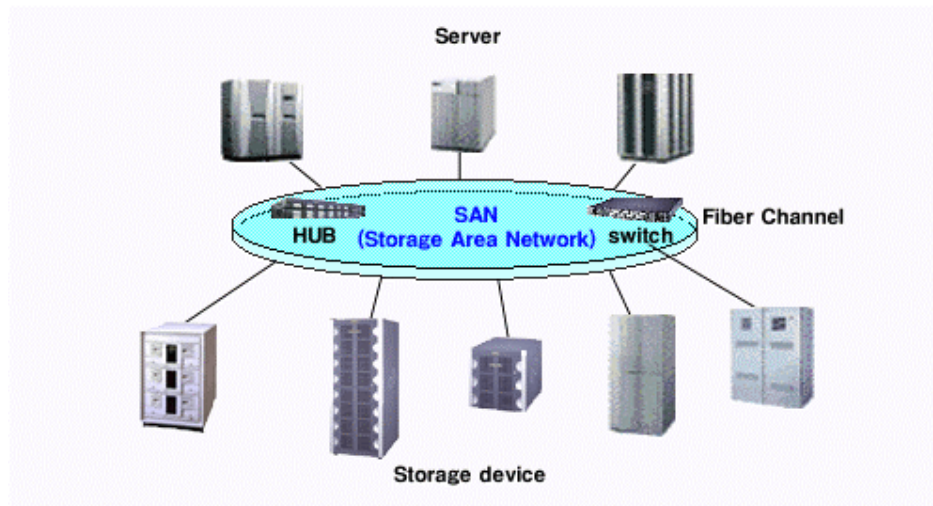
GDS's mirroring function protects data from hardware failures by maintaining replicas of disk data on multiple disks. This allows users to continue to access to the disk data without stopping the application in the event of an unexpected trouble.

Figure: Mirroring of Disk Units



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

Figure: SAN (Storage Area Network)

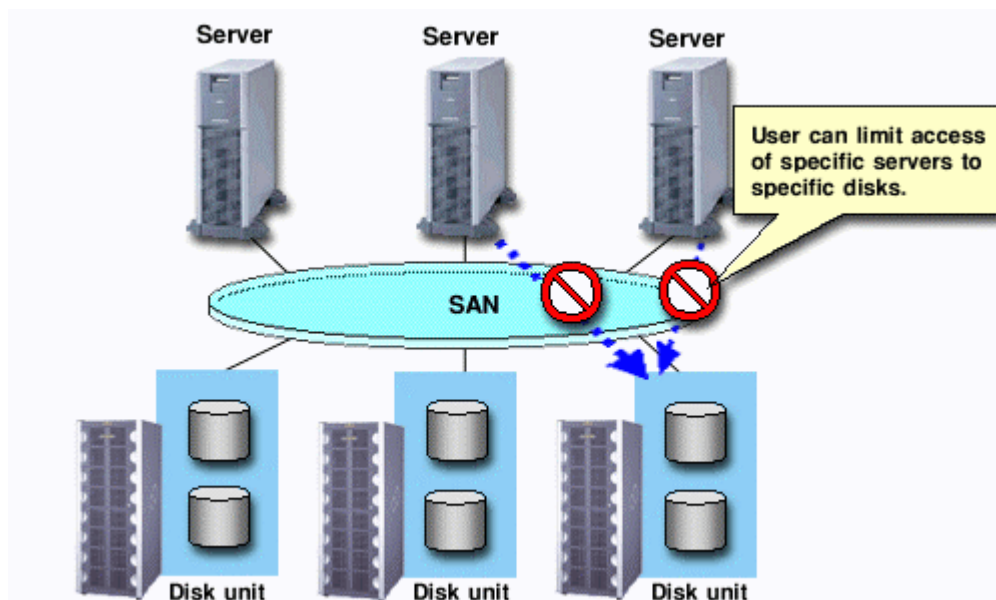


Multiple servers can be connected to multiple disk units in a SAN environment. Disk-stored data can be accessed from those servers. This allows simultaneous access to file systems or databases and improves the efficiency of data duplication between the servers and backup procedures. On the other hand, 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 done 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



GDS provides a snapshot function that allows users to create replications of disk data (snapshots) at a certain moment to another area for backup purpose. Snapshots will allow you to conduct operations, such as non-disruptive backup and data analysis, concurrently with the core business. When using a ETERNUS3000 or GR700 series storage system with Advanced Copy function, snapshots can be created instantly without imposing a load on the server or the SAN. The GDS Snapshot function helps you configure backup administration that will not affect the core business.

1.2 Functions for High Availability

Failure of hardware related to disk unit interrupts data access, causing applications or system to stop. In this section, functions which protect data from unexpected hardware problems, and improve system availability is explained.

1.2.1 Disk Mirroring

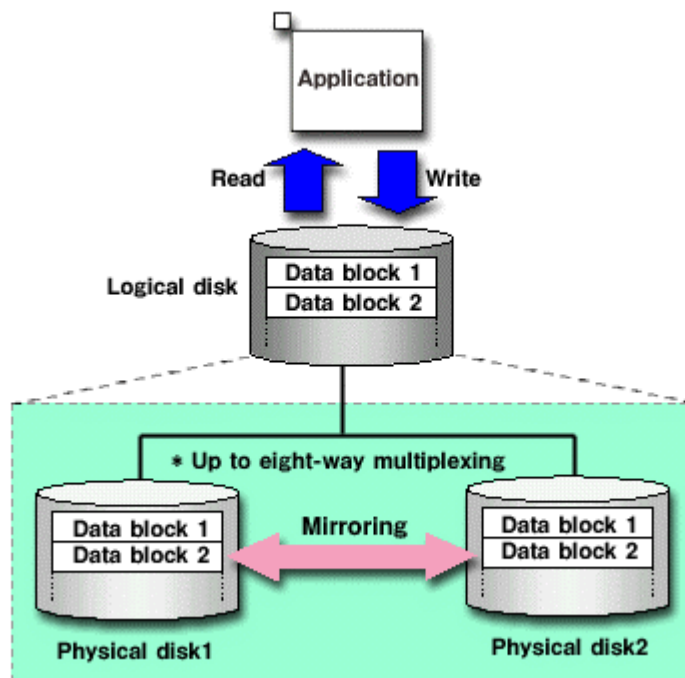
Disk mirroring utilizes multiple disks in order to maintain data synchronization.

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

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.


Figure: Disk Mirroring



1.2.1.1 System Disk Mirroring

System disk failure may cause the entire system to stop. Since you cannot reboot your computer, the operating system must be reconstructed, resulting in a long interruption. Mirroring the system disk ensures that you can continue normal operation, even if a failure occurs with one of the disks.

Figure: System Disk Mirroring

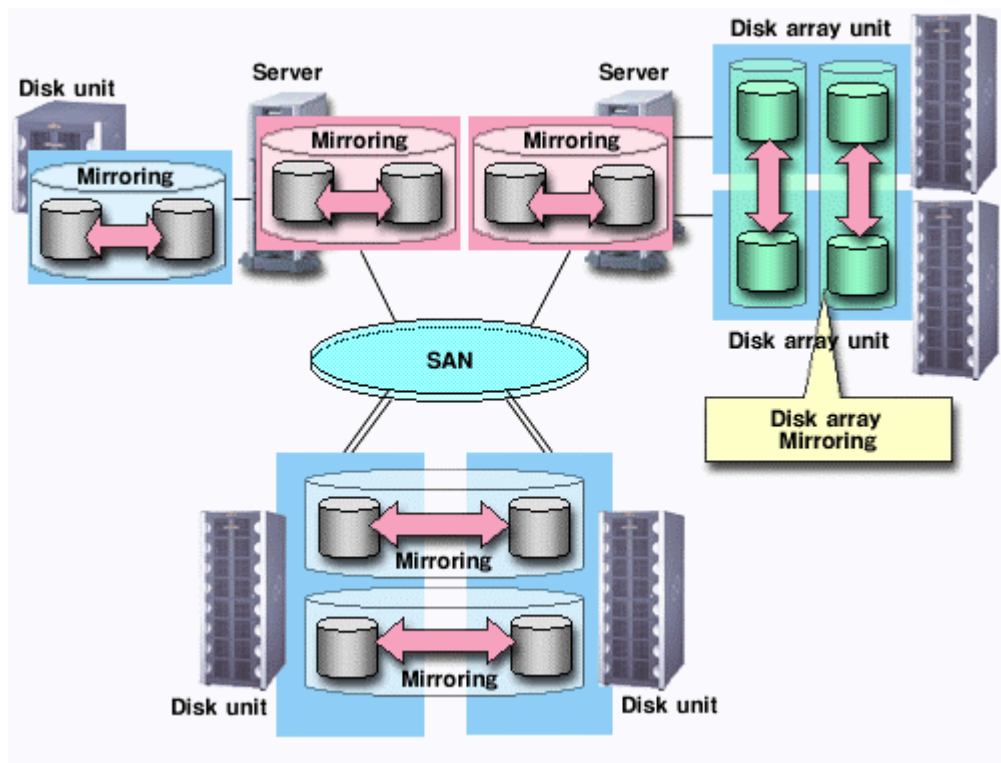


1.2.1.2 Mirroring between Disk Array Unit

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

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

Figure: Mirroring between Disk Array Unit




1.2.1.3 Shared Disk Mirroring

PRIMECLUSTER GDS can operate in conjunction with cluster system to mirror shared disks connected to multiple servers (also referred to as node).

This is called shared disk mirroring, and to differentiate from local disk mirroring which mirrors disks connected to a single node.

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

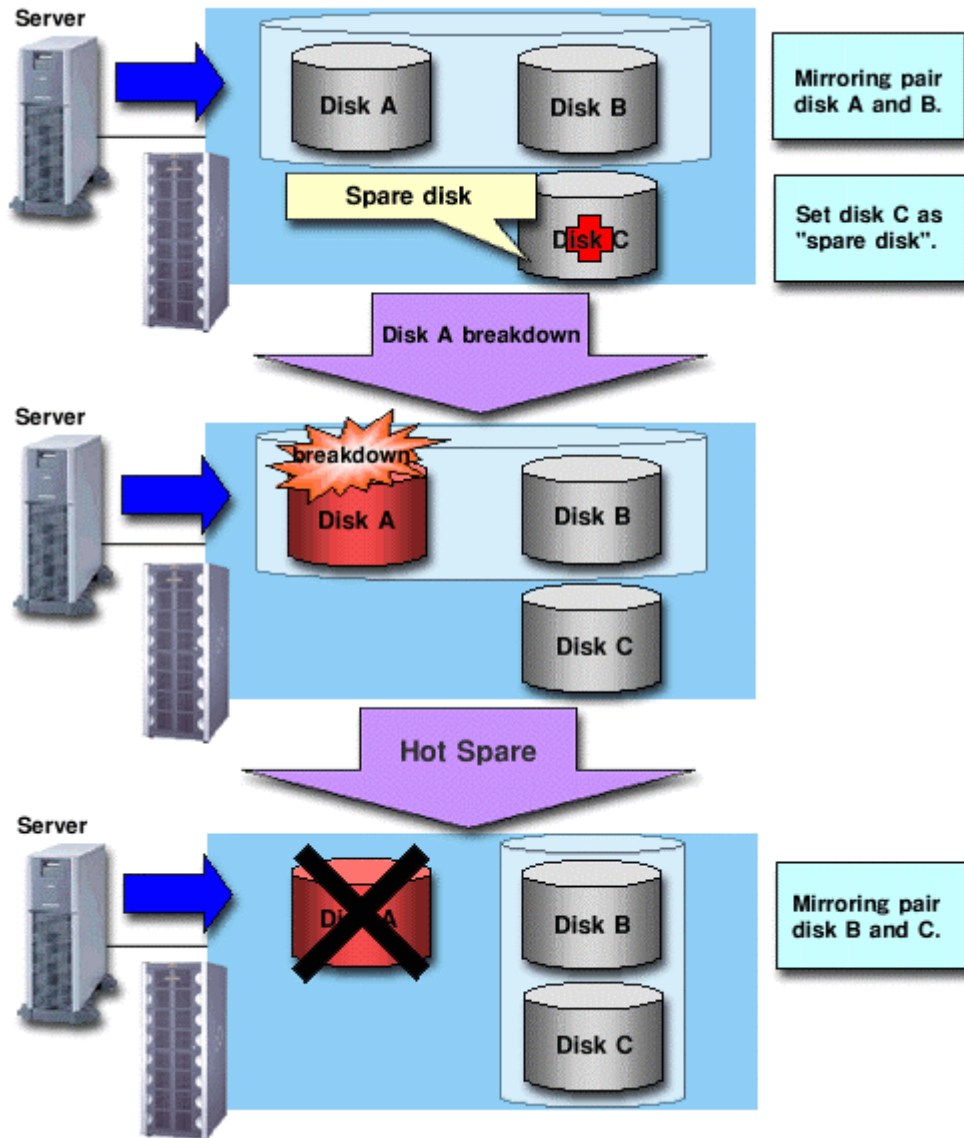
Figure: Shared Disk Mirroring



1.2.2 Hot Spare

Hot spare reserves a spare disk apart from the mirror disks, and automatically restores mirroring when one of the mirror disks breakdown. Even if the administrator does not notice the breakdown immediately, the application will continue to access data safely, until the defective disk unit is swapped with a new one.


Figure: 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. Just select the defective disk on the GUI display, and preparation before the swap, and restoration after the swap will be completed. If the spare disk has been operating in place of the defective disk unit, the spare disk will disconnect automatically, returning to its original state.

Figure: Hot Swap



Note Conditions for Hot Swap

The GDS's function for hot swapping is available only when using a disk unit supporting hot swapping.

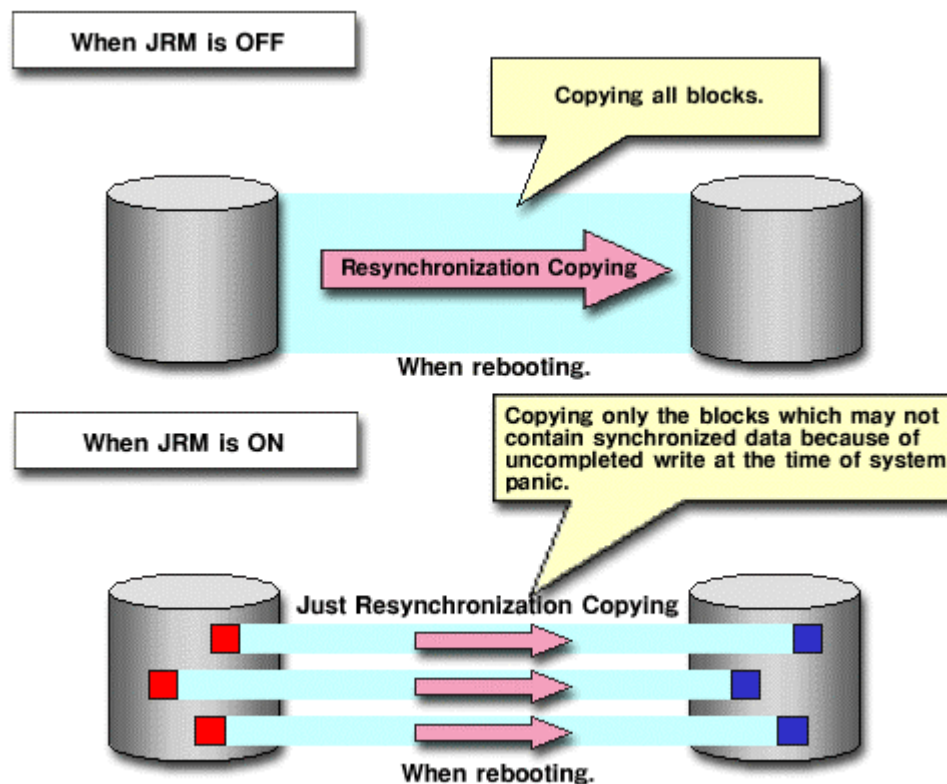
Do not remove or insert a disk when using a disk unit not supporting hot swapping. It may cause breakdown, or damage data.

1.2.4 Just Resynchronization Mechanism (JRM)

When a system failure occurs by a panic and the like, the user must resynchronize the disks after rebooting the computer (after service failover in case of cluster system). Although the user can run an application during this copying process, redundancy will be lost, and load will be generated 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.2.14 Just Resynchronization Mechanism (JRM)."

1.3 Functions for High Manageability

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

Efficient management function is essential in the SAN environment, since one disk will be directly shared by multiple servers.

Functions improving disk manageability are explained below.

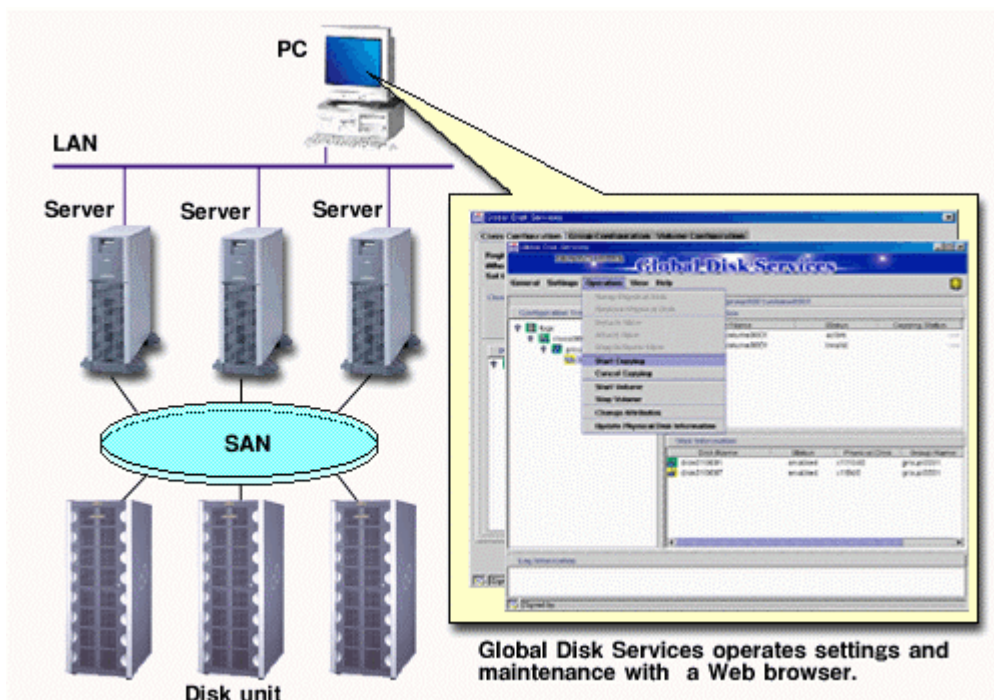
1.3.1 Operation Management Interface

By registering a disk unit with GDS, system administrator 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 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 the inexperienced Solaris(TM) Operating Environment 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



1.3.2 Centralized Disk Management

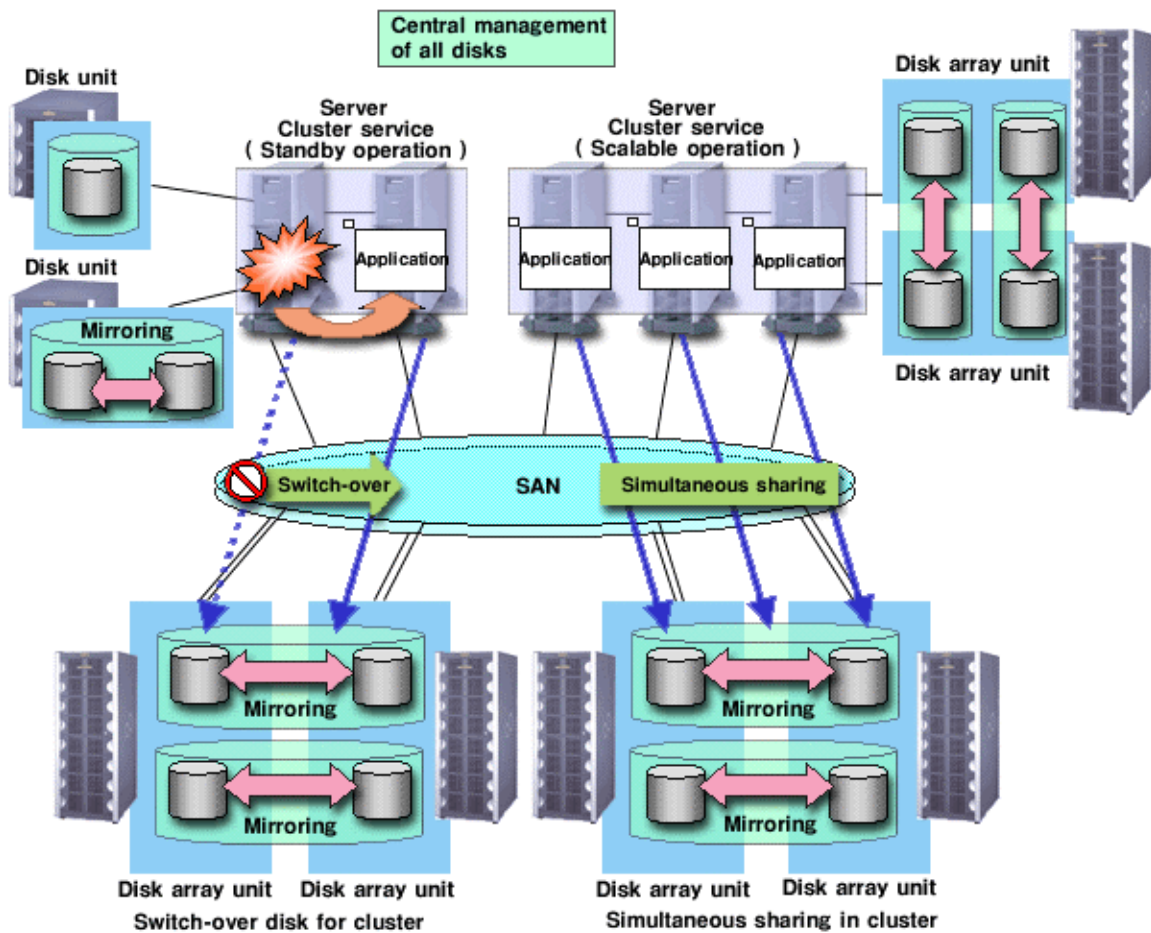
By registering all disk units connected to the server with GDS, the system administrator will be able to centrally manage all operations on the disk via GDS's operation management interface. Since disk units managed by GDS are virtualized as logical volumes, there is no need for the application to handle physical disks.

Various disks, regardless to their physical configuration (single disk, disk array, multi path), connection configuration (local or shared), and usage (system disk, local disk, switch-over disk for cluster) can be managed centrally.

Whether to mirror the managed disk can be determined as necessary.

Mirror volume is suitable in improving disk data availability and manageability using the mirroring function. On the other hand, single volume is useful in improving manageability using the volume management function. For example, a user may want to manage a disk array with sufficient availability, without carrying out mirroring.

Figure: Central Management of Disks



1.3.3 Name Management

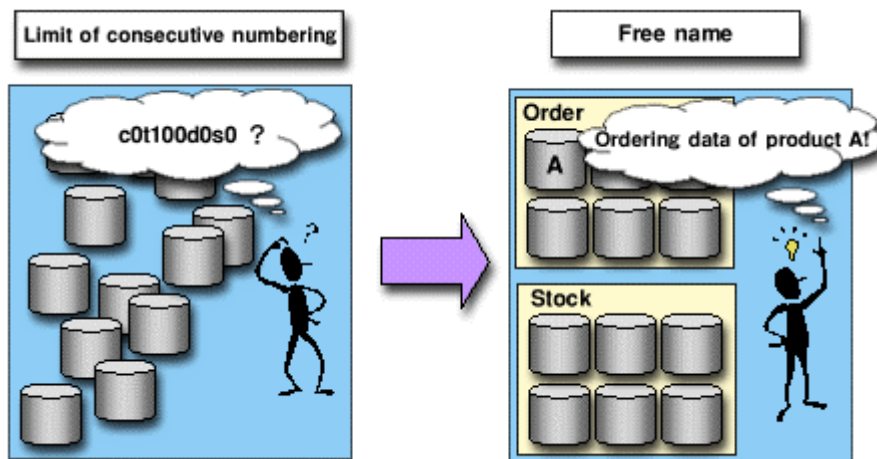
In Solaris(TM) Operating Environment, disks are usually named in "cXtXdX" format, and the administrator differentiates the disks by assigning consecutive numbers.

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

With GDS, the administrator can freely name objects such as physical disks and logical volumes.

Names that are easy to remember, for example, those associated with hardware configuration or data contents can be assigned. Once an object is named, the name remains the same even if the physical configuration is changed. In other words, the user does not have to make any changes to applications that recognize the name.

Figure: Free Device Name



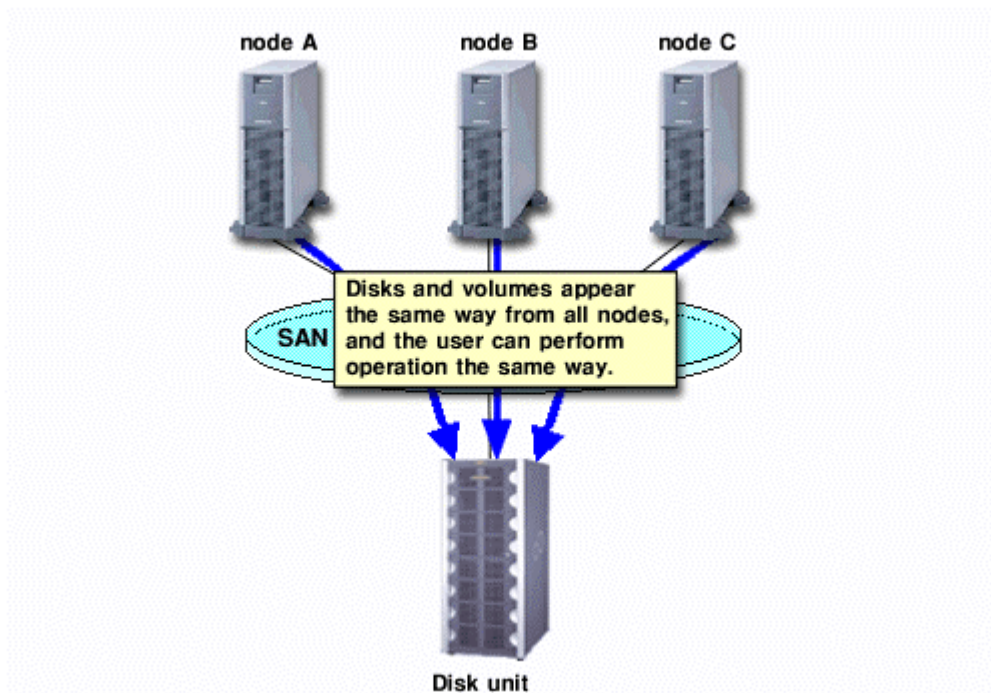
1.3.4 Single System Image Environment

In a SAN environment, disk unit can be accessed from multiple servers. PRIMECLUSTER 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 volume using the same name from all nodes.
- Application programs can access the disk simultaneously from all nodes. For details, refer to "Note".
- Application programs can perform operations from all nodes, such as changing the configuration on objects (disks, volumes, etc). Changes are reflected to 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 an application that is shared simultaneously.

1.3.5 Access Control

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

Class Scope

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

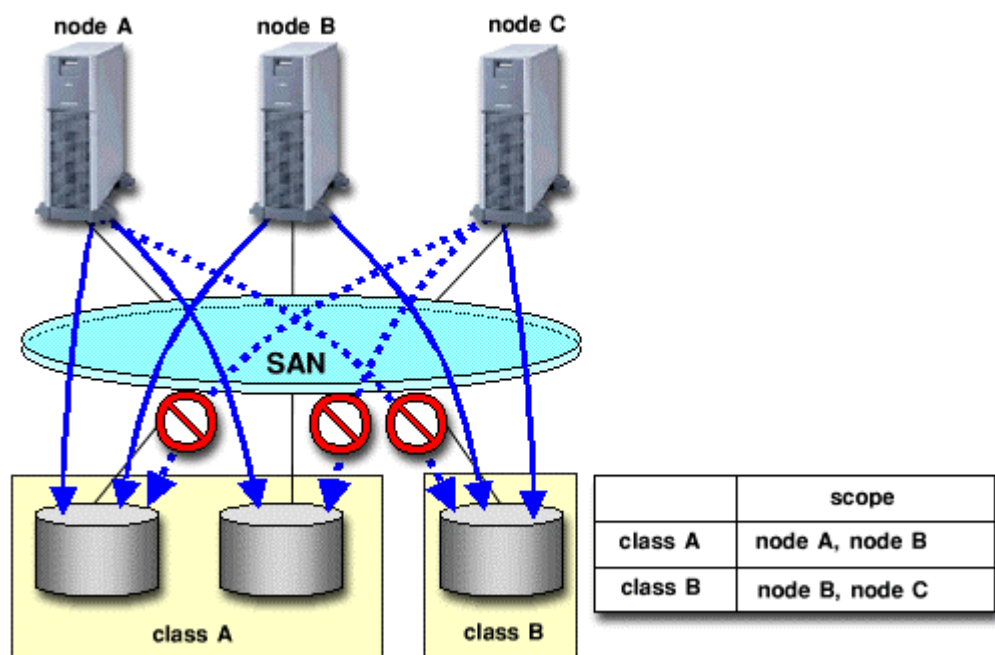
A class is like a container that holds multiple disks.

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

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

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

Figure: Class Scope




Starting and Stopping Volume

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

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

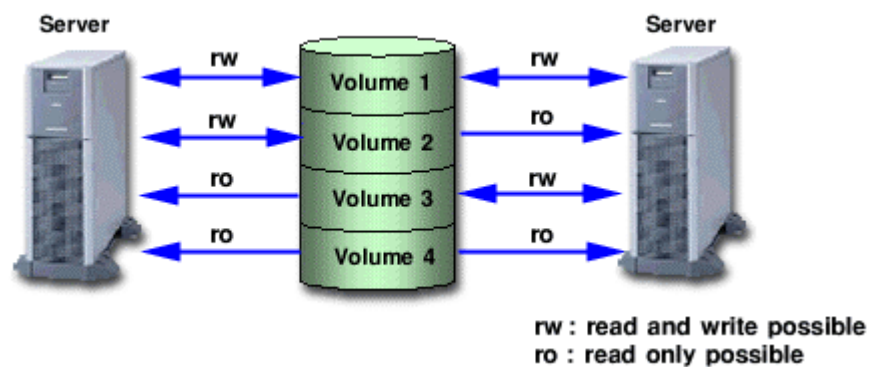
Figure: Starting and Stopping Volume



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.

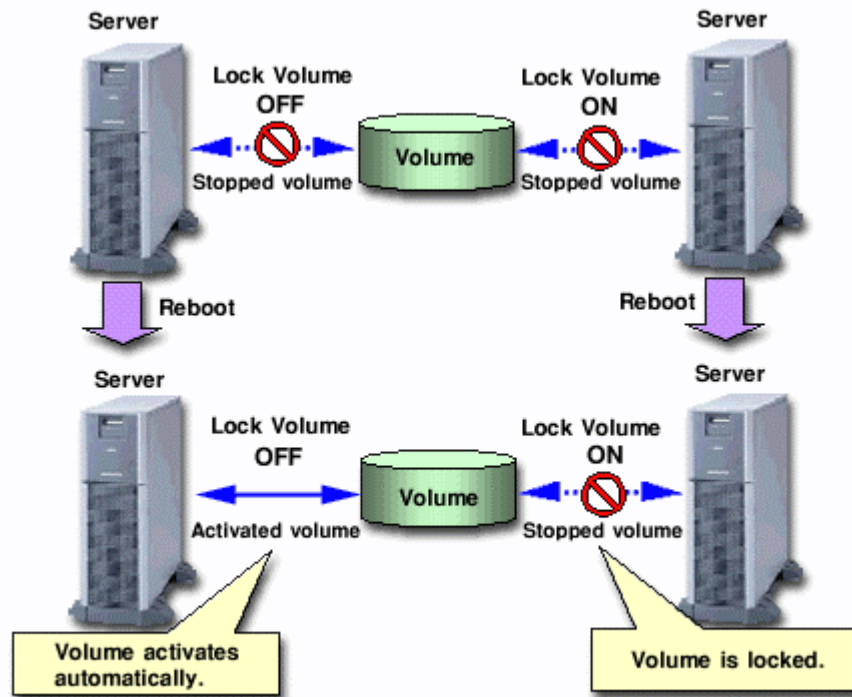
Figure: Access Mode



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.

In order to adapt to such needs, GDS provides functions that realize flexible disk configuration and performance improvement, such as the creation of large-capacity volumes and I/O load balancing.

Since these functions are configured by structural objects similar to those in a mirroring function, users will be able to operate and manage volumes in a consistent manner with other volume management tasks including mirroring.

1.3.6.1 Logical Partitioning

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

In a Solaris(TM) Operating Environment, physical disks can only be partitioned to a maximum of seven disk slices (physical slices).

GDS allows users to divide physical disks and objects equivalent to disks into a maximum of 254 logical devices.

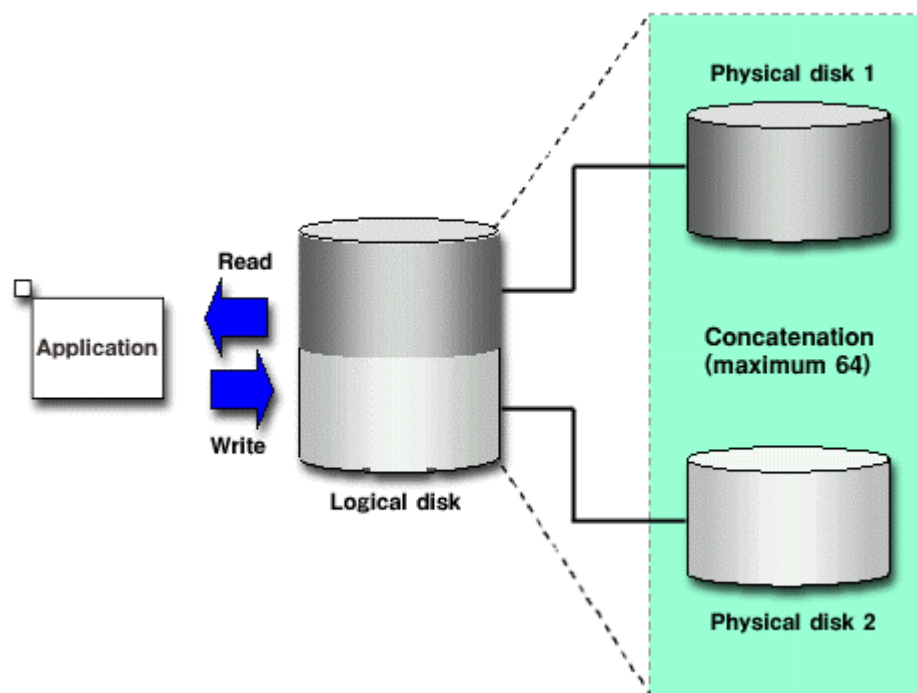
For details, see "2.5 Logical Volume" and "D.4 sdxvolume - Volume operations".

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



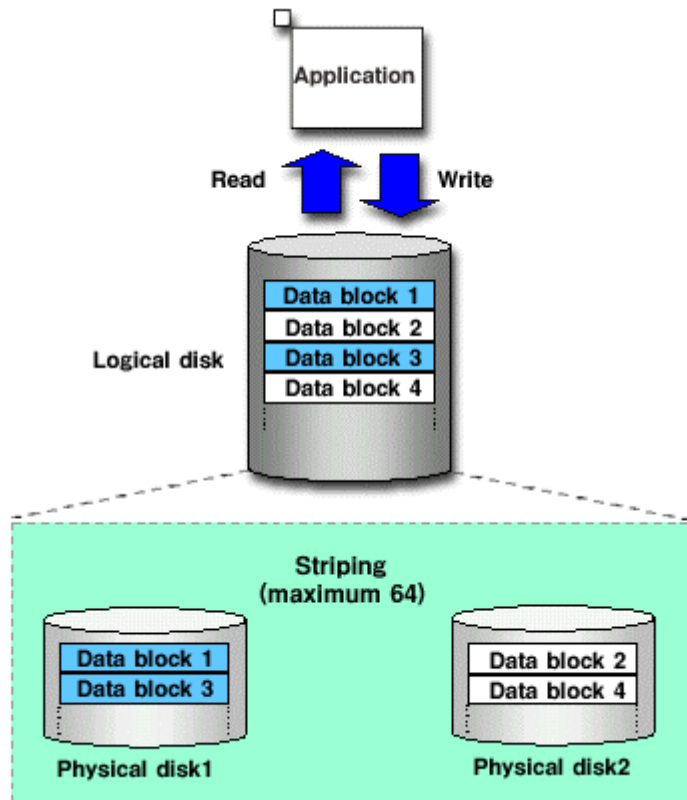
1.3.6.3 Disk Striping

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 together.

Figure: Combining Striping and Mirroring



1.3.7 Snapshot by Detaching Slice

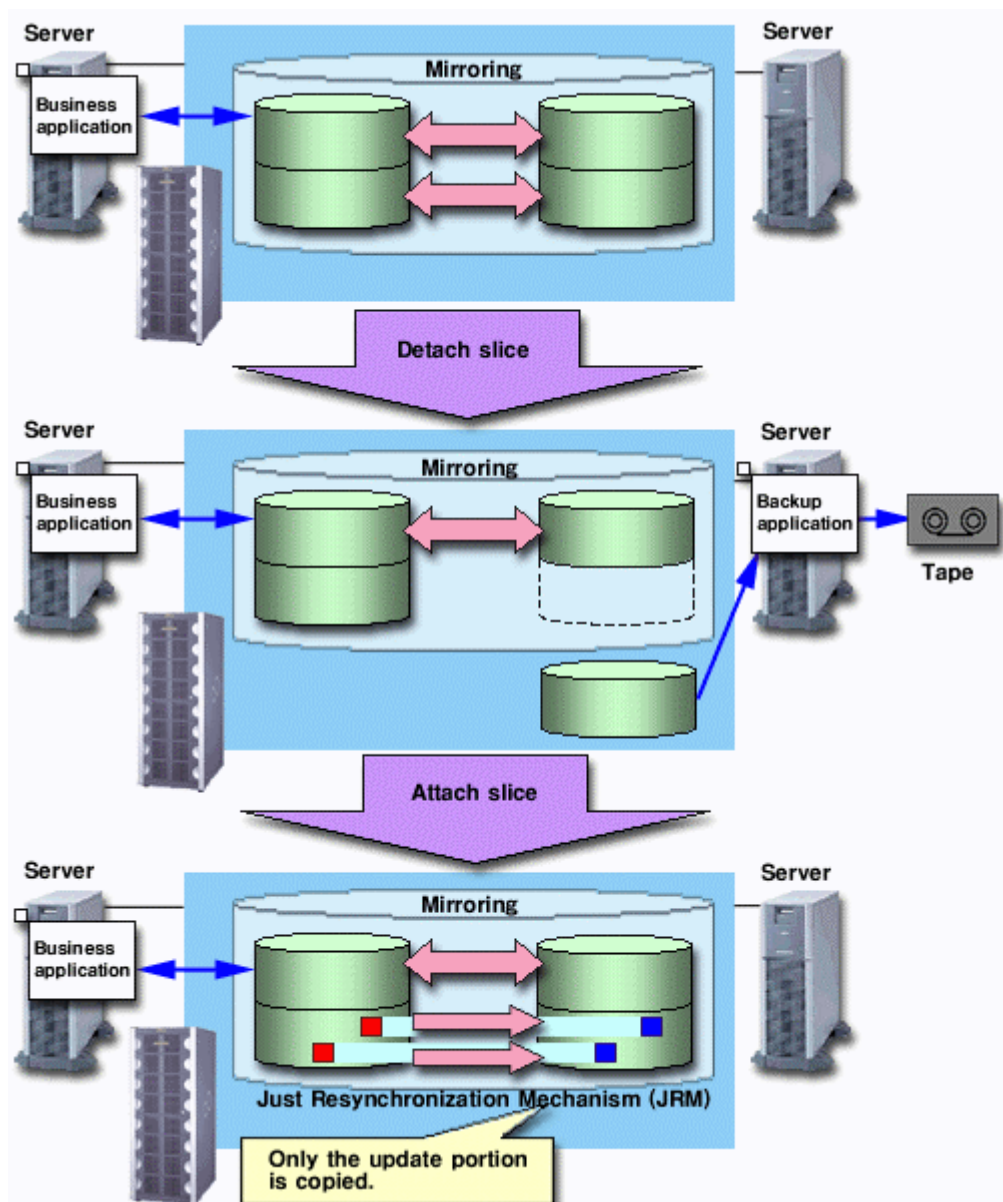
By temporarily detaching a slice from the mirror volume, the user can create a snapshot (replication) of the volume. This replication will contain volume data up to the point when the slice was detached.

A detached slice is a logical device separate from the volume, and another application using the detached slice can run simultaneously with the business application using the volume. For example, a backup application can be executed concurrently with the business application.

When multiple nodes are sharing a volume, the user can run the business application which uses the volume on one of the nodes, and another application which uses a detached slice on another node, so that the processing load between the nodes can be distributed.

After the application which uses the detached slice is completed, in order to create a snapshot again, it is necessary to re-attach the detached slice to the volume. At this point, copying will be performed to synchronize data on the slice and volume. This copying uses JRM (Just Resynchronization Mechanism) for slices, and realizes high-speed resynchronization by copying only the portion that was updated while the slice was detached.

Figure: Snapshot Created by Detaching Slice





Note **Ensuring Integrity of Snapshot Data**

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

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

After creating the snapshot, start the application again.

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

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



Note **JRM for Slices**

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

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

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



Information **What is JRM?**

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



Information **Three Types of JRM**

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

1.3.8 Proxy volume

A snapshot of the volume for the business application (a replication at a certain moment) can be created on another volume. The former is referred to as a master volume and the latter as a proxy (or an alternate) volume.

You can create a snapshot, or a copy of the master volume using the proxy volume.


Usage of this alternate volume can resolve various problems.

For example, stopping a system with a large amount of data for backup purpose 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 the system.


Figure: Problems Inherent to a 24-hour Operation System



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

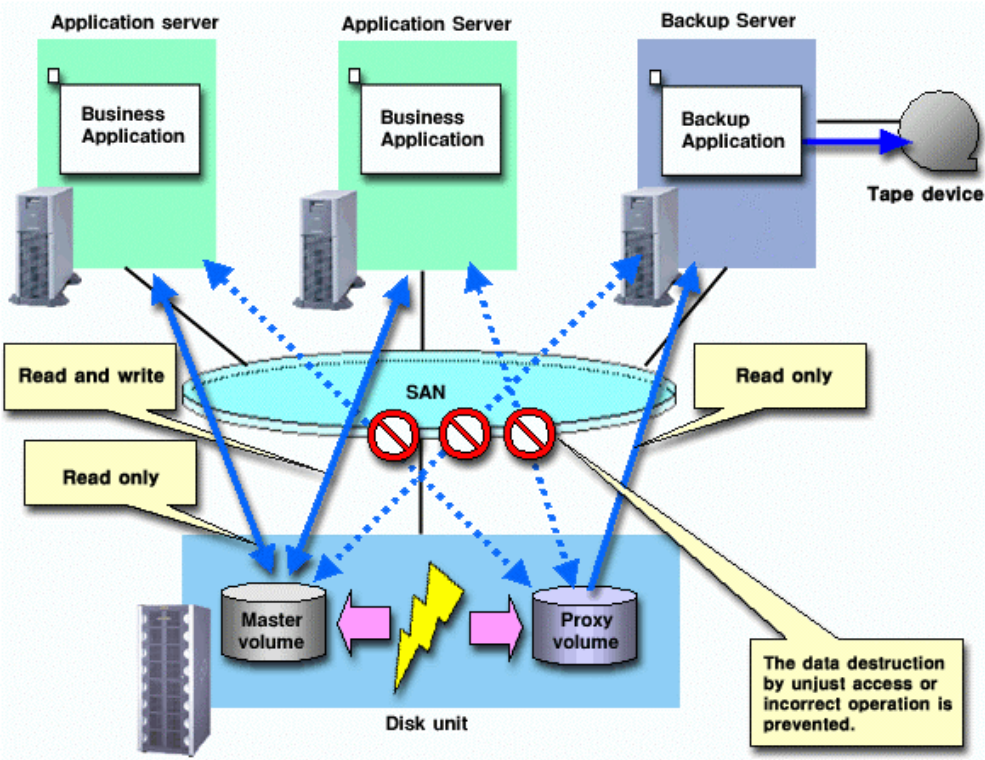
By using the proxy volume, users can conduct various tasks such as non-disruptive backup, data analysis, verification, and data protection from disaster, without affecting the main business application.

Figure: Backup Using the Proxy Volume



It allows users to utilize any logical volumes as a master or proxy volume, including a single server local volume and shared volumes in a cluster system.
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



The snapshot functions using a proxy volume are explained here.

1.3.8.1 Snapshot Created by Synchronization

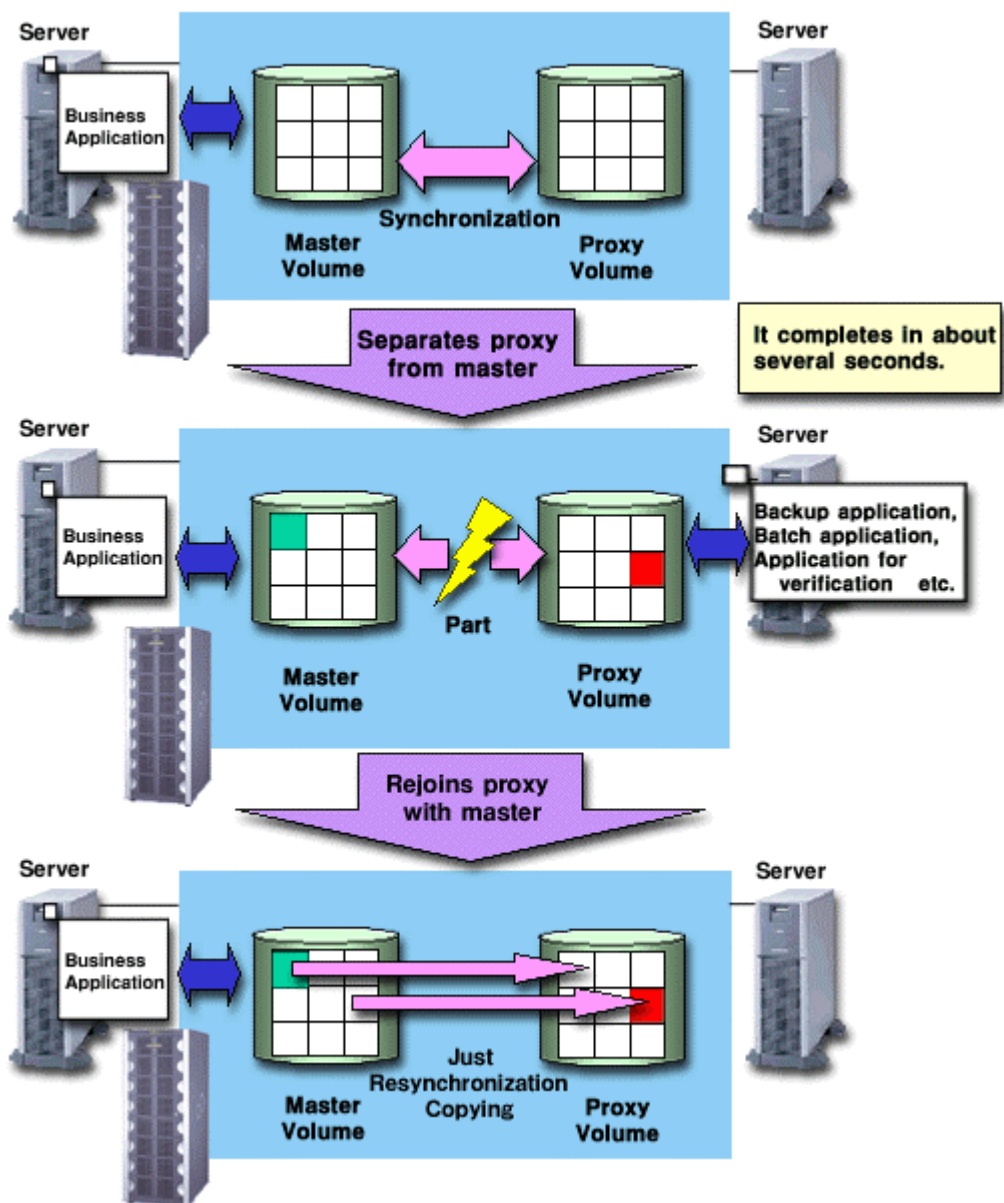
A master volume snapshot can be created by synchronizing data of the master volume and the proxy volume in advance and then separating the synchronized proxy volume from the master volume at any moment.

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

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

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

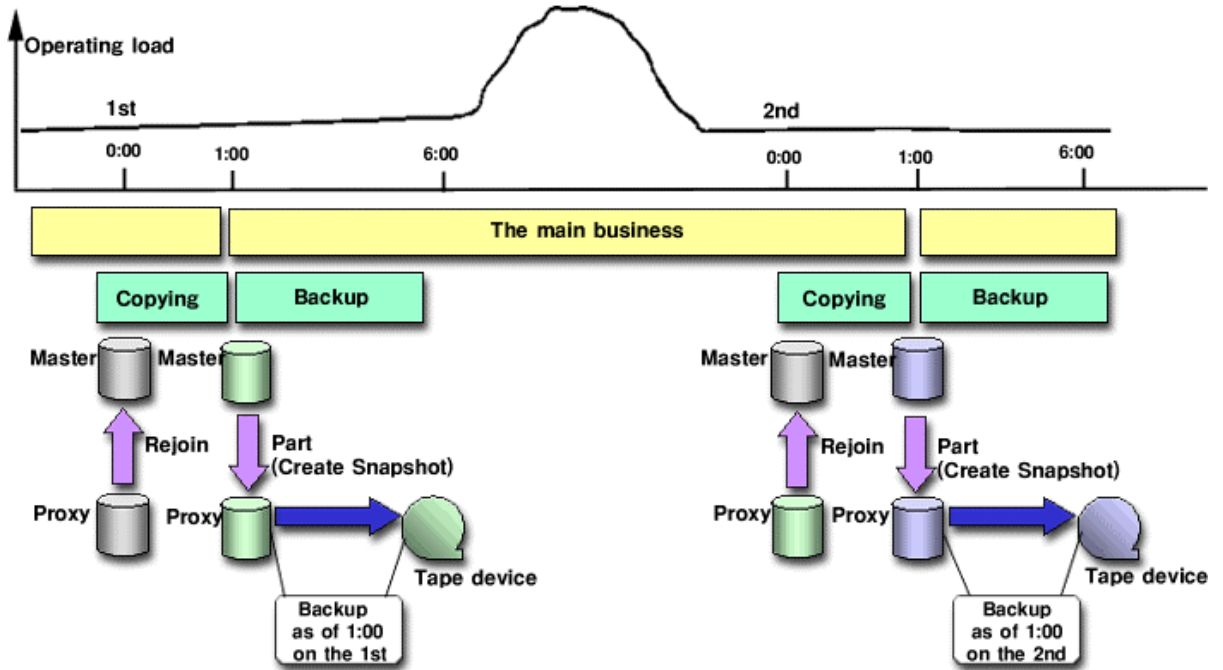
Figure: Snapshot Created by Synchronization



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

For example, in a situation where the disk load is higher during the day than at night, and saving the data to tape takes less than 5 hours, and the resynchronization copying less than an hour, the routine backup schedule will be as follows.

Figure: Example of a Daily Backup Schedule



Note Ensuring integrity of snapshot data

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

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

After creating the snapshot, start the application again.

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

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

For details, see "D.16 Online Backup Making Using Command."



Note **JRM for Proxies**

JRM for proxies speeds up the resynchronization process when joining a parted proxy again to the master. GDS records the changes made on the master and proxy on the memory while the proxy is parted. The resynchronization copy performed when re-joining the parted proxy copies only the updated portions to realize high-speed resynchronization.

Unlike JRM for volumes or slices, JRM for proxies is always enabled and cannot be turned off.

However, if there is an occurrence of a system shutdown while the proxy is parted, just resynchronization is not performed when the proxy is later re-joined. Resynchronization is performed by copying the entire data, not only the updated portions.

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

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



Information **Three Types of JRM**

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



Note **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.

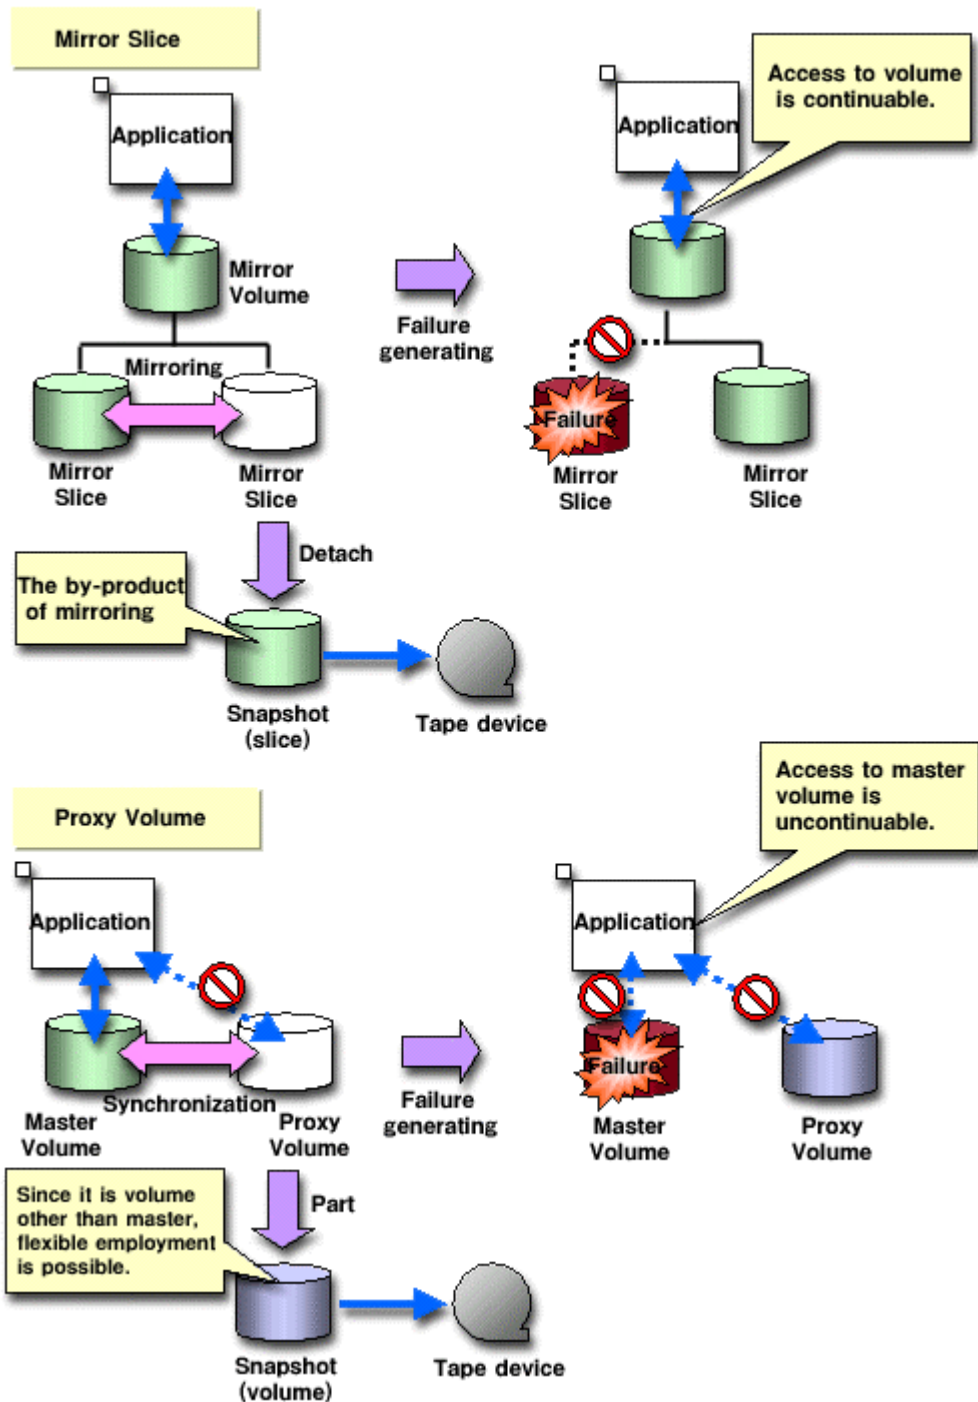
Mirrored slices are equals, and their purpose is to maintain data redundancy. So even if an abnormality is detected in one of the slices, the volume can be accessed as long as the other slice is normal.

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 is 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 business application.

In short, while the purpose of GDS Snapshot is to create a snapshot with a proxy volume, the snapshot created by detaching a slice is a by-product of mirroring. Therefore, GDS Snapshot provides more flexible disk configuration and supports various tasks.

Figure: Difference between a Mirrored Slice and Synchronized Proxy Volume

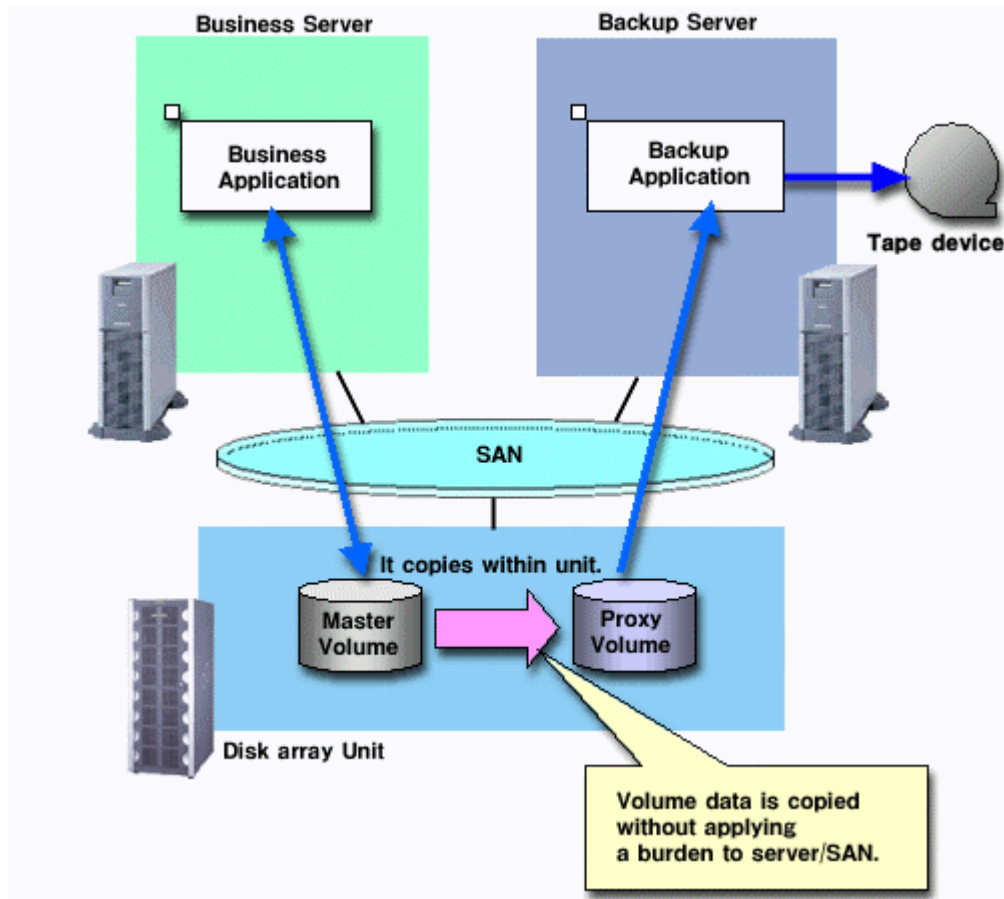


1.3.8.2 Snapshot Function with No Load to the server/SAN

By cooperation with Advanced Copy function of ETERNUS3000 or GR700 series storage system, or EMC TimeFinder and SRDF for Symmetrix storage device, processes such as copying for resynchronization, copying for maintaining synchronization, and recording the portion upgraded while the proxy was being parted, are performed within the disk device. Therefore, snapshots can be created without imposing a load on the server accessed by the business 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



1.3.8.3 Instant Snapshot Created by OPC

By liaising with OPC (One Point Copy) function provided by ETERNUS3000 or GR700 series storage system, users can create snapshots instantly at any given moment.

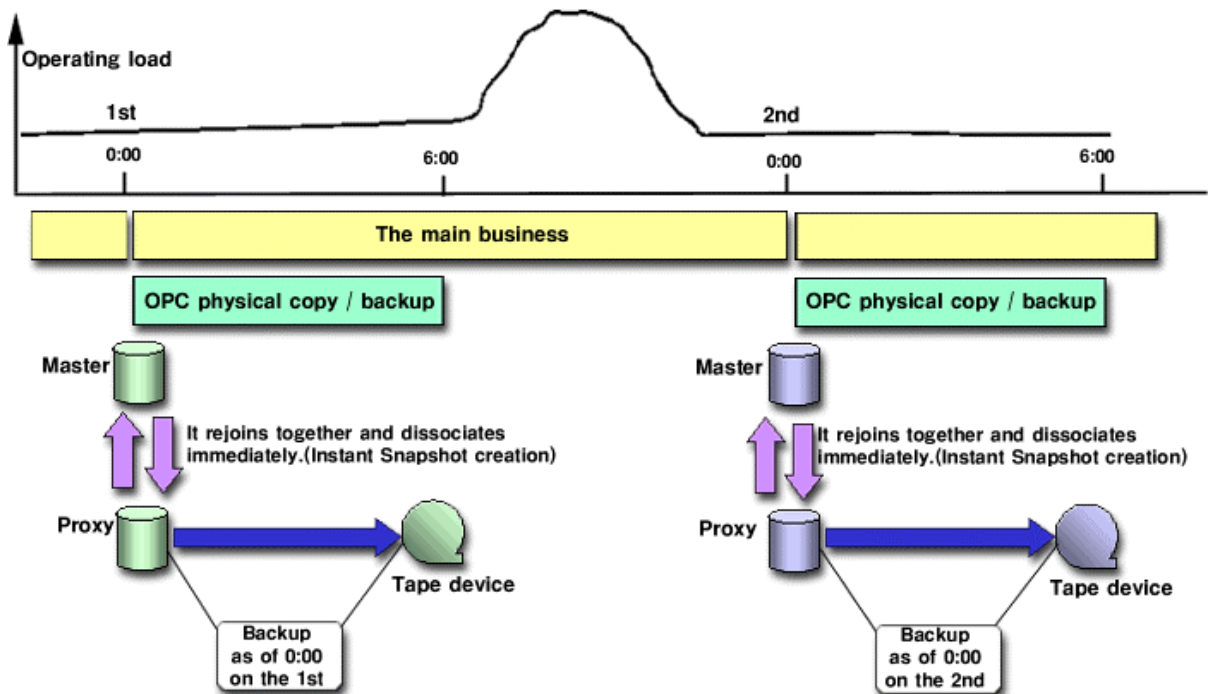
Unlike snapshot created by synchronization, instant snapshot created by OPC does not require to synchronize the volumes in advance, so scheduling is unnecessary. However, as copying process will take place within the disk array after the snapshot has been created, I/O performance of the disk will be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



Note Instant Snapshot Created by OPC

Instant snapshot created by OPC is only available with the following storage systems supporting the function: ETERNUS3000, GR720, GR730, and GR740.

Figure: Instant Snapshot Created by OPC



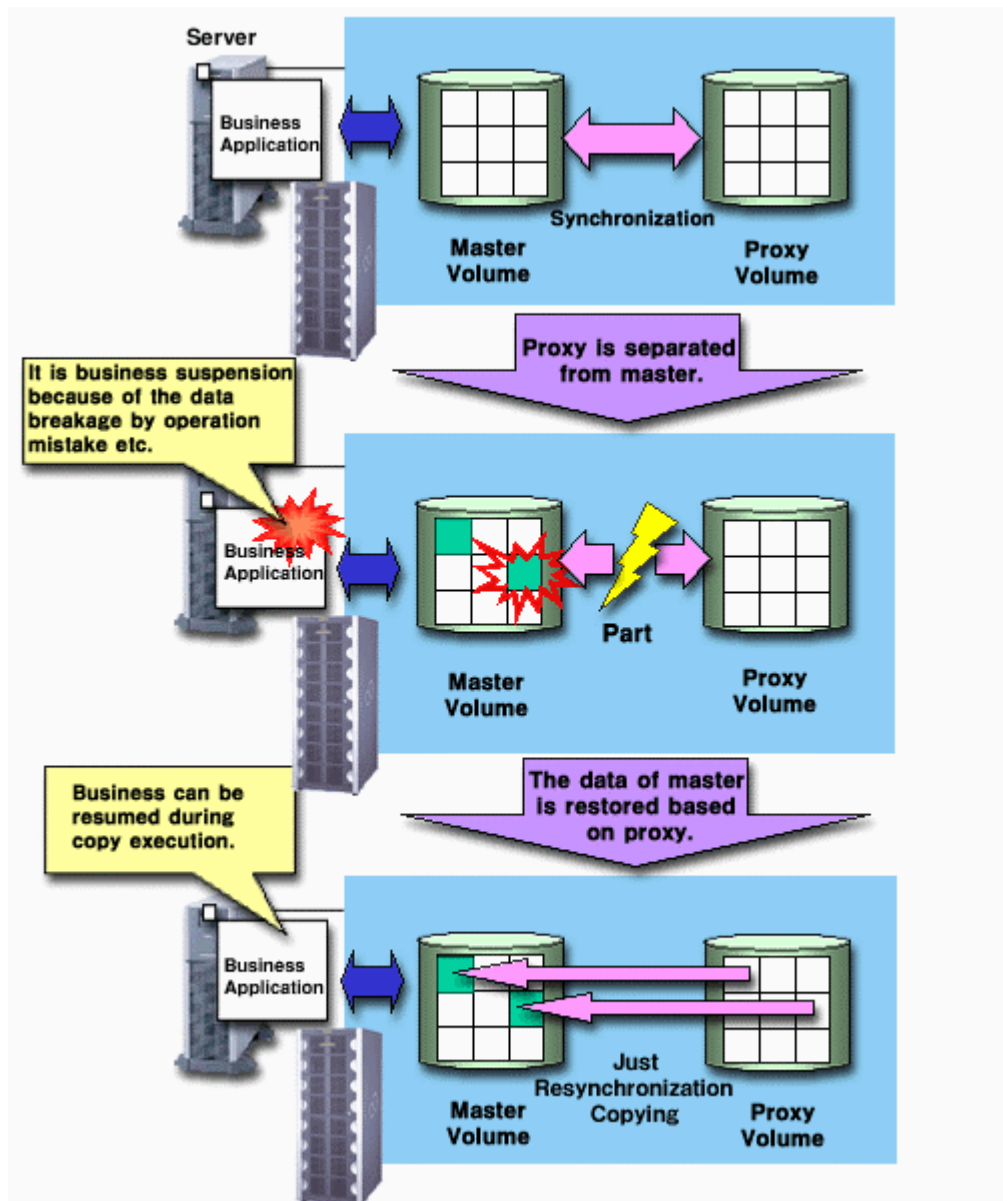
1.3.8.4 Instant Restore

If for example, a misoperation damages the 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 business application immediately without waiting for the completion of the copying process. Synchronization copying process will take place in parallel to other tasks.

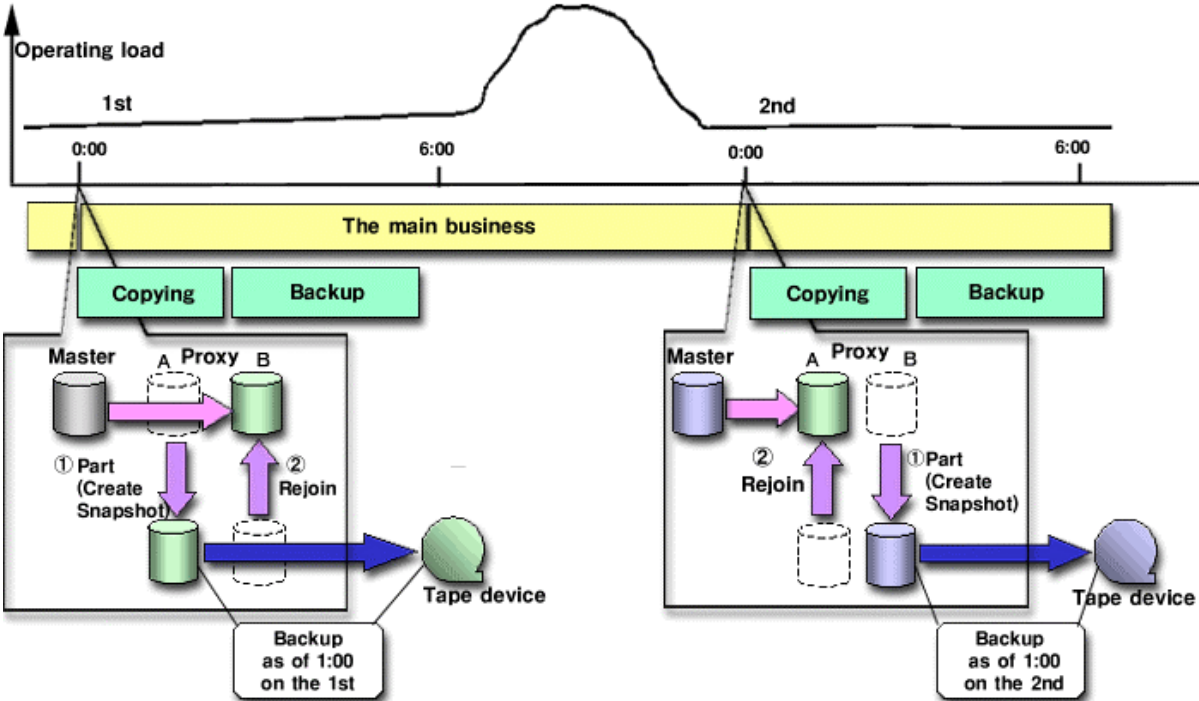
Users may use this function with any disk unit, as it is not dependent on a specific function of a disk unit.

Figure: Instant Restore



For example, in a situation where the disk load is higher during the day than at night, and saving the data to tape takes less than 5 hours, and the resynchronization copying less than an hour, if you intend to always perform instant restore from the disk (proxy volume), and not from a tape, the routine backup schedule will be as follows.

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

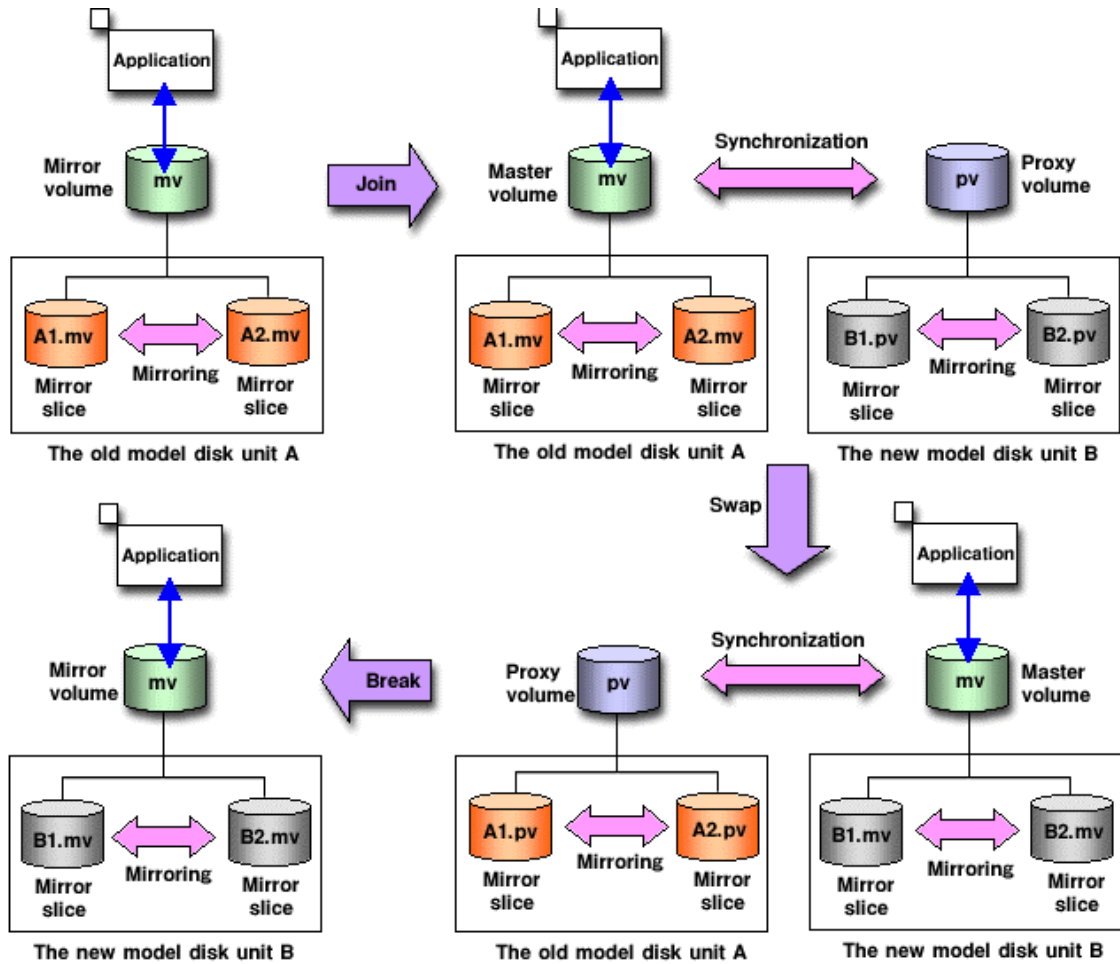


1.3.8.5 Online Disk Migration

When the master and proxy volumes are synchronized, users can migrate from the physical disk used as the master volume without affecting the business 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



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.

2.1 SDX Object

Virtual resources managed by GDS are called , 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.

You can freely name all objects including the class. All names should be unique, as they are managed using the given name throughout the system.

Attributes

A class has the following attributes.

Name

This attribute identifies the class within a system.

Type

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

Root

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

Local

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

Shared

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

Scope

When a class is a "shared" type, it displays the node set sharing the class.

Hot spare

It displays the operation mode of the hot spare.

Operation

Create

Registering a disk with the `sdxdisk -M` command and specifying a new class name will automatically create a new class. For details on GDS Management View, see "5.2.3.1 Class Configuration."

Delete

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

Status Display

Class status can be displayed with the `sdxinfo` command. For details on GDS Management View, see "5.3.1 Confirming Configuration and Monitoring Status."

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."

Status

Closed

All objects within a class are unavailable for reference. For details, see "F.1.4 class status abnormality."

Reference

Rules


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

Important Points

- A.2.1 Physical Disk Driver

- A.2.2 Managing System Disks
- A.2.4 Booting from a CD-ROM Device
- A.2.9 System Disk Mirroring
- A.2.23 Configuration Operation in Cluster System

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 specifies the type of disk. You can set to one of the followings.

Keep

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

Single

A volume can be created without connecting to a group.

Spare

The disk will be used as a spare disk.

Undefined

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

Operation

Create

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

Delete

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

Connect

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

Disconnect

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

Swap

The disk will be prepared for swapping with the `sdxswap -O` command.
For details on GDS Management View, see "5.3.3 In Case of Disk Abnormality."

Restore

Swapped disk will be restored with the `sdxswap -I` command.
For details on GDS Management View, see "5.3.3 In Case of Disk Abnormality."

Status Display

The disk status will be displayed with the `sdxinfo` command.
For details on GDS Management View, see "5.3.1 Confirming Configuration and Monitoring Status."

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 the Management View is not supported.

Status

ENABLE

Normal status of a disk.

DISABLE

The disk cannot be accessed.

SWAP

The disk cannot be accessed, but may be swapped.

Reference

Rules

- A.1.3 Number of Disks
- A.1.6 Number of Spare Disks
- A.1.7 Number of Keep Disks

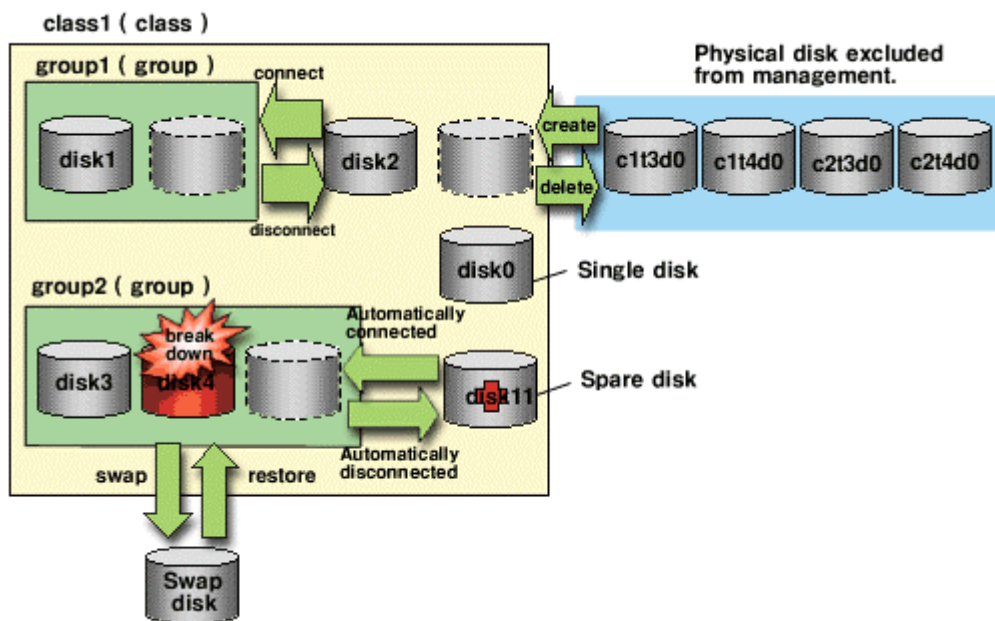
Important Points

- A.2.3 Restraining Access to Physical Special File
- A.2.5 Initializing Disk
- A.2.6 Disk Size
- A.2.8 Spare Disk
- A.2.10 Keep Disk
- A.2.11 GFS Partition on a Keep Disk

Information

- A.3.1 Spare Disk Size

Figure:SDX Disk



2.4 Disk Group

Disk group acts as container that holds the disks that will be mirrored, striped or concatenated. It may be simply referred to as "group."

Among group attributes is type attribute. Type attribute value may be "mirror," "stripe" or "concat." By connecting more than one disk to a group, disks within that group will be mirrored, striped or concatenated, 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 total of eight disks and lower level groups combined may be connected to the group. 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, you can concatenate up to 64 disks.

Stripe Width

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

Operation

Create

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

For details on GDS Management View, see "5.2.3.3 Group Configuration."

Delete

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

For details on GDS Management View, see "5.5.3 Removing a Group."

Connect

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

For details on GDS Management View, see "5.2.3.3 Group Configuration."

Disconnect

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

For details on GDS Management View, see "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 Confirming Configuration and Monitoring Status."

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

Rules

A.1.4 Number of Groups

A.1.8 Creating Group Hierarchy

Information


A.3.2 Guidelines for Mirroring

A.3.3 Guidelines for Striping

A.3.4 Guidelines for Concatenation

A.3.5 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 four kinds of volumes: a single volume which is created in a single disk, a mirror volume which is created in a mirror group, a stripe volume which is created in a stripe group and a volume created in a concatenation group.

Since a single volume is created in a single disk, its data will not be redundant. Single volumes are used when it does not need 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.

Since mirror volumes are created within mirror groups, as long as more than one disk or one lower level group is connected, data will be made redundant with the mirroring function.

Stripe volume can distribute the I/O load among multiple disks with its striping function. Its data will not be redundant.

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

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

In a Solaris(TM) Operating Environment, physical disks can only be partitioned to a maximum of seven disk slices (physical slices). GDS allows users to create both a volume with a corresponding physical slice, and a volume without a corresponding physical slice. So if you total the volumes with and without physical slices, single disks and groups can be partitioned to a maximum of 254 volumes.

Attributes

A volume has the following attributes.

Name

This attribute identifies the volume within a class.

JRM

This attribute sets the just resynchronization mechanism. The value can be set to one of the followings.

on

JRM is active.

off

JRM is inactive.

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 is set to read and write possible.

ro

The default is set to read only possible.

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 VTOC. The value can be set to one of the followings. However, note that physical slice attribute for volumes that are created in stripe group or concatenation group must be set to "off."

on

When the volume is a single volume, a slice in the single disk will be registered with the VTOC. When the volume is a mirror volume, and if there are disks directly connected to the mirror group, the slices on the disks will be registered with the VTOC. However, if a mirror volume is created in a mirror group to which only lower level groups are connected, the mirror volume will not have any physical slice regardless whether this attribute set to "on."

off

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

Operation

Create

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

For details on GDS Management View, see "5.2.3.4 Volume Configuration."

Delete

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

For details on GDS Management View, see "5.5.2 Removing a Volume ."

Start

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

For details on GDS Management View, see "4.2.3 Operation ."

Stop

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

For details on GDS Management View, see "4.2.3 Operation ."

Status Display

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

For details on GDS Management View, see "5.3.1 Confirming Configuration and Monitoring Status ."

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 "5.3.4 Copying Operation ."

Restore

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

The operation from the Management View is not supported.

Status

ACTIVE

Data is normal, and accessible.

STOP

Not accessible, but can be activated and changed to ACTIVE status. However, in case of proxy volume, it has to be parted from the master volume before activating.

INVALID

Data is invalid, and not accessible.

Reference

Rules

A.1.5 Number of Volumes

Important Point

A.2.7 Volume Size

A.2.22 Volume Access Mode

Figure: Logical Volume

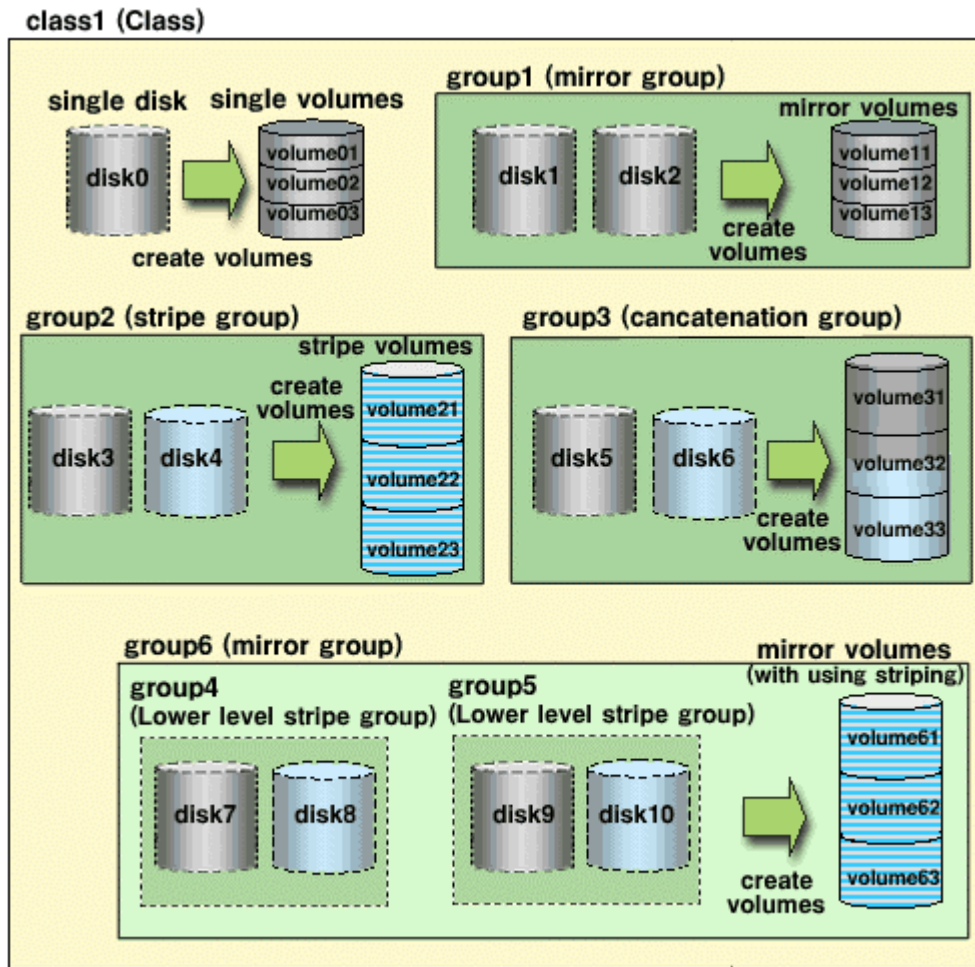



Figure: Logical Partitioning (Physical Slice Attribute)



2.6 Logical Slice

Logical slices are the smallest objects managed by GDS. 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.

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 sets the just resynchronization mechanism. The value can be set to one of the followings.

on

JRM is active.

off

JRM is inactive.

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

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 Backup."

Attach

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

For details on GDS Management View, see "5.3.2 Backup."

Activate

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

For details on GDS Management View, see "5.3.2 Backup."

Stop

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

For details on GDS Management View, see "5.3.2 Backup."

Take over

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

The operation from the Management View is not supported.

Status Display

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

For details on GDS Management View, see "5.3.1 Confirming Configuration and Monitoring Status."

Change Attributes

Attribute values of the detached slice can be changed with the `sdxattr -S` command.

The operation from the Management View is not supported.

Status

ACTIVE

Data is normal, and accessible.

STOP

Data is normal, but not accessible.

INVALID

Data is invalid, and not accessible.

COPY

Synchronization copying is in progress in order to validate the data.

TEMP

The slice is temporarily detached, and can be accessed in isolation.

TEMP-STOP

The slice is temporarily detached, but cannot be accessed in isolation.

NOUSE

The slice is not accessible for a special reason.

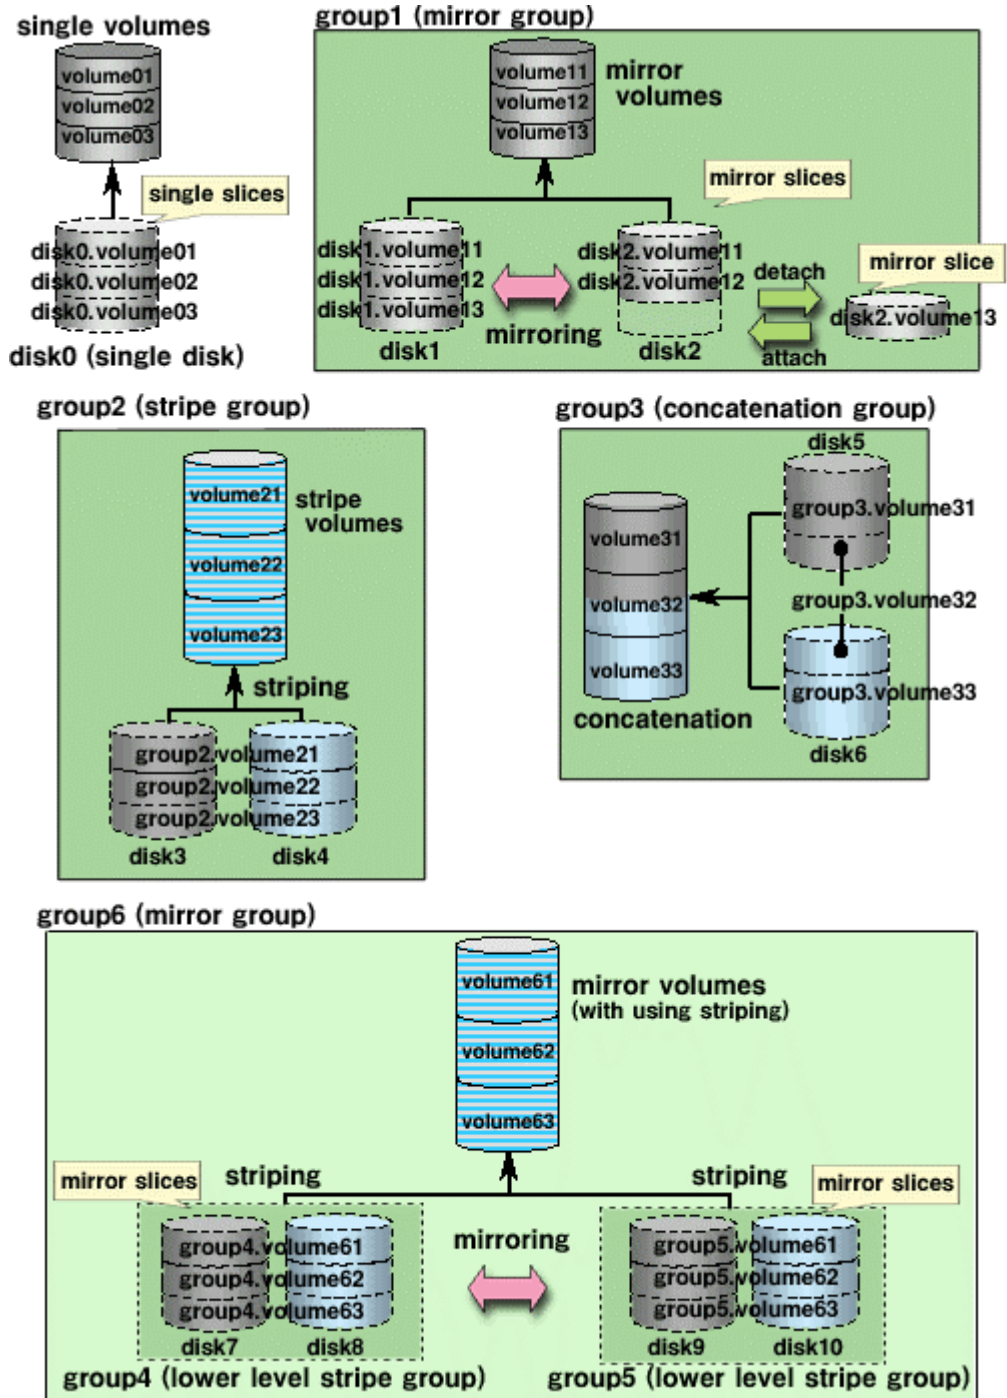
Reference

Important Points

A.2.12 Creating a Snapshot by Detaching Slice

A.2.14 Just Resynchronization Mechanism (JRM)

Figure: Logical Slice



2.7 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 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 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 operation from the Management View is not supported. You must use respective commands.

Join

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

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.

Part

Master-proxy synchronization will be canceled and the proxy will be accessible as another device 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.

Rejoin

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

Restore

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

Swap

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

Status

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)

The status can be checked in the PROXY volume field using the `sdxinfo` command.

Reference

Rules

- A.1.9 Exception to Proxy Configuration
- A.1.10 Number of Proxy Volumes
- A.1.11 Proxy Volume Size
- A.1.12 Proxy Group Size

Important Points

- A.2.13 The Difference between a Mirror Slice and a Proxy Volume
- A.2.14 Just Resynchronization Mechanism (JRM)
- A.2.15 Object Operation When Using Proxy
- A.2.16 Snapshot by Cooperation with Advanced Copy of ETERNUS3000 or ETERNUS GR700 series
- A.2.19 Snapshot by Cooperation with EMC TimeFinder or EMC SRDF
- A.2.17 Instant Snapshot Created by OPC
- A.2.20 Ensuring Integrity of Snapshot Data

Figure: Operating Proxy Volume

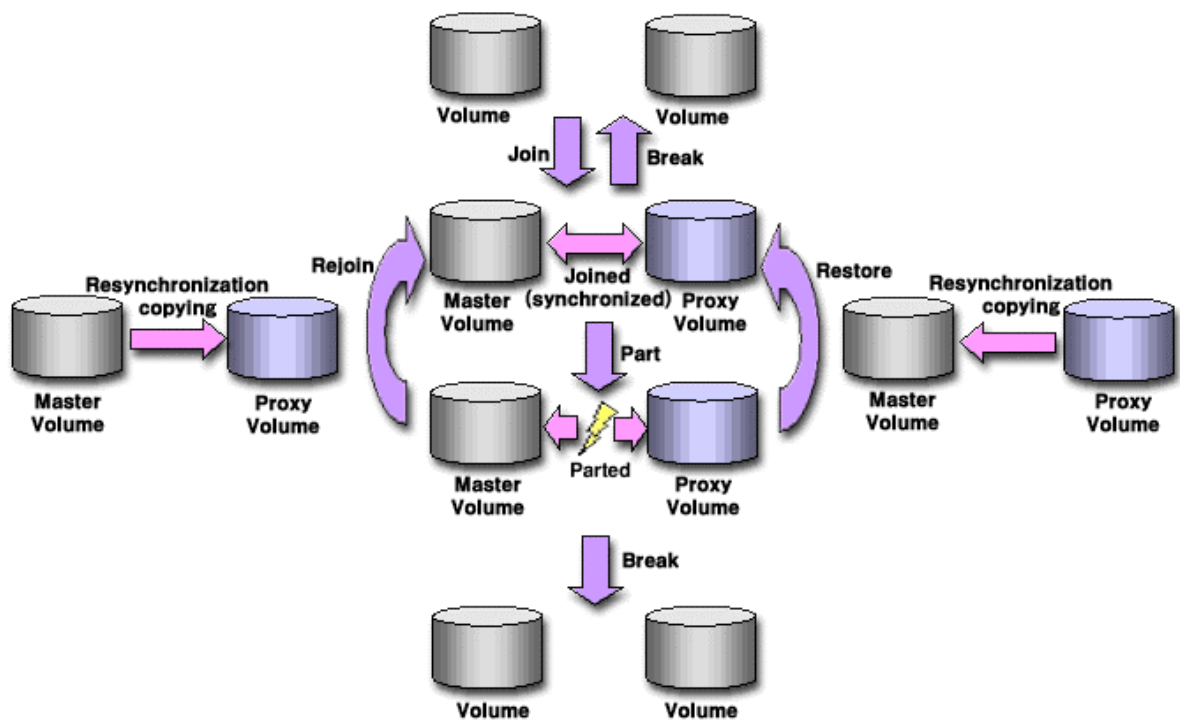
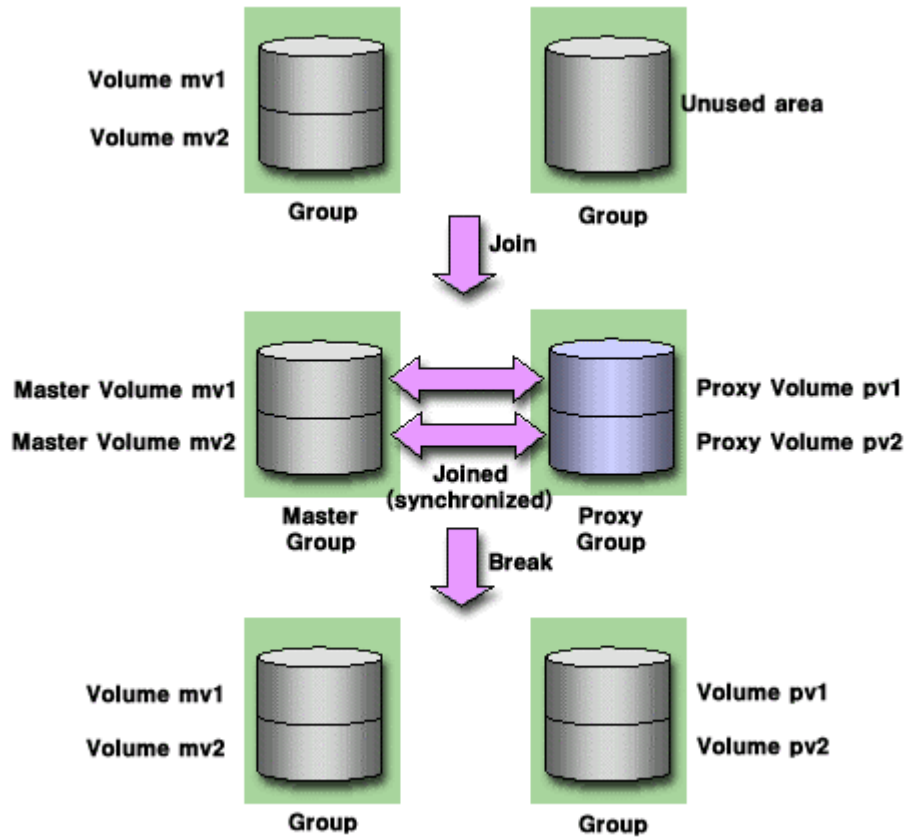


Figure: Operating Proxy Group



* Part, Rejoin, Restore are same as "Cart: Operation of Proxy Volume".

Figure: Swapping Master Volume and Proxy Volume

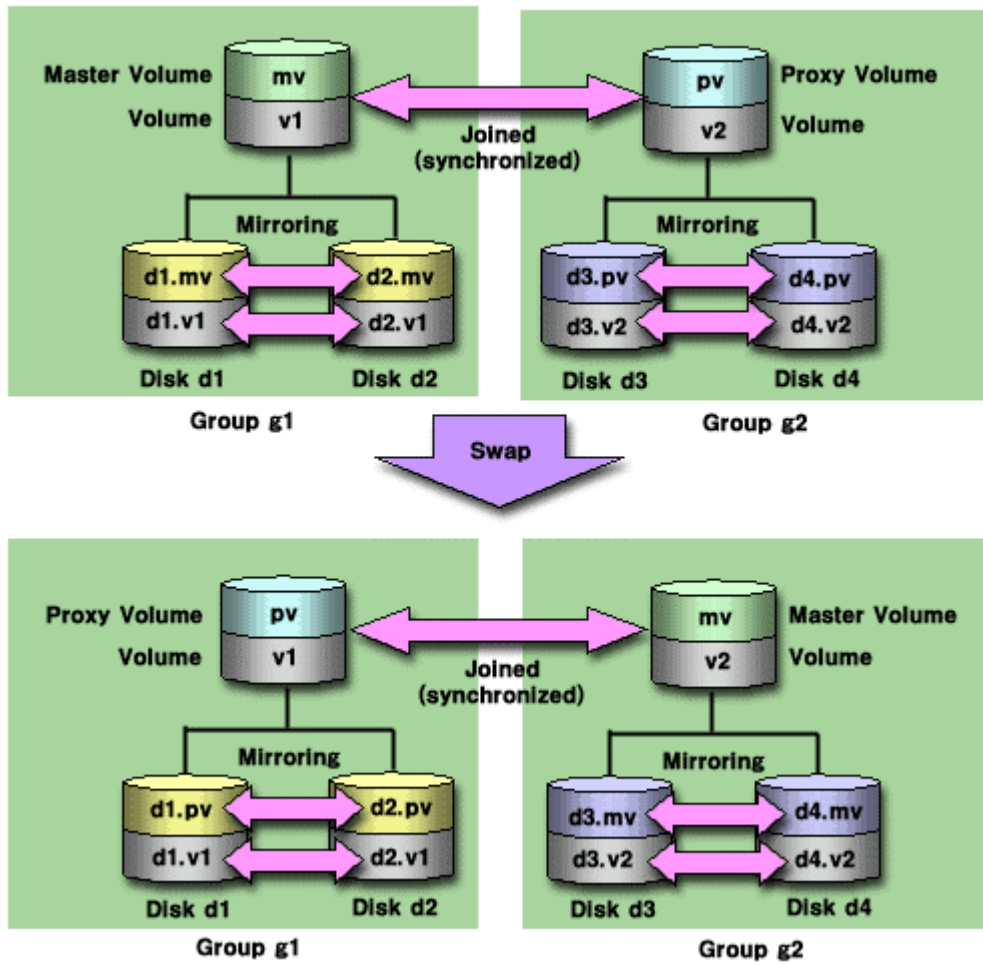
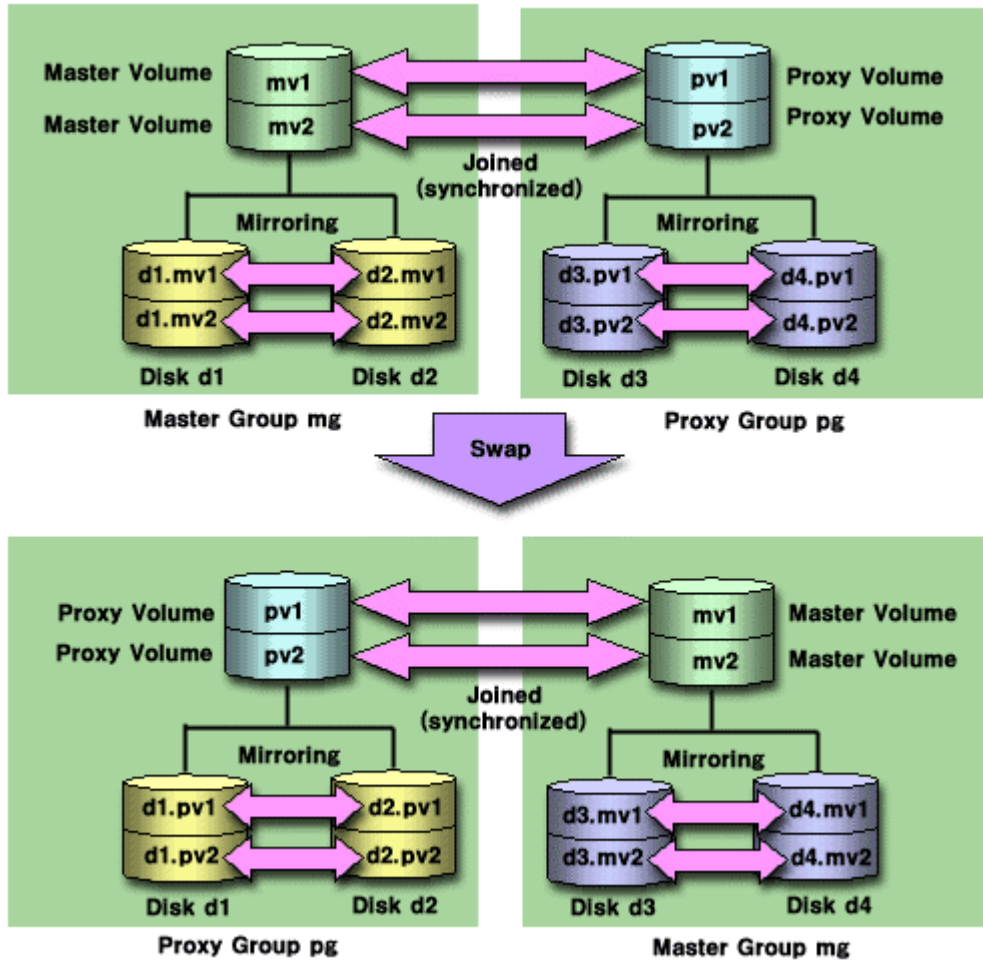


Figure: Swapping Master Group and Proxy Group



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 will manage, and register the user names, at all nodes where GDS will be installed.

3.1.1.1 User Group Types

There are two types of user groups.

wvroot

This is the administrator group for Web-Based Admin View. This user type can perform all operations.

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 "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 responsible users for the entire system should be registered to this group.

3.1.2 Setting up the Client Environment

Client hardware, OS and Web browser supported by GDS Management View is as follows.

PC Client

CPU

Pentium II 233MHz and above

RAM

minimum of 128MB

OS

Windows(R) 98, Windows(R) 98 Second Edition, Windows NT (R), Windows(R) 2000, Windows(R) Me, Windows(R) XP

Web Browser

Netscape Communicator 4.5 to 4.7X, Netscape 6.23, Internet Explorer 5.0 or higher.
(Web browsers need Java Plug-in.)

Java Plug-in

Java(TM) 2 Runtime Environment Standard Edition V1.2.2_008 or higher
Java(TM) 2 Runtime Environment Standard Edition V1.3.1_02 or higher

Display

800 * 600 dots and above, possible of displaying high color (16bit,65,536 colors)

Solaris OE Workstation

CPU

UltraSPARC 200MHz and higher

RAM

minimum of 130MB

OS

Solaris 8 OE, Solaris 9 OE

Web Browser

Netscape Communicator 4.7X
(Web browser needs Java Plug-in.)

Java Plug-in

10.1.2.2,REV=1999.12.06.10.02 or higher.
20.1.2.2,REV=1999.10.14.18.04 or higher.
Java(TM) 2 Runtime Environment Standard Edition V1.3.1_02 or higher



Note HotJava browser support

HotJava browser is not supported. Users must use either Netscape Navigator or Internet Explorer.



Note Web Browser Version

Depending on the Web browser version, it may not operate normally.
Please confirm the version of your Web browser.

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

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

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

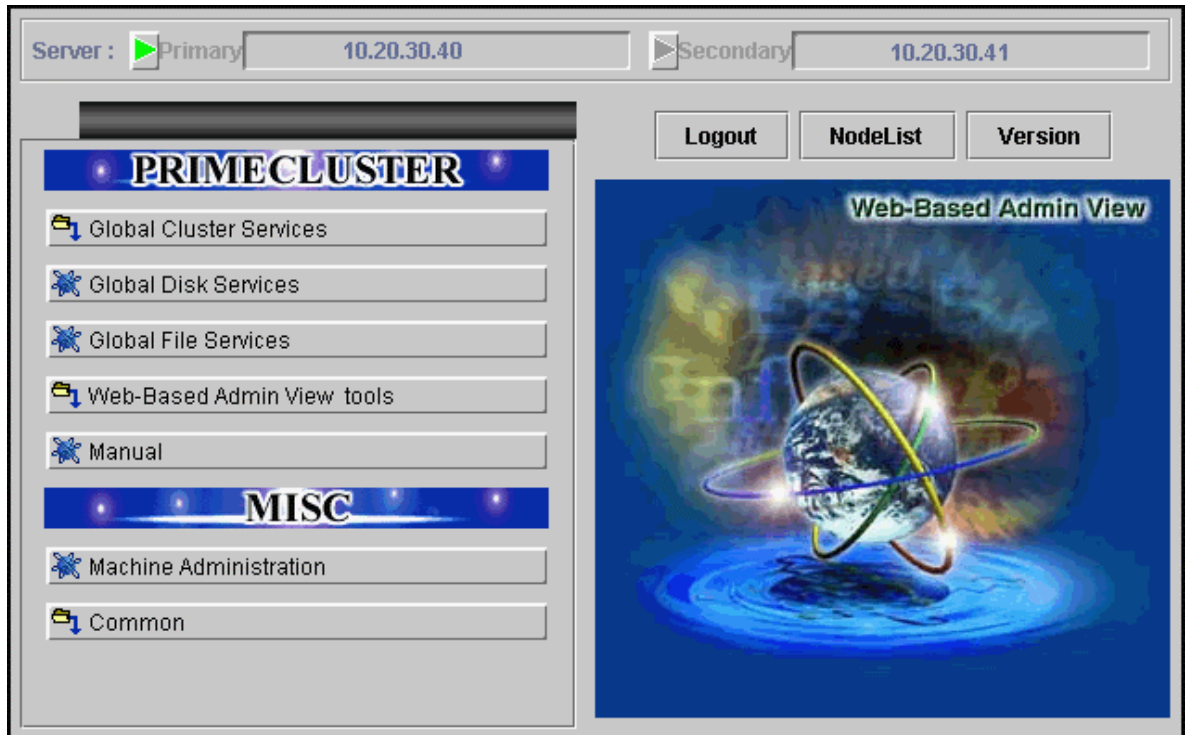
3.2.2 Web-Based Admin View Operation Menu Functions

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

Menu	Outline
Global Disk Services	Start the GDS Management View.

Refer to "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 front. As any errors come into focus, be sure to make the picture viewable.

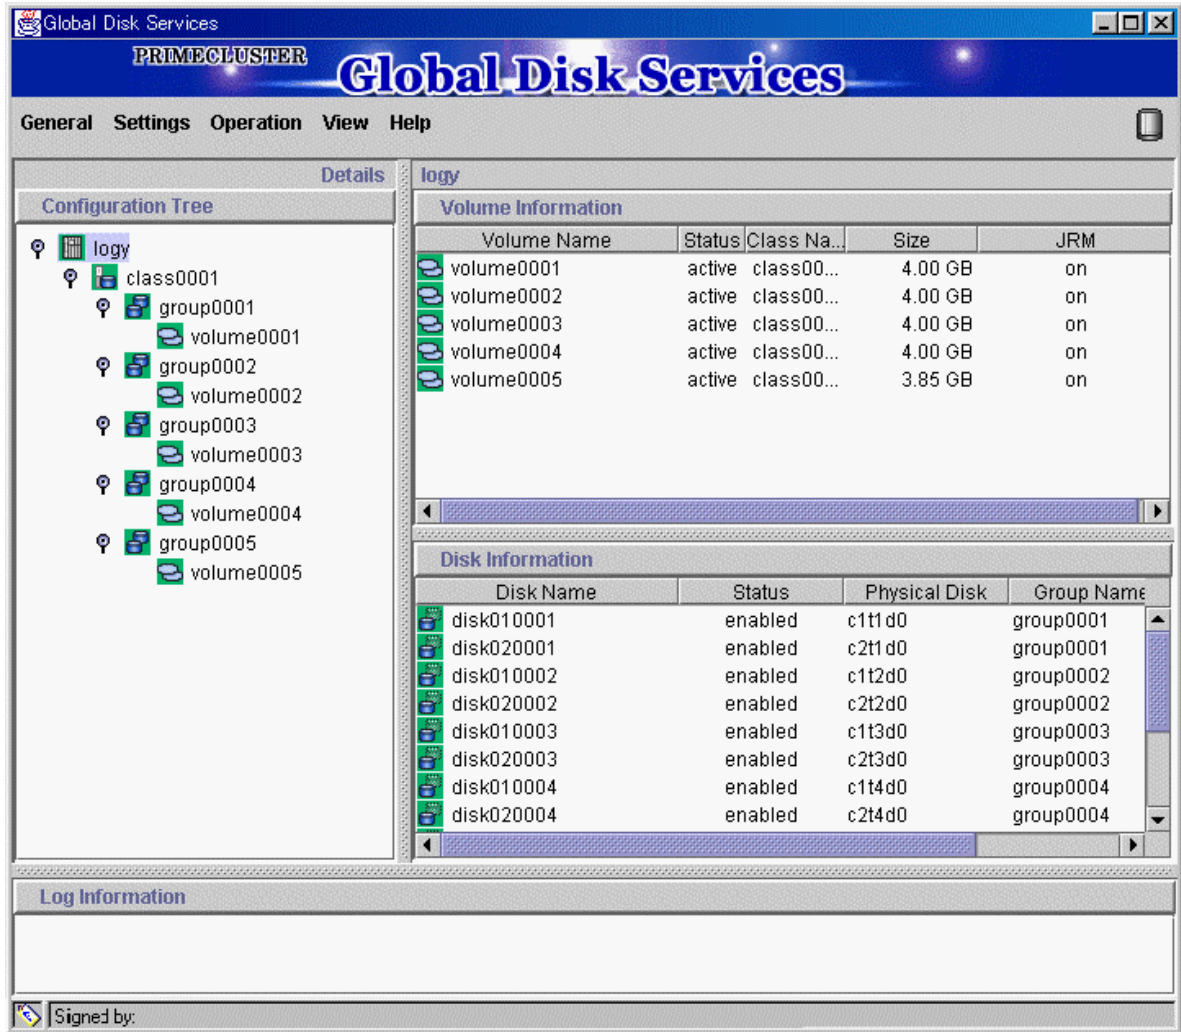
3.2.3 Starting GDS Management View

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 "Chapter5 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



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

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

- Press <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 press 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 for a while. Since it will soon disappear, there is no need to do anything.

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

4.1 Screen Configuration

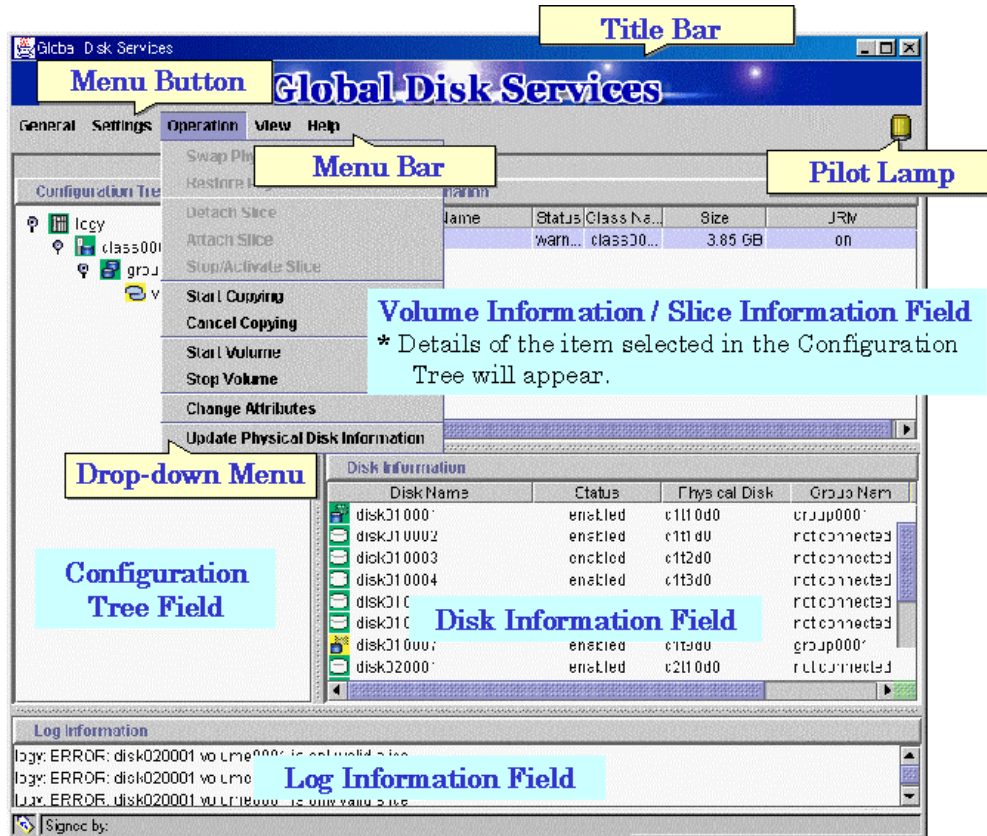
Main Screen

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 or a volume, status confirmation, and disk swap.

Screen configuration of 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.

Volume Information / Slice Information Field

The contents will change depending on which object is selected on the Configuration Tree.

When node, class, or group is selected:

Displays the name of the subordinate volume, its status, and the disk class the volume belongs to.

When volume is selected:

Displays the slice, its status and size, which configures the volume.

Disk Information field

Displays SDX disk name and its status, physical disk (cXtXdX), and mirror disk group.

Log Information Field

Displays error message output by 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.

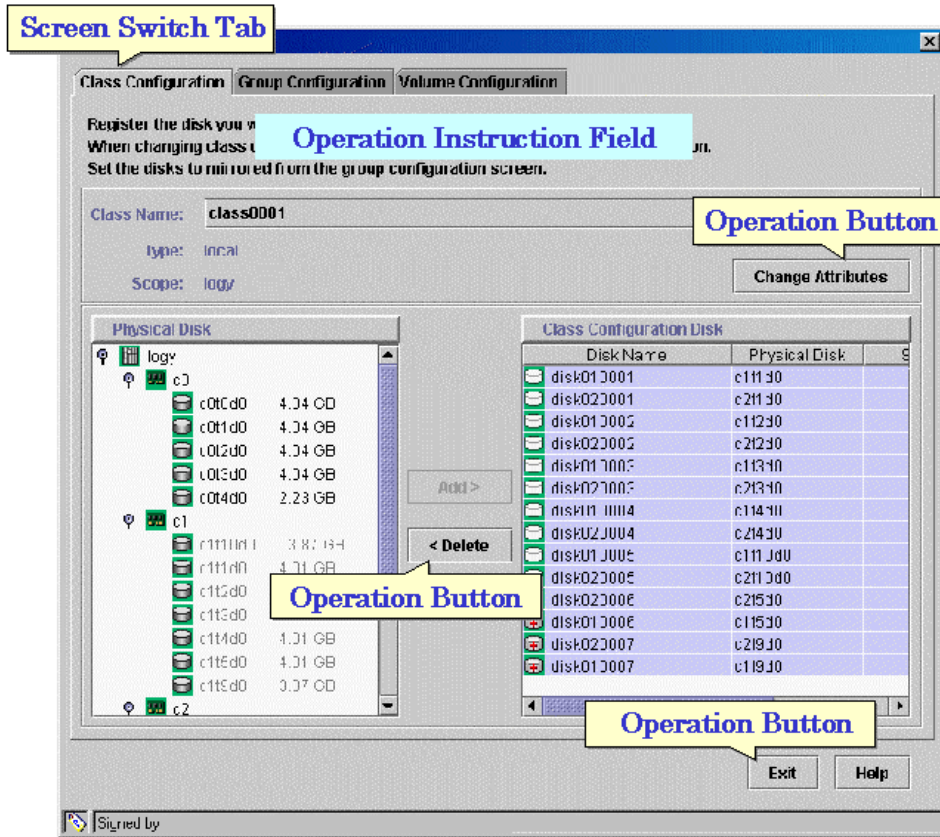
Pilot Lamp	Status	Meaning
 (Gray lit up)	Normal	Normal status.
 (Red blinking)	Abnormal	Critical abnormality (e.g. closed class, invalid volume)
 (Red lit up)	Abnormal	When red blinking warning lamp is single-clicked.
 (Yellow blinking)	Degradation	Volume operating at degradation. (e.g. invalid slice,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

GDS Configuration Settings screen

Select Configuration from <Settings> menu, and "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 configurations, see "5.2.3 Operating from the Settings Menu".

4.2 Menu Configuration and Functions

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

This section explains the menu configuration and functions.

4.2.1 General



Change Monitoring Intervals

Sets the monitoring interval (minutes) of objects.

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



Note

When Changing Monitoring Intervals

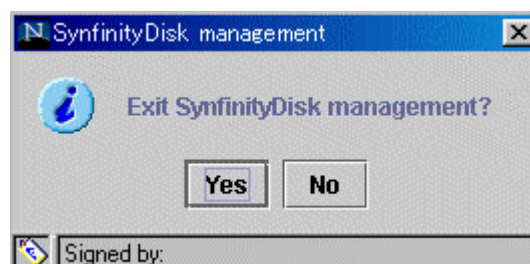
Monitoring interval (minutes) set with Change Monitoring Intervals is not saved.

The content of the settings will be invalid once you exit GDS Management View. After restarting the GDS Management, you must set the monitoring interval (minutes) again.

Exit

Exits GDS Management.

Figure: [General]: [Exit] Screen



4.2.2 Settings

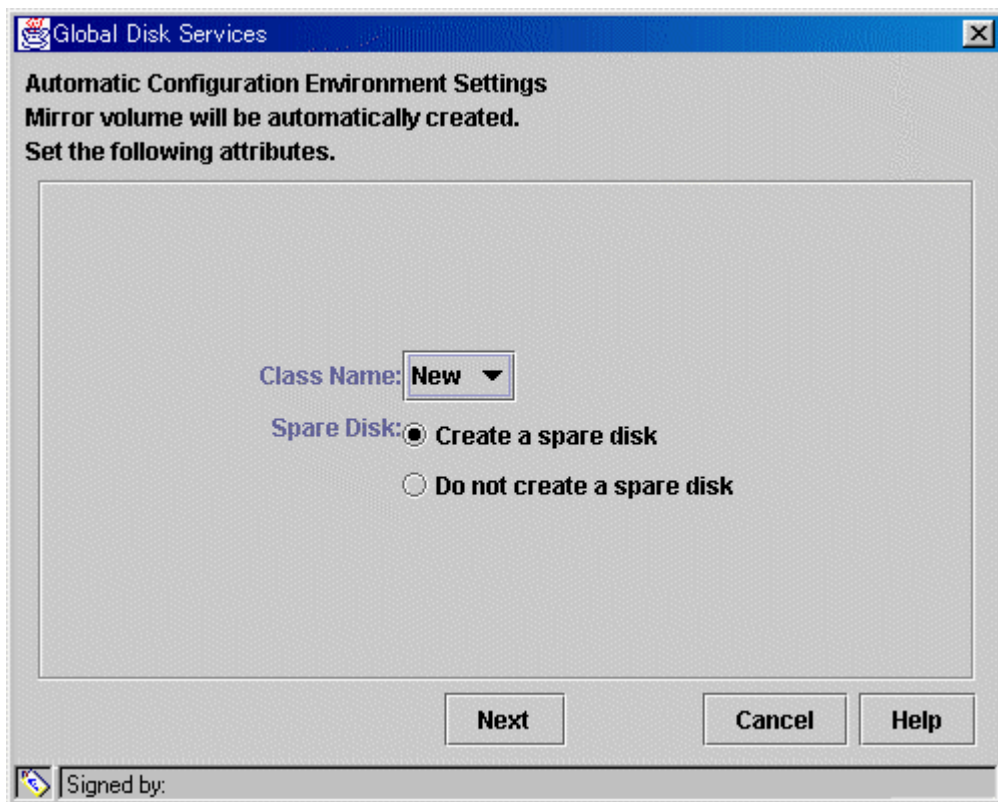


Automatic Configuration

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

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

Figure: [Settings]:[Automatic Configuration] Screen

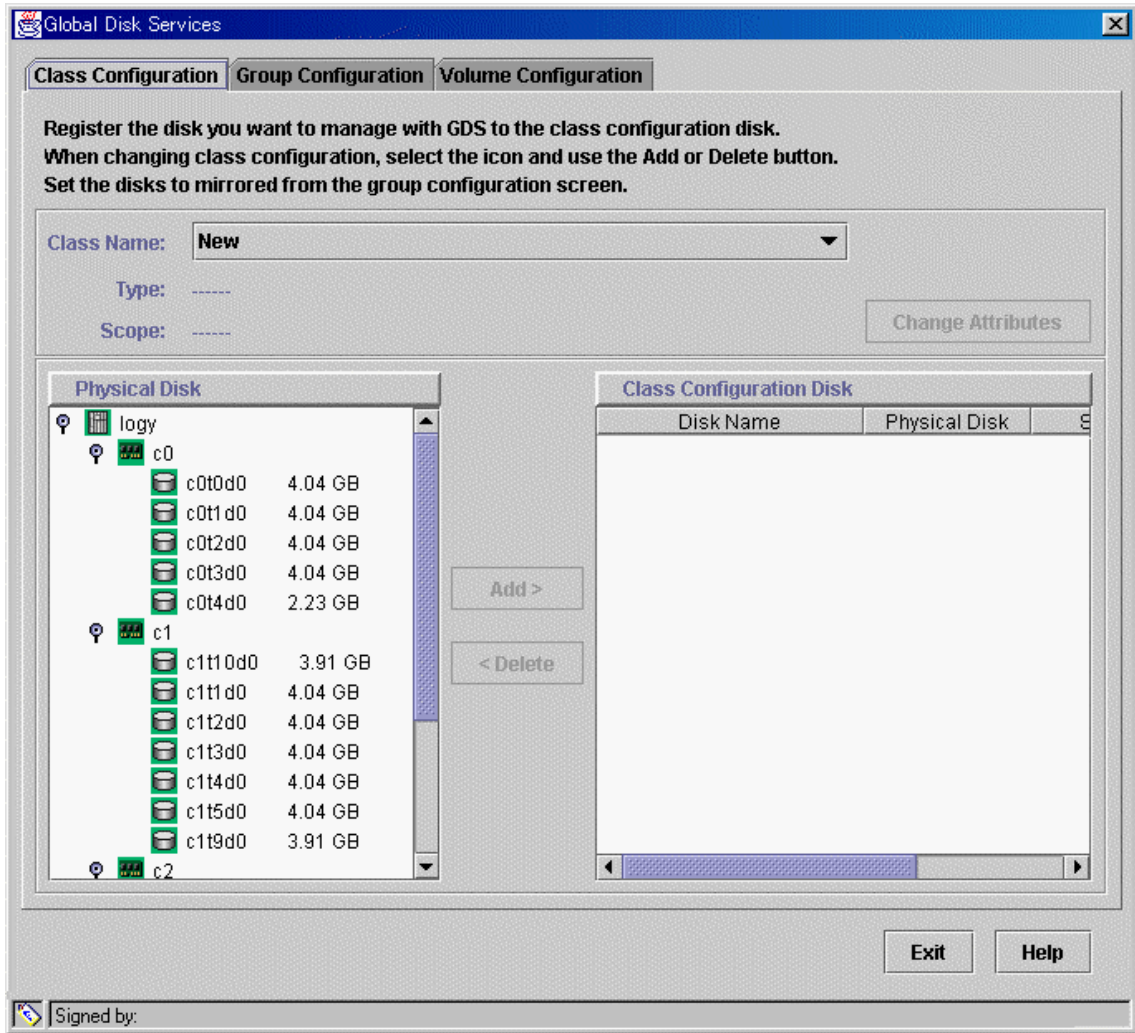


Class Configuration

Sets the class configuration.

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

Figure: [Settings]:[Class Configuration] Screen

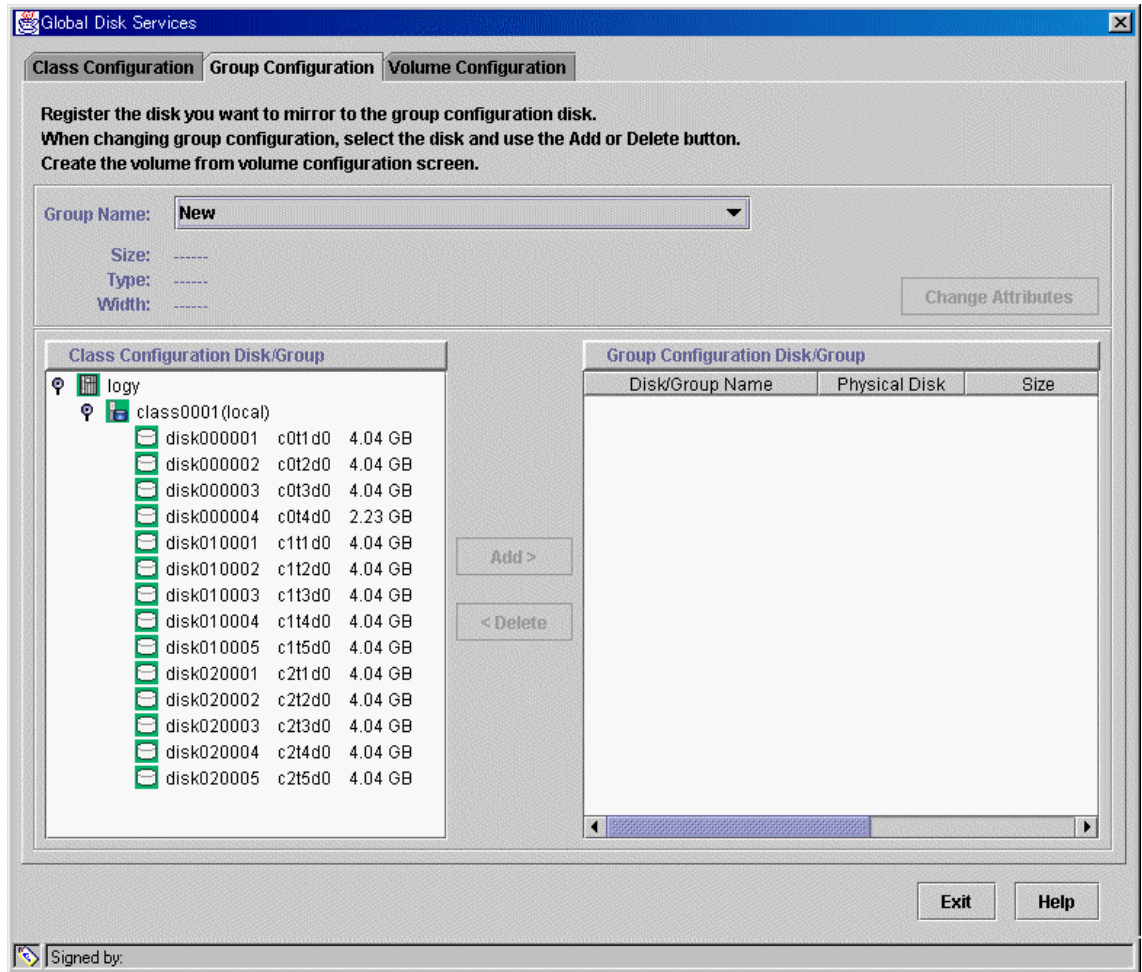


Group Configuration

Sets the group configuration.

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

Figure: [Settings]:[Group Configuration] Screen

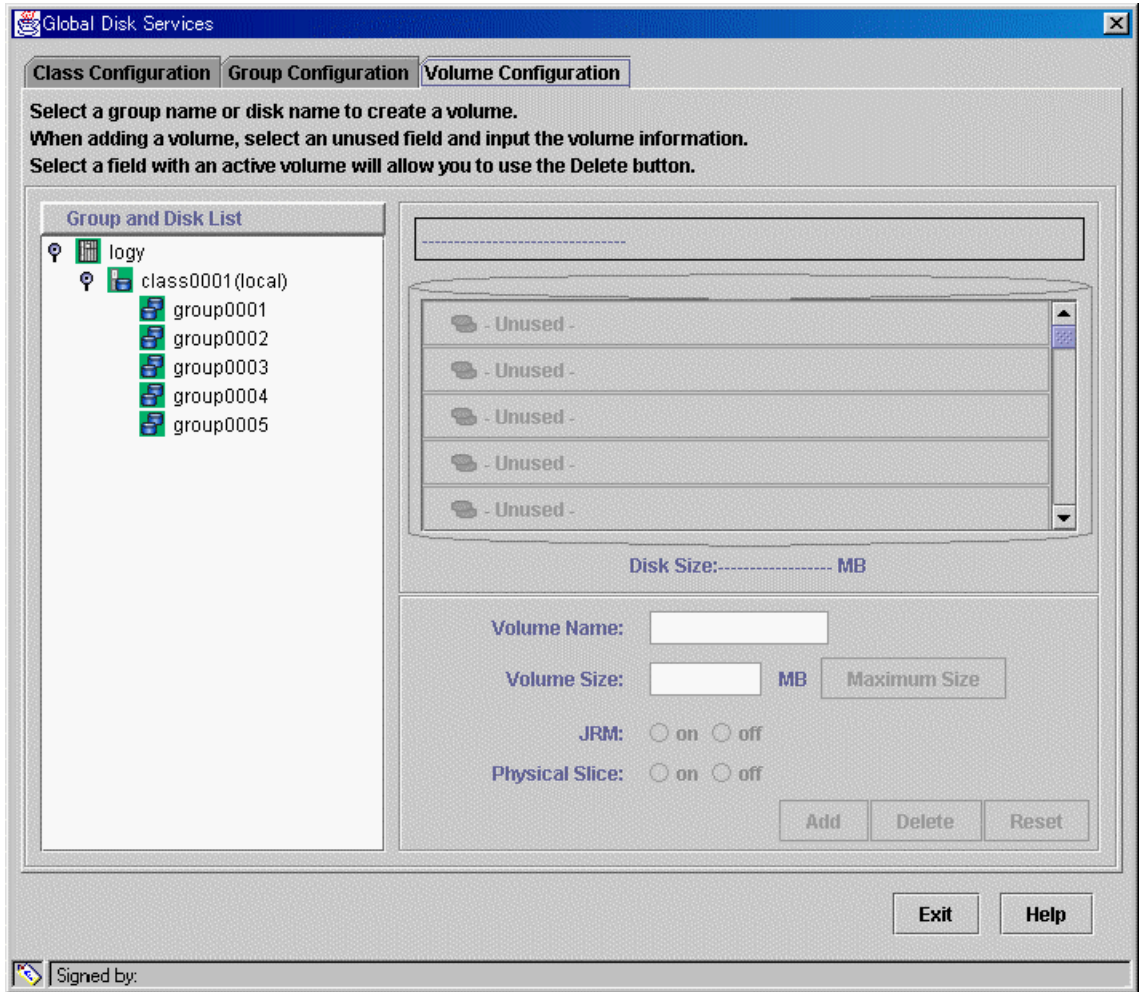


Volume Configuration

Sets the volume configuration.

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

Figure: [Settings]:[Volume Configuration] Screen

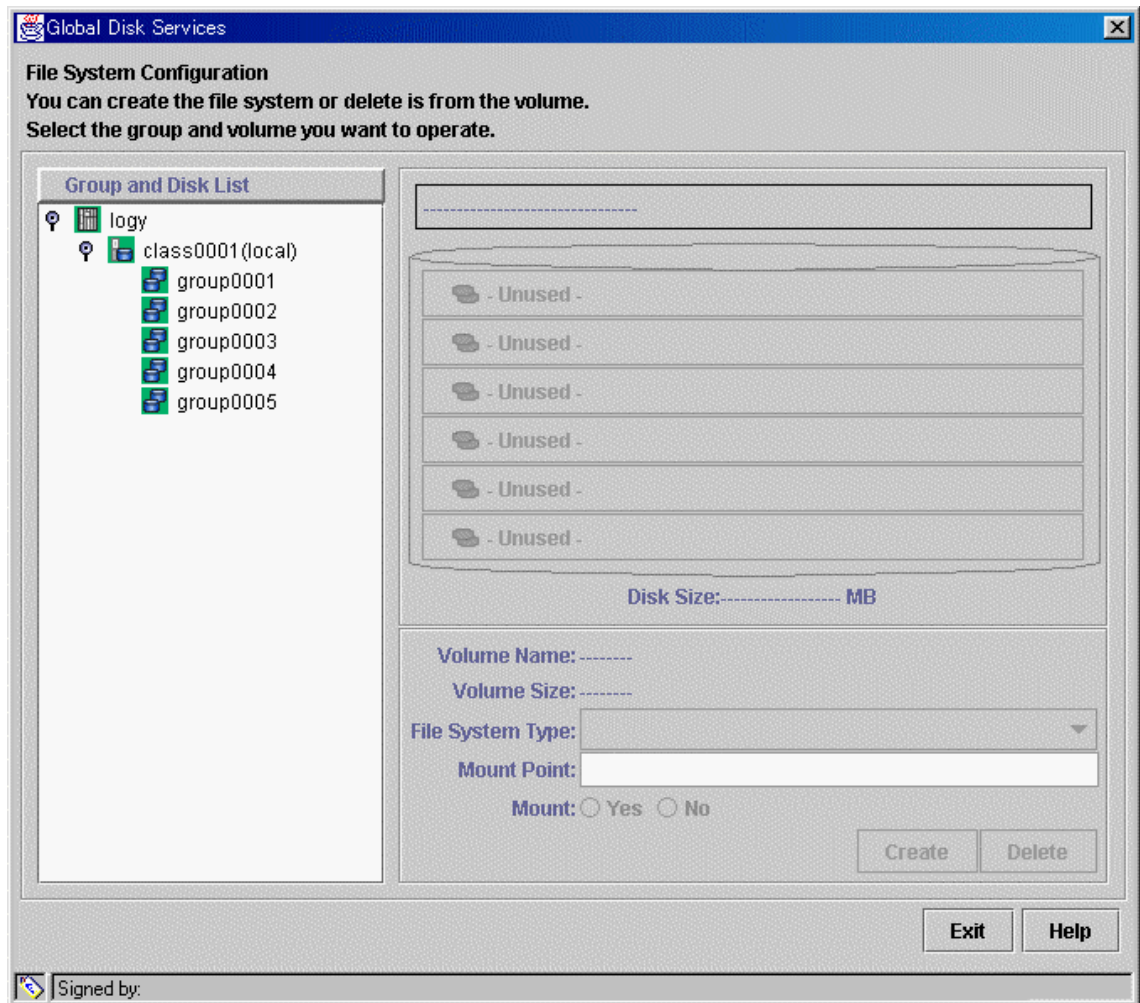


File System Configuration

Sets the file system configuration.

For details, see "5.2.4 File System Configuration".

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

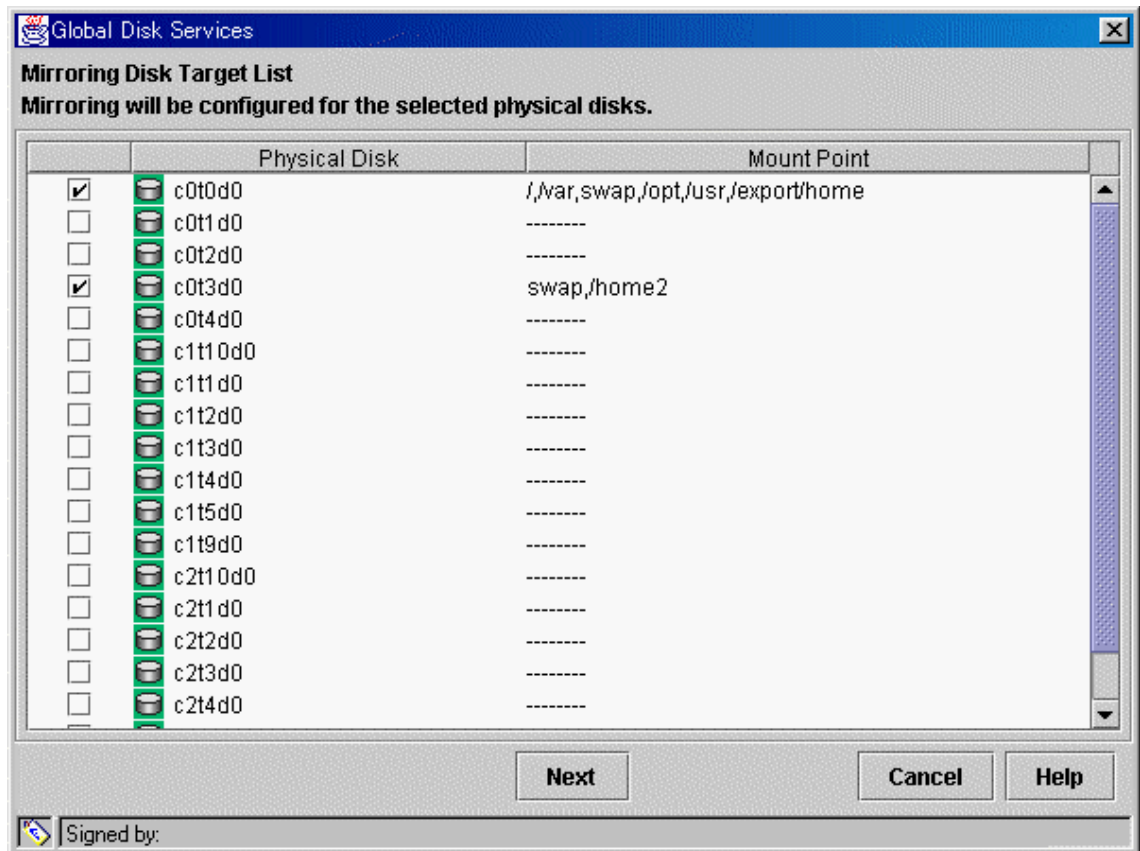


System Disk Settings

Mirrors the system disk.

For details, see "5.2.1 System Disk Settings".

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

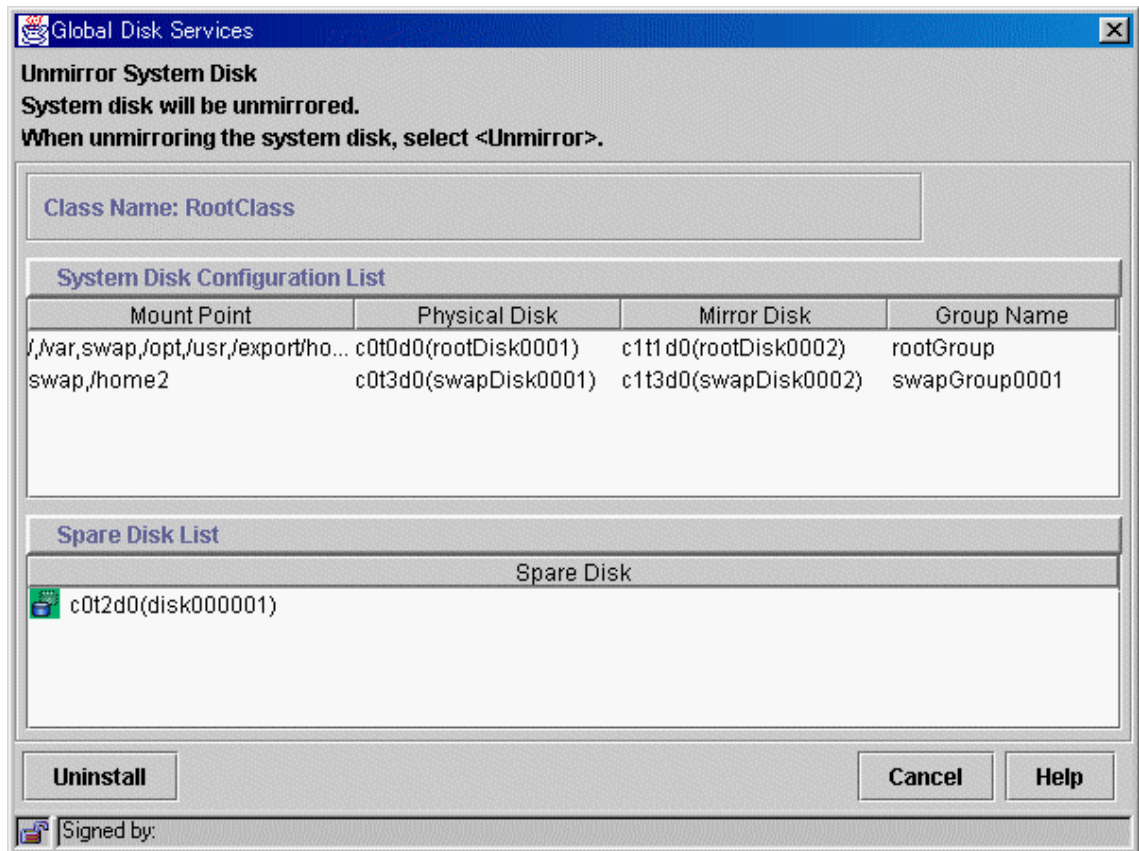


Unmirror System Disk

Unmirrors the system disk mirroring.

For details, see "5.5.5 Unmirroring the System Disk".

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 "5.3.3 In Case of Disk Abnormality".

Restore Physical Disk

Places the swapped physical disk on-line for restoration, after swapping the physical disks.
For details, see "5.3.3 In Case of Disk Abnormality".

Detach Slice

Detaches one of the slices from volume to prepare for backup.
The detached slice will become accessible as a separate volume.
For details, see "5.3.2 Backup".

Attach Slice

Slice detached by [Detach Slice] is attached to volume again.
For details, see "5.3.2 Backup".

Stop/Activate Slice

Stop Slice

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

Activate Slice

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 Backup".

Start Copying

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

For details, see "5.3.4 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 "5.3.4 Copying Operation".

Start Volume

Starts the stopped volume.

Stop Volume

Stops the volume.

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 situations given below.

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

4.2.4 View



Abnormal Object

Only displays objects with abnormalities; not all objects.

Update Now

Usually, GDS Management screen updates the information on object status at an interval specified by [Change Monitoring Intervals] on [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.
The status and the icons of objects are shown below.

1. Node

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

2. Adapter




Icon	Status	Meaning
 (Green)	Normal	-






Note Adapter Icon Status

When an adapter fails, the physical disk under the adapter will be displayed as abnormal, and not the adapter icon. Adapter icon will always be green.



3. Class (local)

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




4. Class (shared)

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




5. Class (root)

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




6. Group (mirror)

Icon	Status	Meaning
 (Green)	Normal	-
 (Yellow)	Abnormal	Abnormality within the group.
 (Red)	Abnormal	Closed class within a class to which the group belongs.




7. Group (stripe)

Icon	Status	Meaning
 (Green)	Normal	-
 (Yellow)	Abnormal	Abnormality within the group.
 (Red)	Abnormal	Closed class within a class to which the group belongs.






8. Group (concat)

Icon	Status	Meaning
 (Green)	Normal	-
 (Yellow)	Abnormal	Abnormality within the group.
 (Red)	Abnormal	Closed class within a class to which the group belongs.





9. Physical Disk

Icon	Status	Meaning
 (Green)	Normal	-
 (Yellow)	Abnormal	Object status of volume or slice is in a status other than active.
 (Red)	Abnormal	Abnormality within the physical disk.






10. Disk connected to the group

Icon	Status	Meaning
 (Green)	enabled	Operation enabled.
 (Yellow)	enabled	If the disk is connected to a mirror group, object status of volume or slice is in a status other than active. If the disk is connected to a stripe group or concatenation group, I/O error has occurred on the disk.
 (Red)	disabled	Operation disabled.
 (Red)	close	Closed class within a class to which the mirror disk belongs.
 (Light Brown)	swap	Disk swap possible.





11. Single Disk

Icon	Status	Meaning
 (Green)	enabled	Operation enabled.
 (Yellow)	enabled	Object status of volume or slice is in a status other than active.
 (Red)	disabled	Operation disabled.
 (Red)	close	Closed class within a class to which the single disk belongs.







12. Spare Disk

Icon	Status	Meaning
 (Green)	enabled	Operation enabled.
 (Yellow)	enabled	Object status of volume or slice is in a status other than active.
 (Red)	disabled	Operation disabled.
 (Red)	close	Closed class within a class to which the spare disk belongs.
 (Light Brown)	swap	Disk swap possible.










13. Unused Disk

Icon	Status	Meaning
 (Green)	enabled	Operation enabled.
 (Red)	disabled	Operation disabled.
 (Red)	close	Closed class within a class to which the unused disk belongs.
 (Light Brown)	swap	Disk swap possible.







14. Volume

Icon	Status	Meaning
 (Green)	active	Activating volume.
 (Yellow)	copy	Synchronization copying
 (Yellow)	warning	Object status of some slices is in a status other than active. (except synchronization copying)
 (Black)	stop	Volume stopped.
 (Red)	invalid	Abnormality of volume.
 (Red)	close	Closed class within a class to which the volume belongs.

15. Slice

Icon	Status	Meaning
 (Green)	active	Activating slice.
 (Black)	stop	Slice stopped.
 (Red)	invalid	Abnormality of slice. (Invalid data)
 (Red)	close	Closed class within a class to which the slice belongs.
 (Blue)	copy	Synchronization 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.

16. Proxy volume

Icon	Status	Meaning
 (Green)	active	Activating proxy volume.
 (Yellow)	copy	Synchronization copying
 (Yellow)	warning	Object status of some slices is in a status other than active. (except synchronization copying)
 (Black)	stop	Proxy volume stopped.
 (Red)	invalid	Abnormality of proxy volume.
 (Red)	close	Closed class within a class to which the proxy volume 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 the abnormality is detected.

Chapter 5 Operation

This chapter explains how to operate GDS.

5.1 Operation Outline

This section explains the GDS settings and operation management.



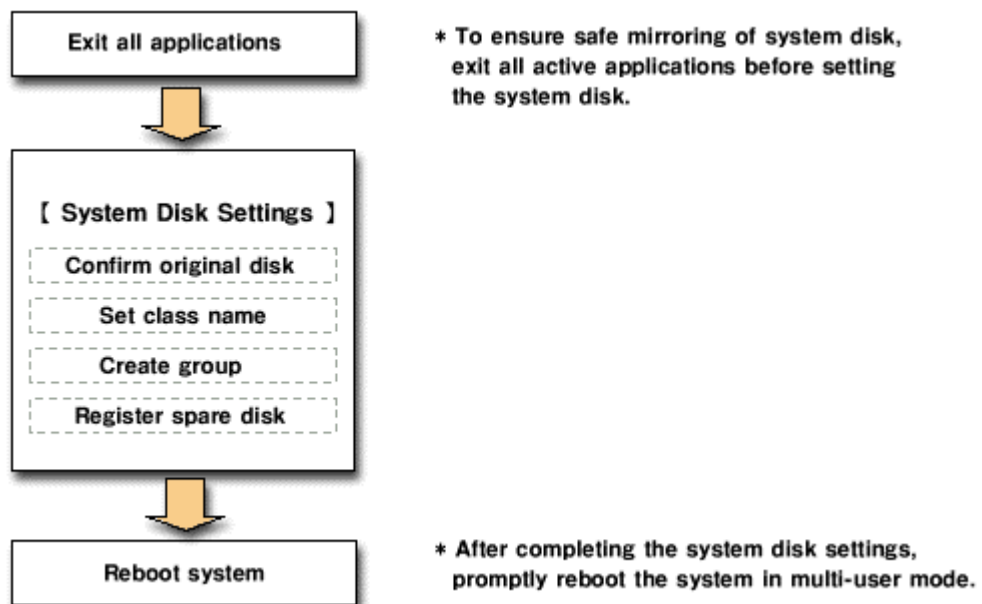
Note In PRIMECLUSTER systems

In order to define the configuration of GDS objects such as classes and volumes, PRIMECLUSTER resources must be registered in advance. For details on registering resources, see "PRIMECLUSTER Cluster Foundation Configuration and Administration".

5.1.1 System Disk Settings

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

Chart: System Disk Settings Operation

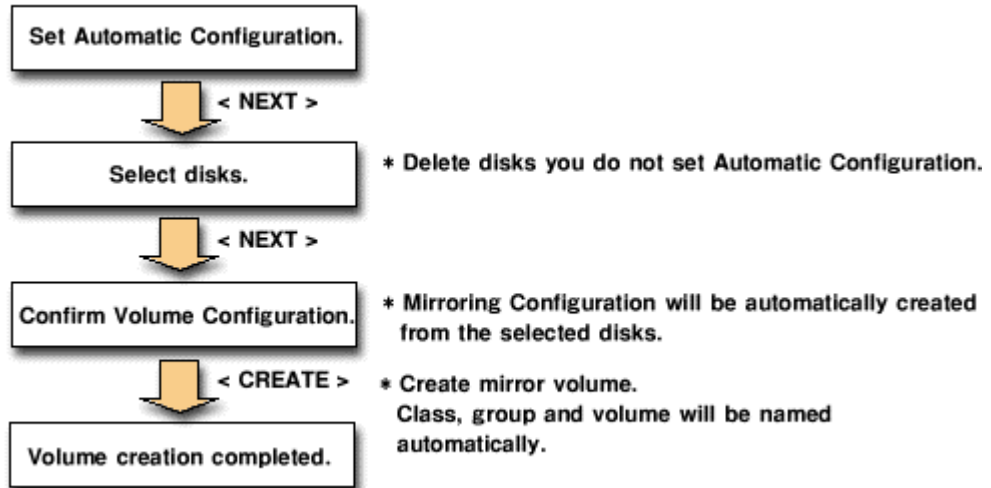


5.1.2 Automatic Configuration

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

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

Chart: Automatic Configuration Wizard Operation



5.1.3 Configuration Settings

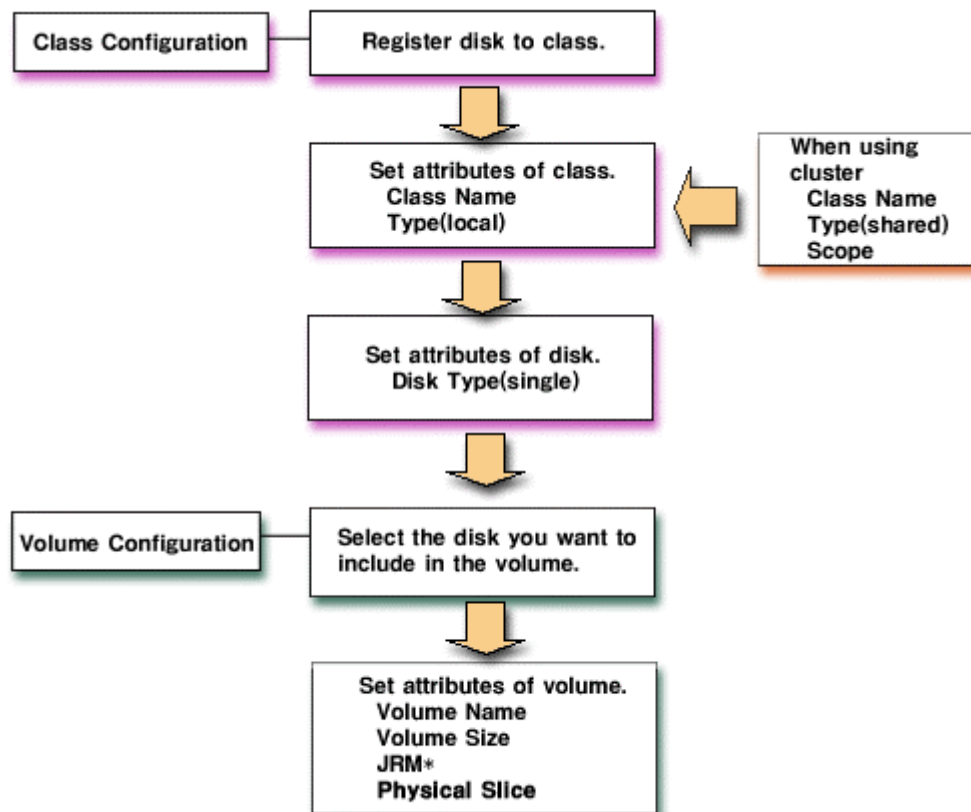
If you are not using [System Disk Settings] or Automatic Configuration Wizard, you must set the configuration using the settings function for class, group and volume. The setting procedures differ depending on the type of volume you are creating.

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

5.1.3.1 Single Volume Configuration Settings

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

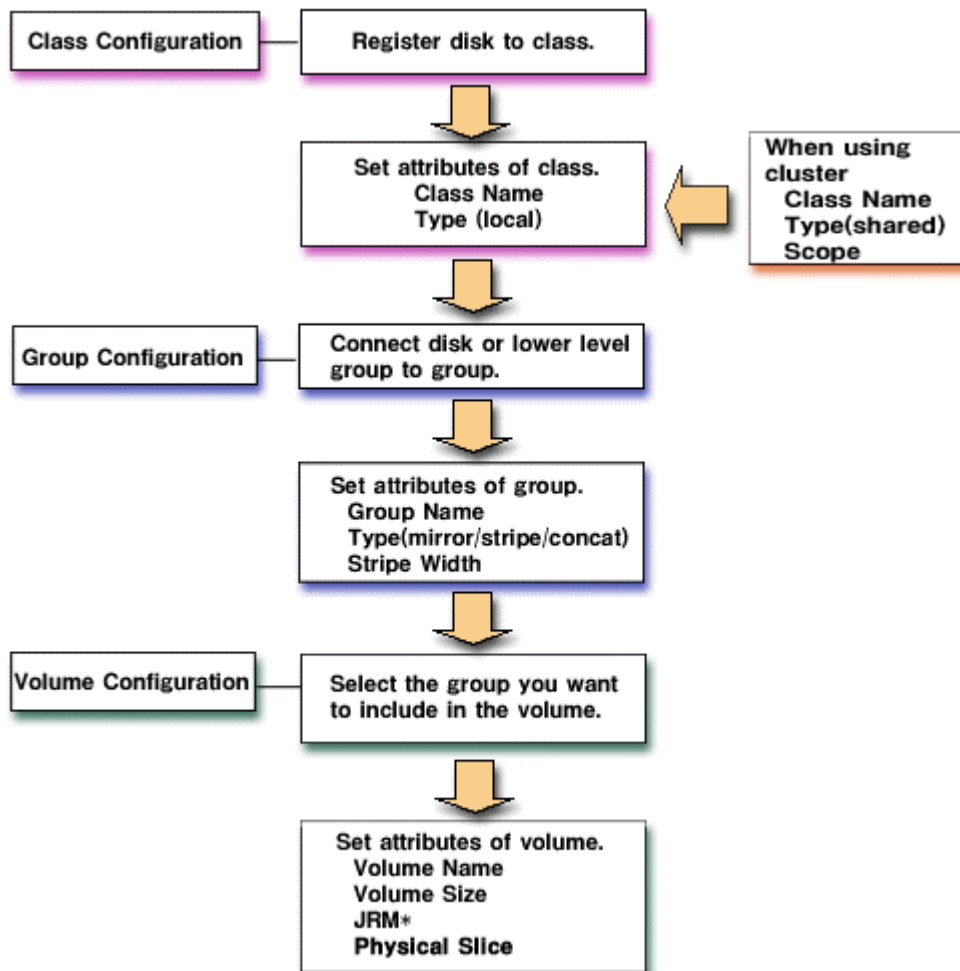
Chart: Single Volume Configuration Settings Operation



5.1.3.2 Other Volume Configuration Settings

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

Chart: Volume Configuration Settings Operation

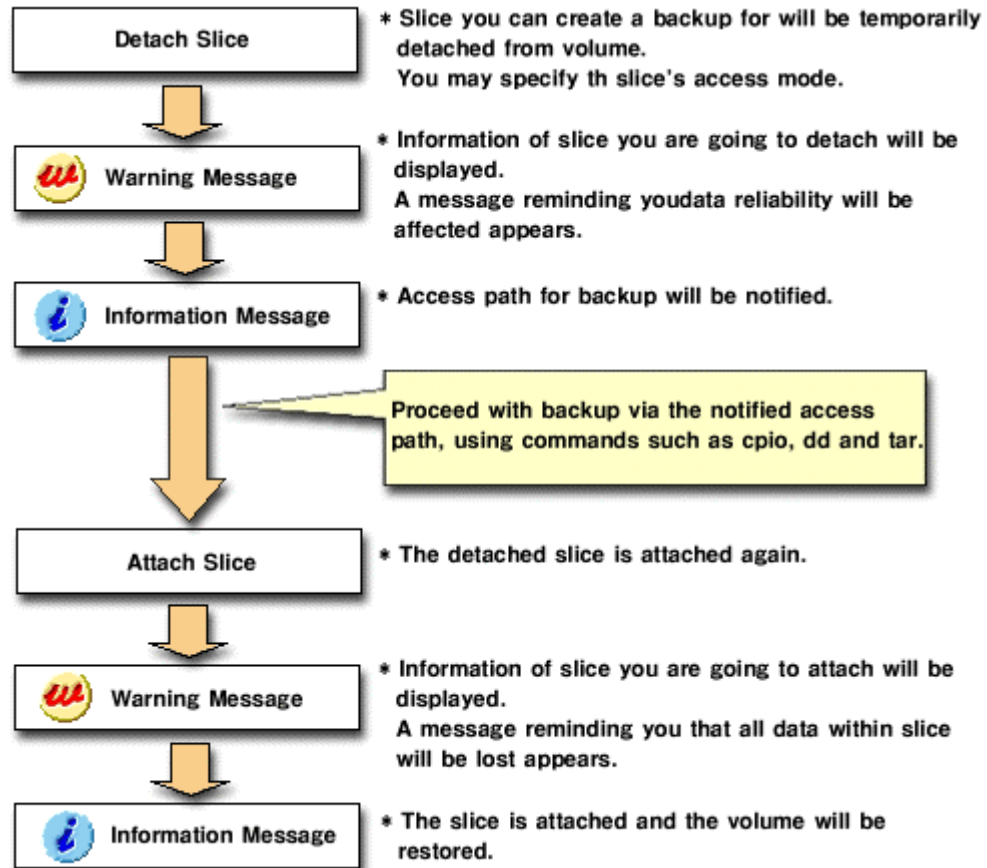


* JRM : Just Resynchronization Mechanism

5.1.4 Backup

The following operations are required in case mirror volume is backed up using the snapshot created by detaching slice.

Chart: Backup Operation



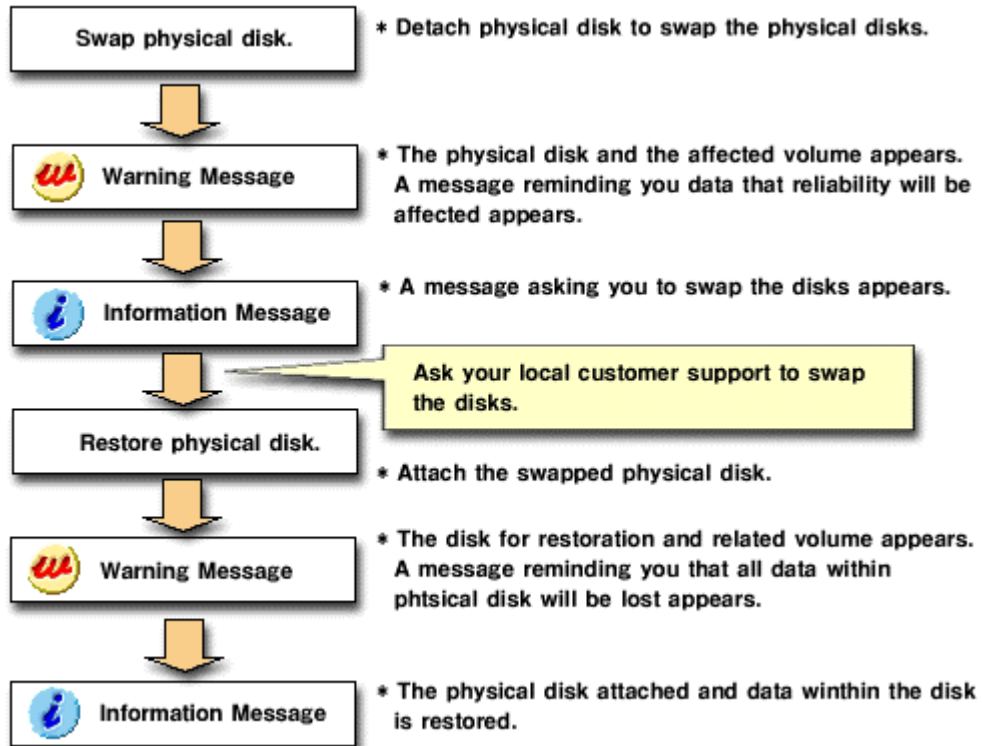
Note Requirement for Applying a Snapshot by Detaching a Slice

A slice can be detached only from a mirror volume with a physical slice. So if a disk is not connected directly to a mirror group, creation of a snapshot by detaching a slice is not available.

5.1.5 In Case of Disk Abnormality

Procedures for swapping the disks if a disk abnormality occurs is shown below.

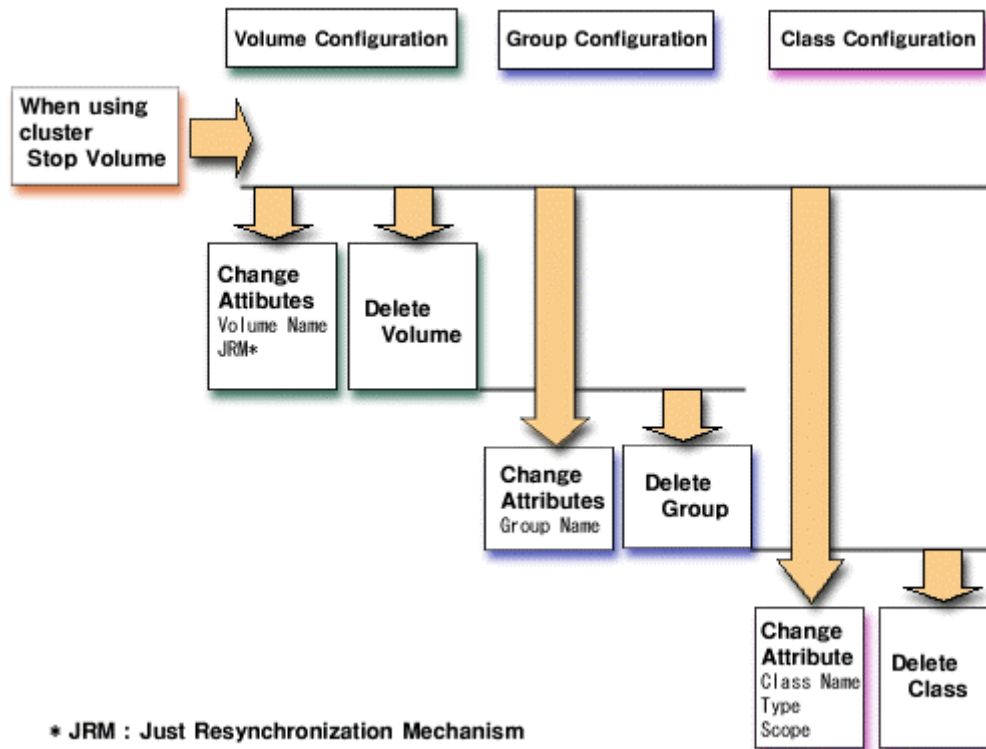
Chart: Disk Swap and Restoration Operation



5.1.6 Configuration Changes

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


Chart: Configuration Changing Operation



5.1.7 Unmirroring the System Disk

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

Chart: 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.

5.1.8 Operation using GDS Management View

GDS Management View supports the following operations. For operations not supported by GDS Management View, you must execute the commands.

Class operation

Operation	Command Operation	GDS Management View Operation
Create	# sdxdisk -M	For details, see "5.2.3.1 Class Configuration" in the "5.2.3 Operating from the Settings Menu".
Remove	# sdxdisk -R or # sdxclass -R	For details, see "5.5.4 Removing a Class" in the "5.5 Removals"/
Change Attribute	# sdxattr -C	For details, see "5.4.1 Class Configuration" in the "5.4 Changes".

Group operation

Operation	Command Operation	GDS Management View Operation
Connect	# sdxdisk -C or # sdxgroup -C	For details, see "5.2.3.3 Group Configuration" in the "5.2.3 Operating from the Settings Menu".
Remove	# sdxdisk -D or # sdxgroup -D	For details, see "5.5.3 Removing a Group" in the "5.5 Removals".
Change Attribute	# sdxattr -G	For details, see "5.4.2 Group Configuration" in the "5.4 Changes".

Volume operation

Operation	Command Operation	GDS Management View Operation
Create	# sdxvolume -M	For details, see "5.2.3.4 Volume Configuration" in the "5.2.3 Operating from the Settings Menu".
Remove	# sdxvolume -R	For details, see "5.5.2 Removing a Volume" in the "5.5 Removals".
Start	# sdxvolume -N	For details, see "Start Volume" in the "4.2.3 Operation".
Stop	# sdxvolume -F	For details, see "Stop Volume" in the "4.2.3 Operation".
Change Attribute	# sdxattr -V	For details, see "5.4.3 Volume Configuration" in the "5.4 Changes".
Change Attribute (JRM)	# sdxattr -V ... -a jrm=[on off]	For details, see "5.4.3 Volume Configuration" in the "5.4 Changes".
Change Attribute (Physical Slice)	# sdxattr -V ... -a pslice=[on off]	For details, see "5.4.3 Volume Configuration" in the "5.4 Changes".
Change Attribute (Lock Volume)	# sdxattr -V ... -a lock=[on off]	Not supported
Change Attribute (Access Mode)	# sdxattr -V ... -a mode=[rw ro]	Not supported
Change Attribute (Physical slice type)	# sdxattr -V ... -a pslice= [on off]	For details, see "5.4.3 Volume Configuration" in the "5.4 Changes".
Start Copying	# sdxcopy -B	For details, see "5.3.4 Copying Operation" in the "5.3 Operation in Use".
Cancel Copying	# sdxcopy -C	For details, see "5.3.4 Copying Operation" in the "5.3 Operation in Use".
Restore	# sdxfix -V	Not supported

Slice operation

Operation	Command Operation	GDS Management View Operation
Detach	# sdxslice -M	For details, see "5.3.2 Backup" in the "5.3 Operation in Use".
Attach	# sdxslice -R	For details, see "5.3.2 Backup" in the "5.3 Operation in Use".
Activate	# sdxslice -N	For details, see "5.3.2 Backup" in the "5.3 Operation in Use".
Stop	# sdxslice -F	For details, see "5.3.2 Backup" in the "5.3 Operation in Use".
Take over	# sdxslice -T	Not supported
Change Attribute	# sdxattr -S	Not supported

Disk operation

Operation	Command Operation	GDS Management View Operation
Swap	# sdxswap -O	For details, see "5.3.3 In Case of Disk Abnormality" in the "5.3 Operation in Use".
Restore	# sdxswap -I	For details, see "5.3.3 In Case of Disk Abnormality" in the "5.3 Operation in Use".
Change Attribute	# sdxattr -D	For details, see "5.4.1 Class Configuration" in the "5.4 Changes".

Proxy operation

Operation	Command Operation	GDS Management View Operation
Join	# sdxproxy Join	Not supported
Break	# sdxproxy Break	Not supported
Part	# sdxproxy Part	Not supported
Rejoin	# sdxproxy Rejoin	Not supported
Restore	# sdxproxy RejoinRestore	Not supported
Swap	# sdxproxy Swap	Not supported

Operation to display the configuration information

Operation	Command Operation	GDS Management View Operation
View	# sdxinfo	For details, see "5.3.1 Confirming Configuration and Monitoring Status" in the "5.3 Operation in Use".

Other operations

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

5.2 Settings

This section explains the GDS setting procedure.

5.2.1 System Disk Settings

System Disk Settings allows you to mirror system disks.

System disk cannot be mirrored under the following conditions. Change the disk configuration before performing the system disk settings.

- When the disk does not have an unused (size=0) slice number.
- When the disk does not have sufficient space area or swap area.

For details on necessary space area (or swap area), refer to "A.2.6 Disk Size".



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

System disk settings for GDS Management View recognizes disk with /, /var, /usr or swap area as the system disk.

When a system disk is recognized as the target of mirroring, the disk will automatically become the original disk, and cannot be excluded from the mirroring configuration. When you have the swap area secured on a separate disk, note that you need a mirror disk.

Also, physical disk that was not recognized as a system disk cannot be mirrored using the [System Disk Settings].



Note **For Safe Mirroring of System Disk**

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

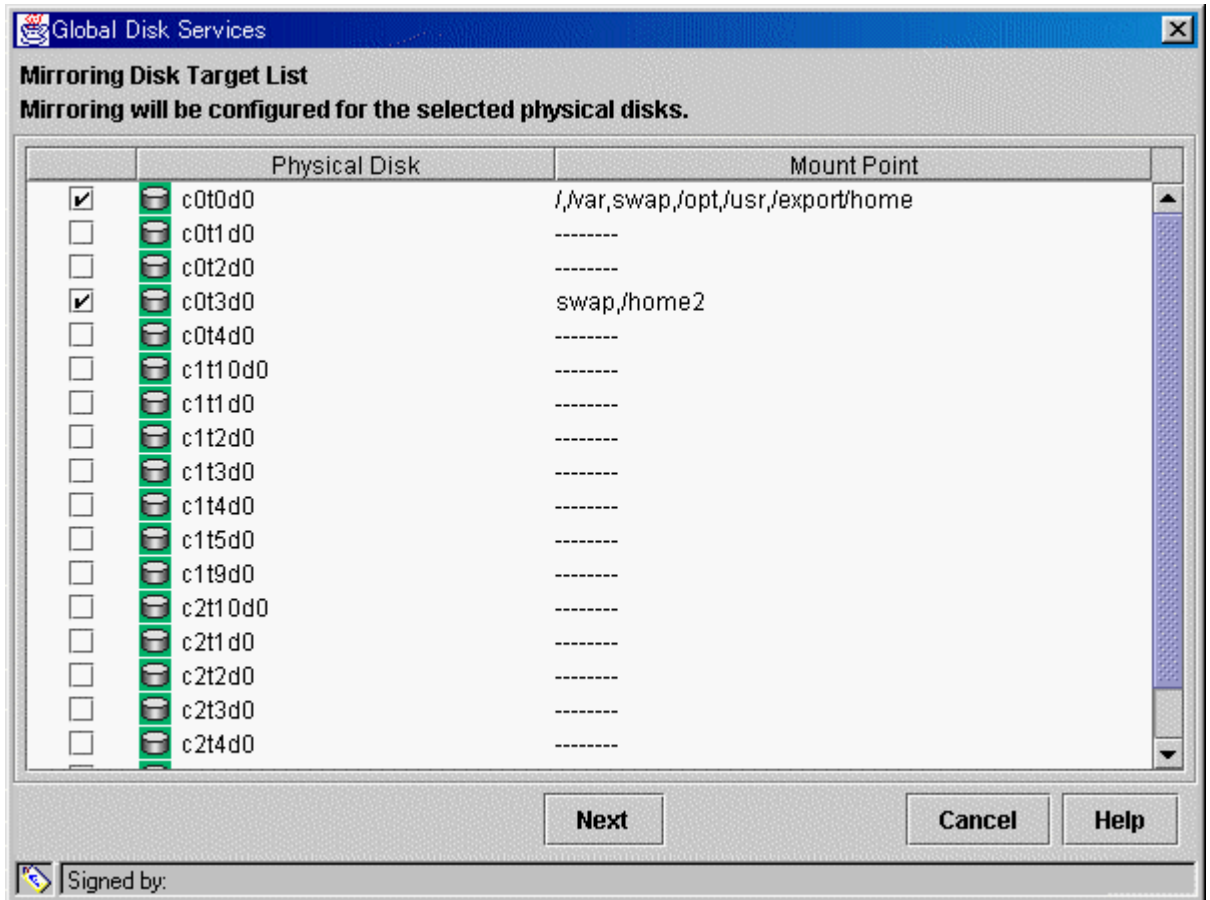
During the mirroring process, there may be a considerable degradation of application response time. After completing the system disk settings, promptly reboot the system in multi-user mode.

How to mirror a system disk is explained.

On the [Settings] menu, click [System Disk Settings].

1. Confirming original disk

Figure: Mirroring Disk Target List

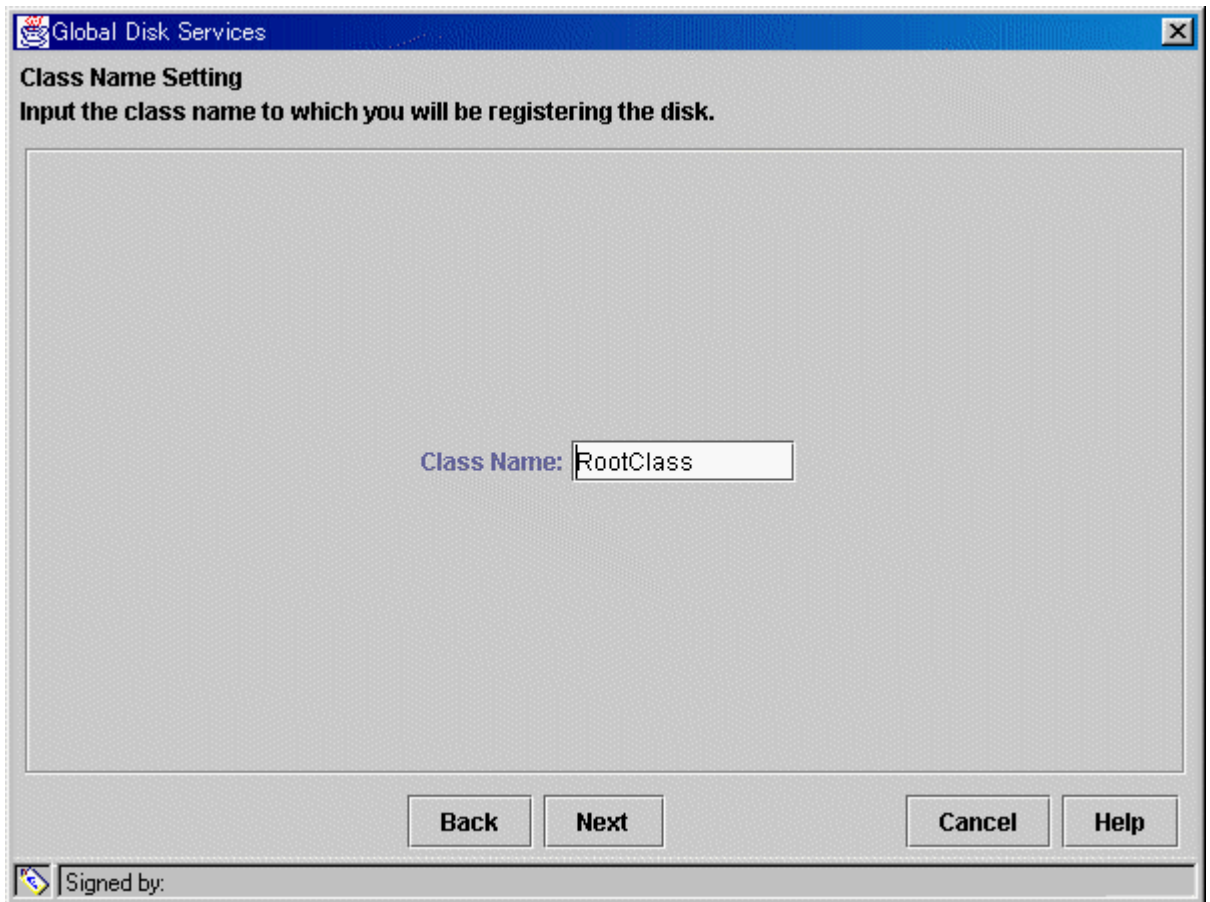


In the "Physical Disk" field, system disks that will be used as original disks will be displayed with a check mark. Disks with mount information set for /etc/vfstab will have the mount point displayed in the "Mount Point" field.

Original disks with check marks in the "Physical Disk" field will be registered with root class as "keep" disk. If you do not want to proceed with mirroring, press <Cancel>. If you want to proceed with mirroring, press <Next>.

2. Creating root class

Figure: Class Name Setting



Type the class name of the root class.



Note Inputting Class Name

Class name will be used for the device path name.

/dev/sfdsk/class name/[r]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 system, class name of the root class should be set different for each node.



See

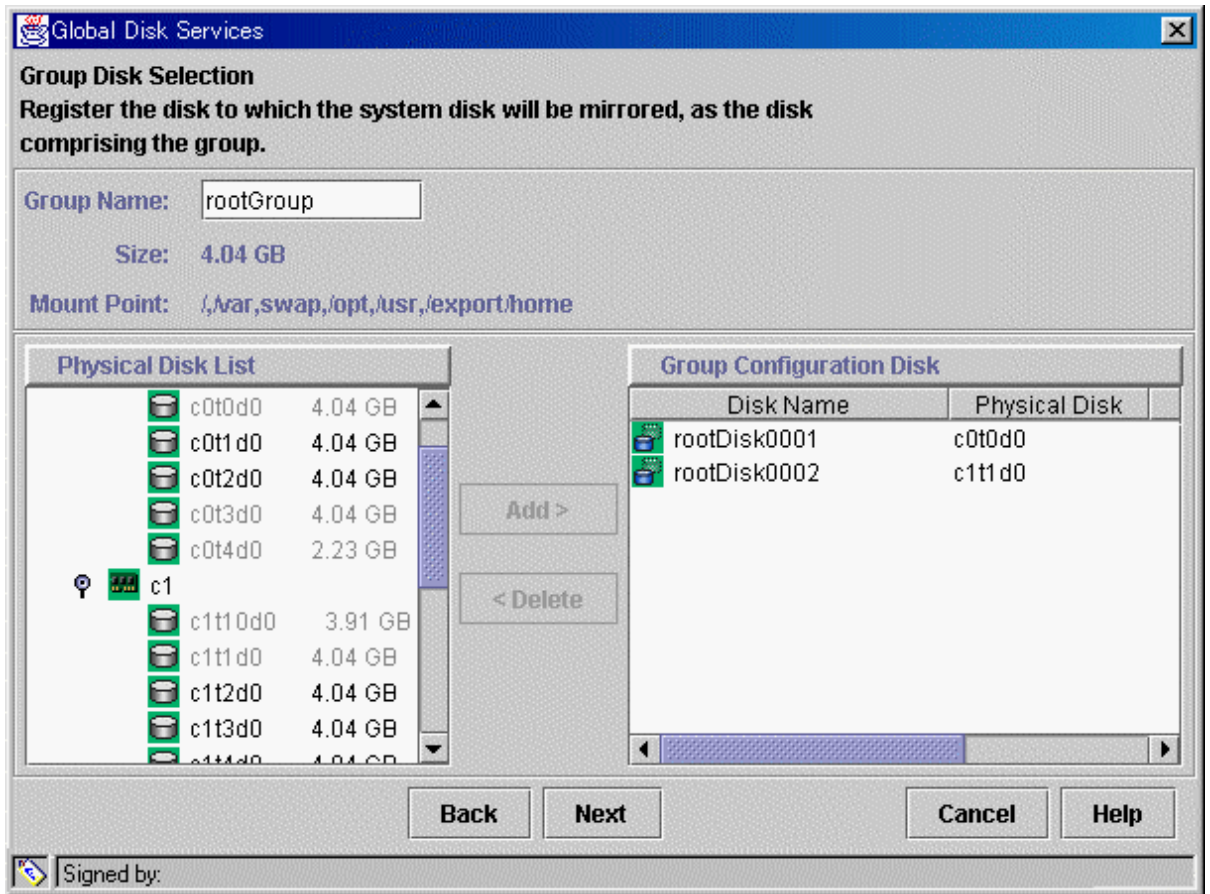
For information on assigning a class name, see "A.1.1 Object Name."

Press <Next> to continue.

3. Creating group

Create the group by selecting the mirror disks. When more than one disk is recognized as the system disk, groups should be created for each 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> list, the original disk you selected will be displayed.

You cannot remove the original disk from the < Group Configuration Disk> list .

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

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

You must select more than one mirror disk.

Double-click the "disk name" in the " Group Configuration Disk " field, and change the name.

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

Once you finish creating all mirror disk 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 press <Add>.

After finishing the registration, press <Next>. When you do not need to register a spare disk, no setting is necessary. Just press <Next>.

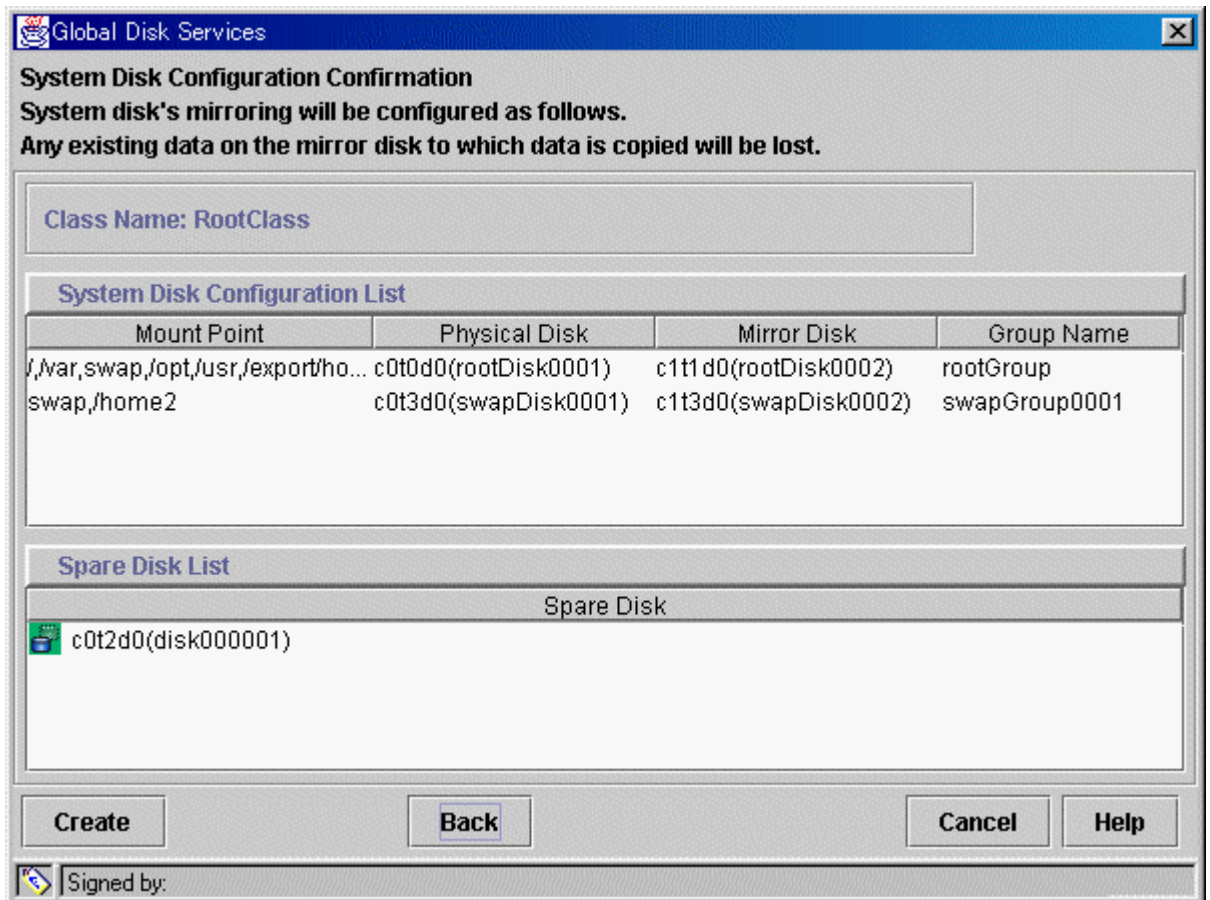


Point Spare Disk Size

The hot spare function will not operate when there is not sufficient space on spare disk to copy the configuration of volume in mirror group. Define the largest-size 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/vfstab, mount point will be displayed in the "Mount Point" field. Press <Create> to continue.



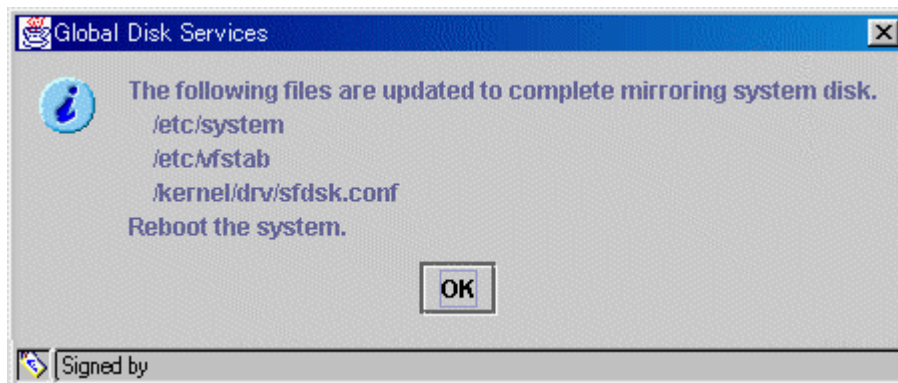
Information Automatically-generated Volume Names

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

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

6. Completing system disk configuration

Figure: Setting System Disk Mirroring Complete



Confirm that the system disk configuration is complete, and press <OK>. After completing the system disk configuration, promptly reboot the system.



Note

When does the mirroring of system disk become valid?

System disk mirroring will take place when the system is activated in multi-user mode.



Note

JRM and System Disk Settings

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



Note

GFS Partition in System Disk

When the system disk contains a partition used by a GFS Local File System, you must make the following modifications on the settings for the relevant partition prior to rebooting the system after completing the system disk configuration.

- Changing the partition information settings
For details on the method of changing the settings, refer to "PRIMECLUSTER Global File Services Configuration and Administration Guide" or the manual page of the `sfxadm(1M)` command.
- Modifying the `/etc/vfstab` file
You must change the device special file name from that for the physical slice to that for the logical volume.

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

5.2.2 Automatic Configuration Wizard

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

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

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

Automatic Configuration Wizard allows you to create a volume by simply pressing <Next>. In order to ensure an environment which allows an application to continuously access data at the time of I/O adapter, controller or cable failure, Automatic Configuration Wizard mirrors disks connected to different I/O adapters.

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

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



Note

When Automatic Configuration Wizard Cannot Be Used

When all target disks are connected to one IO adapter, automatic configuration mode cannot be used.

Also, Automatic Configuration Mode does not support following disk drivers:

mphd, mplb, emcpower, emcp

Users must make the necessary settings from the configuration screen.



Note

Mirroring Expansion Units with Add-On Expansion Mechanism

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



Note

Shared Disk Mirroring

Shared disks in cluster system cannot be mirrored using the automatic configuration feature. Set the configuration from the [Settings] menu.



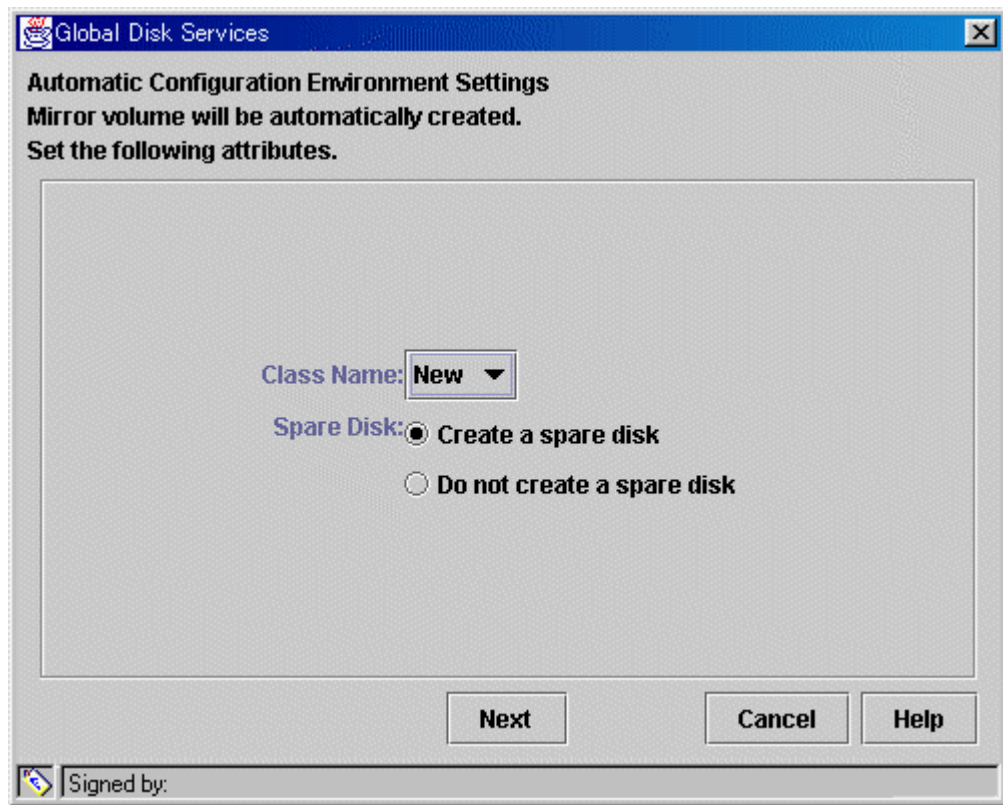
Note

System Disk Mirroring

System disk cannot be configured automatically. Instead, click [System Disk Settings] on the [Settings] menu.

To start Automatic Configuration Wizard, click [Automatic Configuration] on the [Settings] menu.

Figure: Automatic Configuration Environment Settings



Selecting the Configuration Environment for Mirroring

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

a. Selecting the "Class Name"

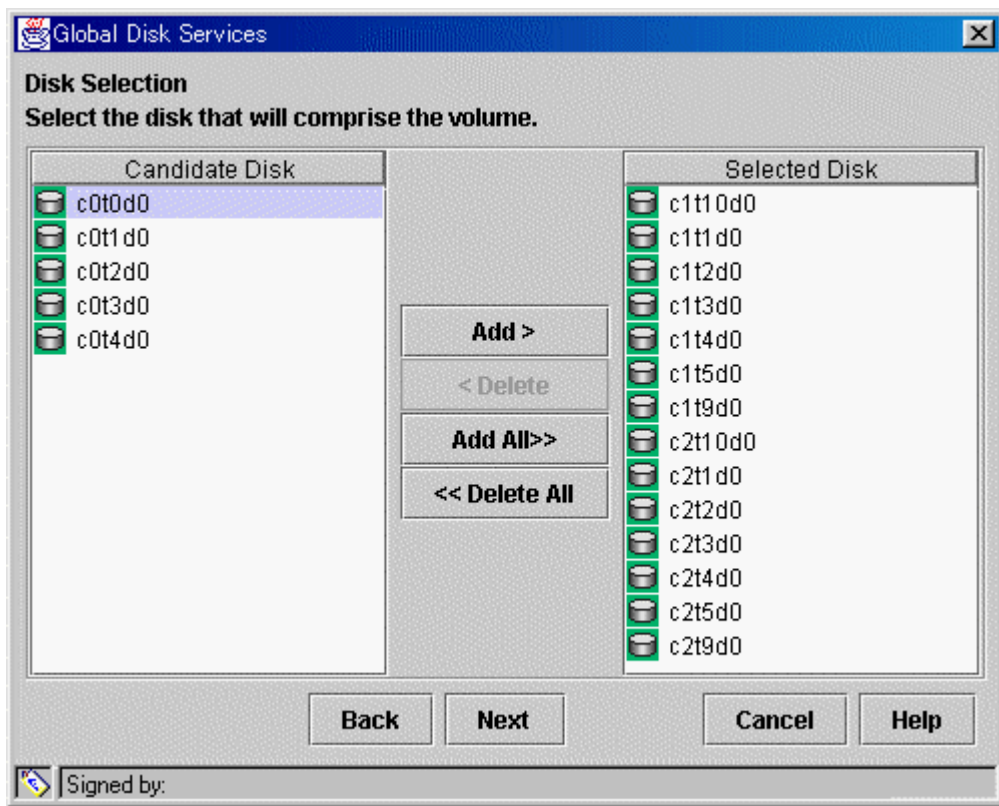
If you pressed <New> in the "Class Name" list, a new class will be created for the volume. If you are creating a volume within an existing class, select the appropriate class from the list. If no classes have been created, the only option available on the list will be <New>.

b. Selecting the "Spare Disk"

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

Press <Next>, and the message below appears.

Figure: Disk Selection



Selecting Disks to Mirror

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

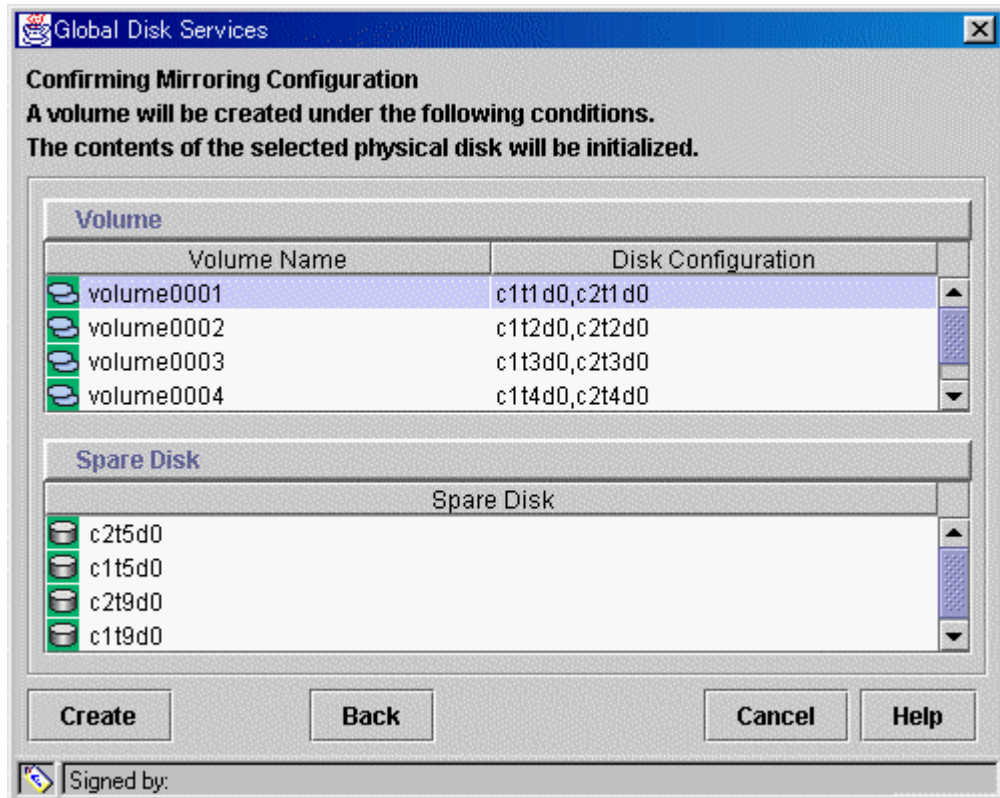
If there are disks you do not want to mirror, select the disk by clicking the icon, and press <Delete>. If there are disks you want to mirror in [Candidate Disk] field, select the disk by clicking the icon, and press <Add>.

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

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

Press <Next>, and the message below appears.

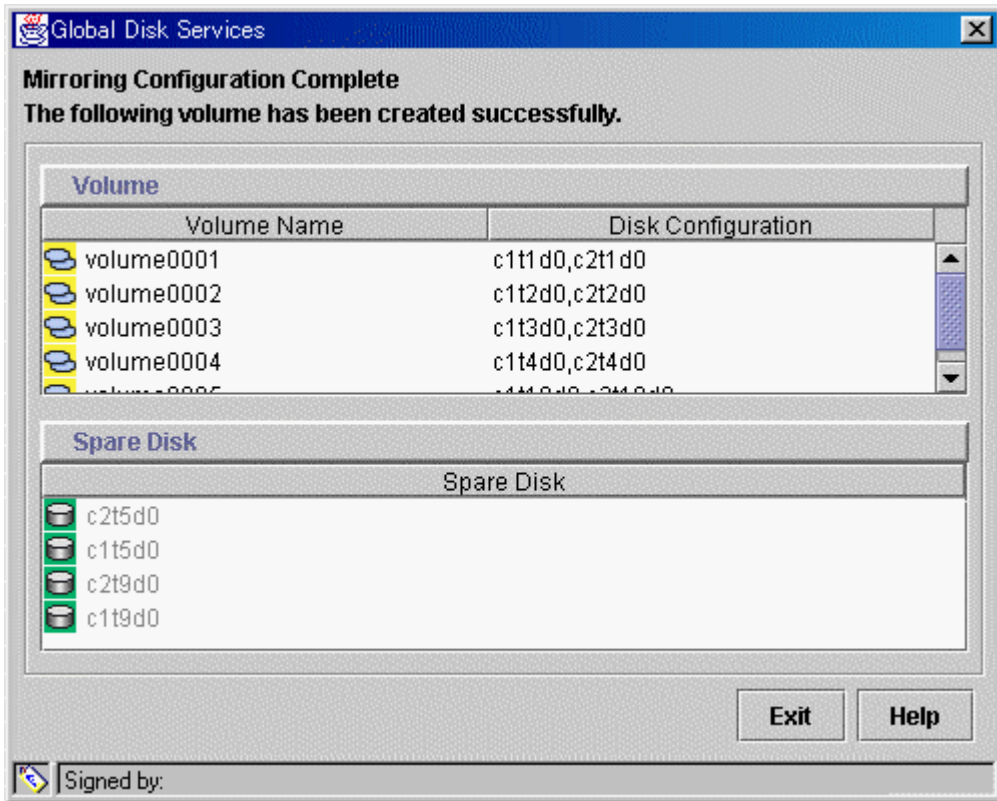
Figure: Confirming Mirroring Configuration



Confirming the Volume Configuration

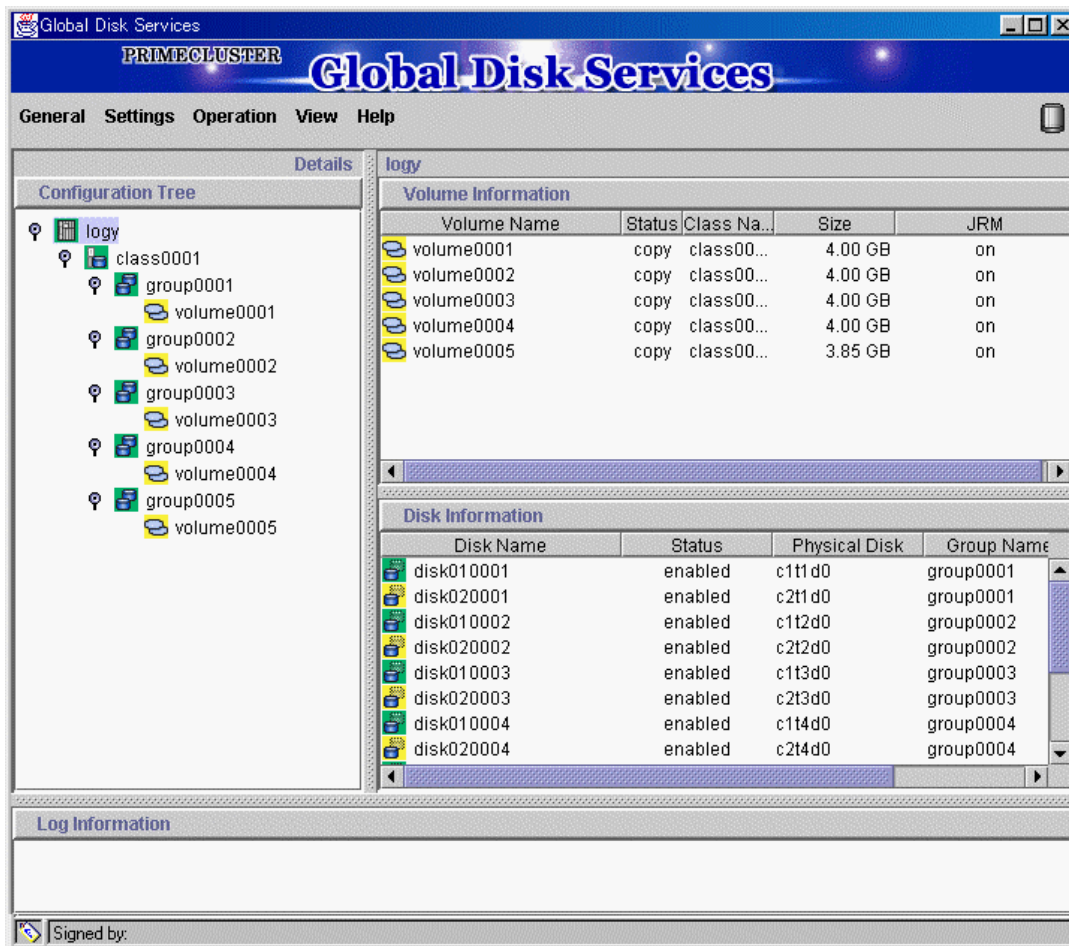
GDS will automatically combine the disks selected in the previous screen and create a volume. Volume name and its disk configuration, and the disk defined as spare disk will appear. Press <Create> to continue.

Figure: Mirroring Configuration Complete



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

Figure: Main Screen (after completing automatic configuration)



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

5.2.3 Operating from the Settings Menu

From the main screen, click [XXXXXX Configuration] on the [Settings] menu. Settings screen for different features appears. Use the screen switch tab to switch between "Class Configuration", "Group Configuration" and "Volume Configuration".



Note Configuring a System of High Reliability

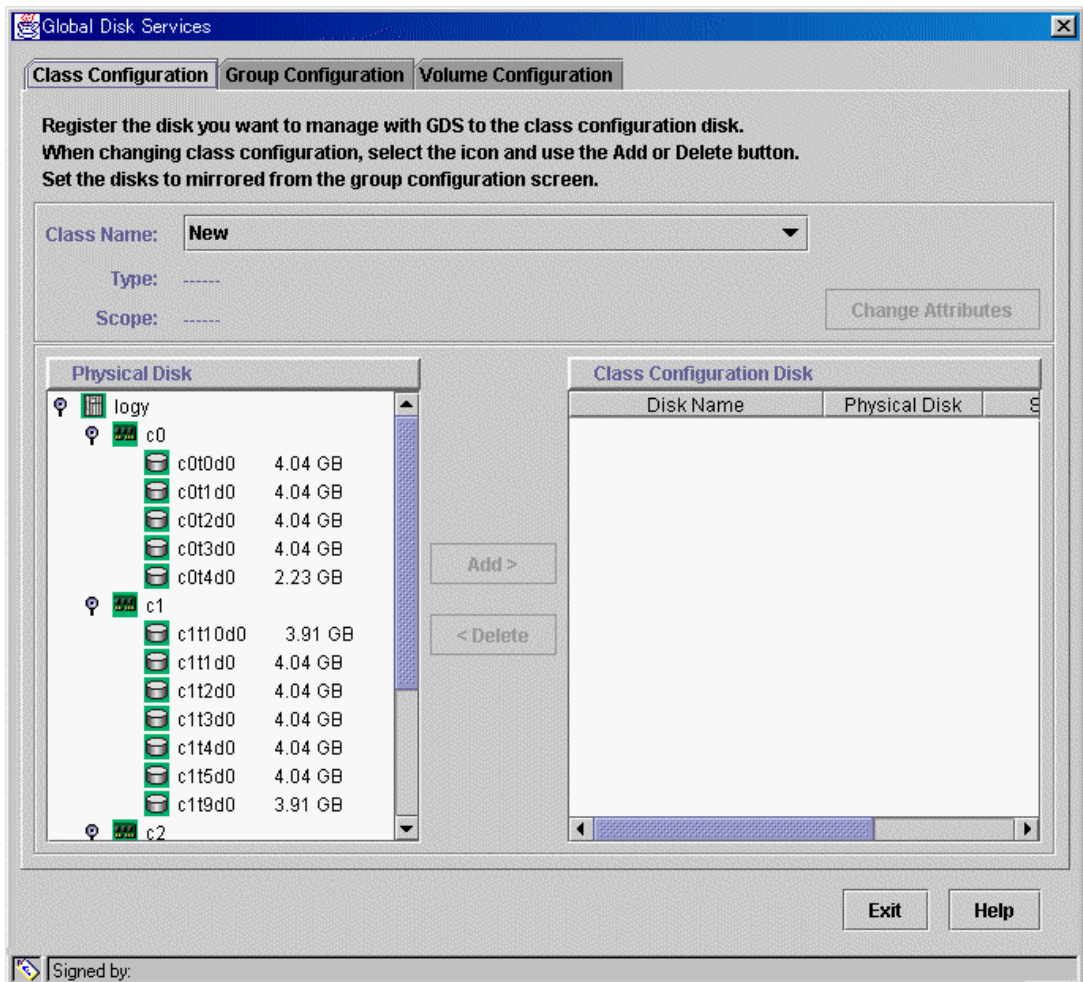
In order to ensure a system of high reliability, mirroring disks connected to separate I/O adapter, controller and cable is recommended.

5.2.3.1 Class Configuration

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

On the [Settings] menu, click [Class Configuration]. Class Configuration screen appears.

Figure: Class Configuration



1. Selecting class name

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

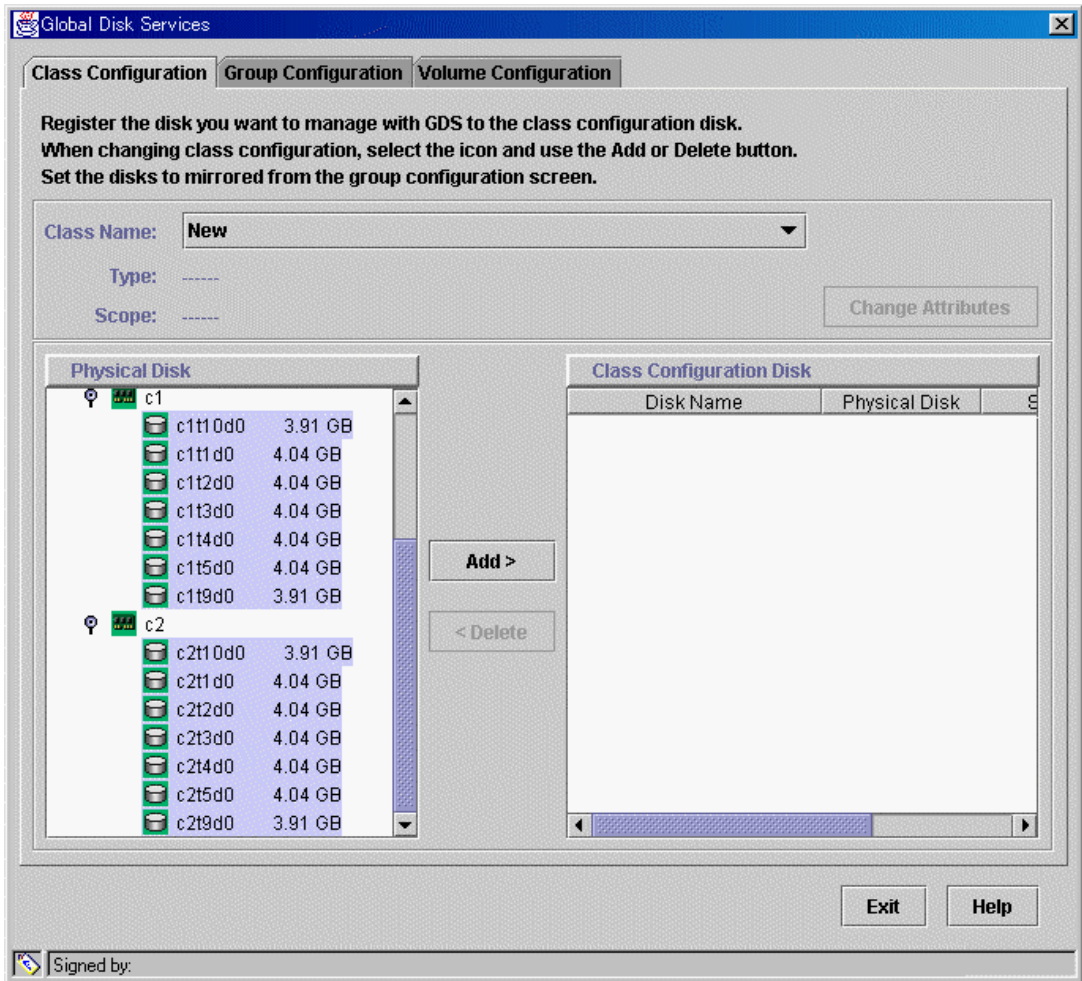
2. Selecting physical disk

In the <Physical Disk List>, 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

Press <Add>, and the message below appears.

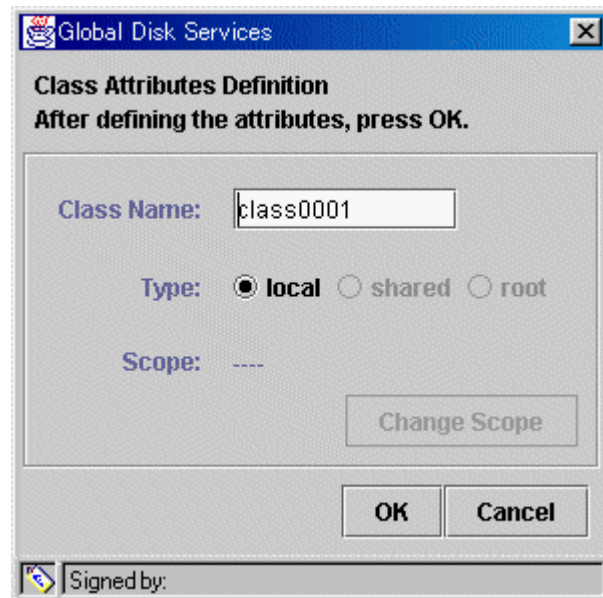
Figure: Warning



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

Press <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 press <OK>.



Note

Inputting Class Name

Class name will be used for the device path name.

/dev/sfdsk/class name/[r]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 press <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, press <Exit> and close Class Configuration screen.



Note

Class in Cluster System

When using cluster system, creating a class adds a class resource, and removing a class removes the resource.

When removing a class resource, remove the class and do not use the PRIMECLUSTER command.



See

For information on how to create a shared type class in cluster system, see "5.2.3.2 Cluster System Class Configuration".

5.2.3.2 Cluster System Class Configuration

When using cluster system, 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 List>, sets the type to "local" by default.

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

You can tell whether a disk is shared or not from the node information (disk information DEVCONNECT field) of sdxinfo command.

2. Displaying "Scope"

Displays connecting nodes that can share a class.

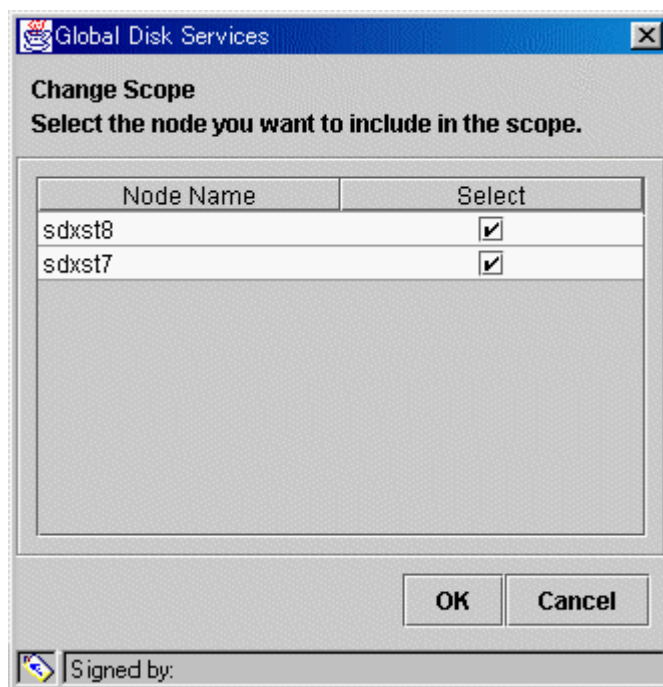
To change scope, press <Change Scope>.

3. <Change Scope> Button

Changes nodes connecting to class.

Press <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).
- Pressing < OK > of the [Change Scope] dialog box determines the class connecting node.
- Pressing <Cancel> of the [Change Scope] dialog box cancels the changing of the connecting node.



Note

Class Resource

- When using 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 command.



See

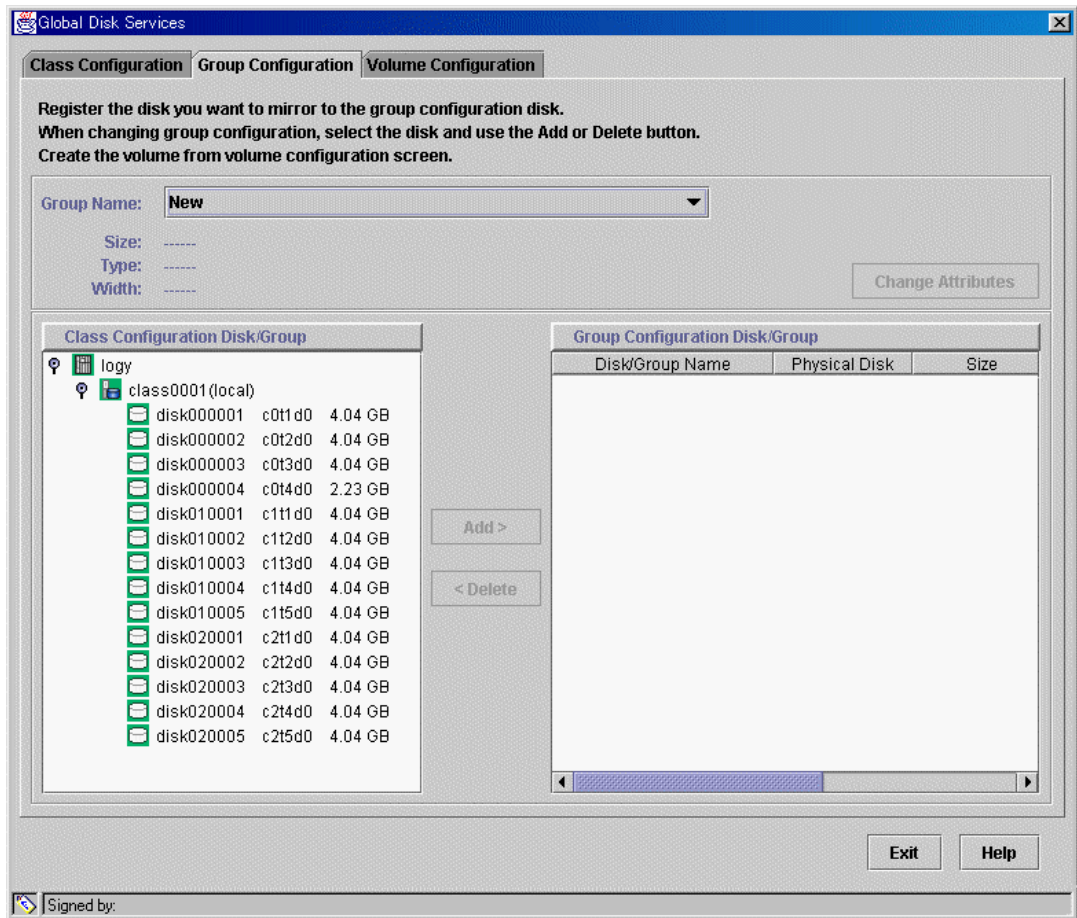
For information on installation and initial setting of cluster system, refer to "PRIMECLUSTER Installation Guide" and "PRIMECLUSTER Cluster Foundation Configuration and Administration".

5.2.3.3 Group Configuration

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

On the [Settings] menu, click [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."

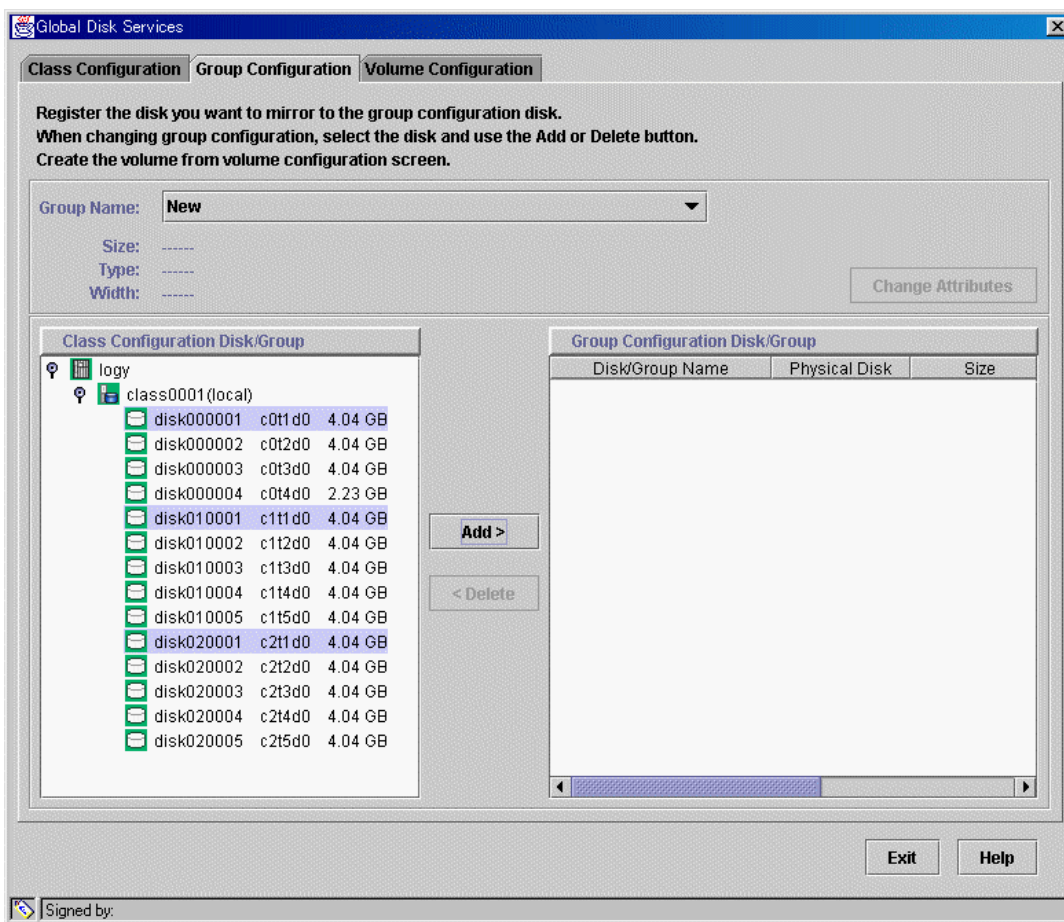
2. Selecting disk/lower level group

In the [Class Configuration Disk/Group] list, 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.

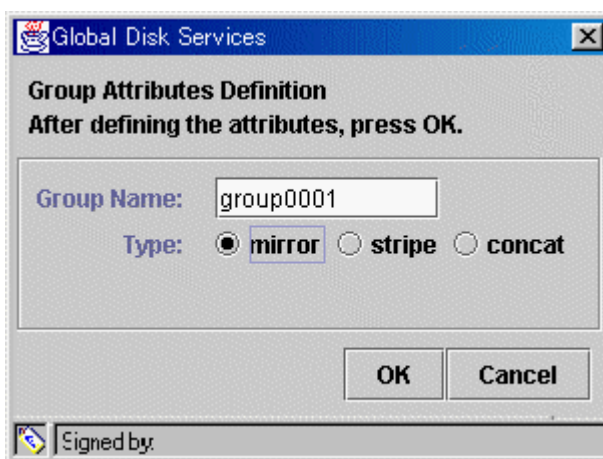
Figure: Selecting Disk/Lower level group to Configure Group



3. Creating a group

Press <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.



See

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

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, press <Exit> and a new group will be created.

If you press <Cancel> in the [Group Attributes Definition] screen, connection of the disk itself will be canceled.

4. Completing group creation

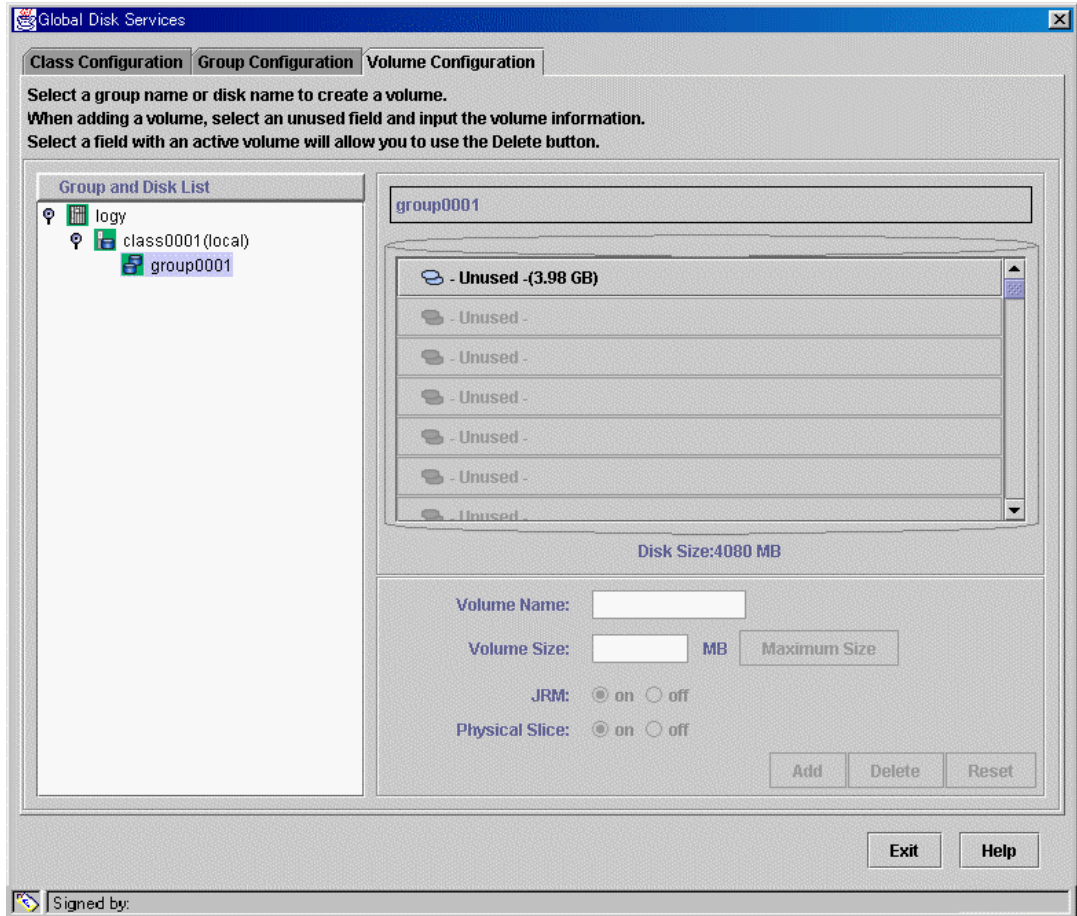
After creating all groups, press <Exit> and close [Group Configuration] screen.

5.2.3.4 Volume Configuration

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

On 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. Type numbers only.



Note

Size and Cylinder Boundary

When the specified size does not match the cylinder boundary, it will automatically adjust by rounding up. For details see "A.2.7 Volume Size".

c) <Maximum Size> button

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

d) Setting "JRM"

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

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

e) Setting "Physical Slice"

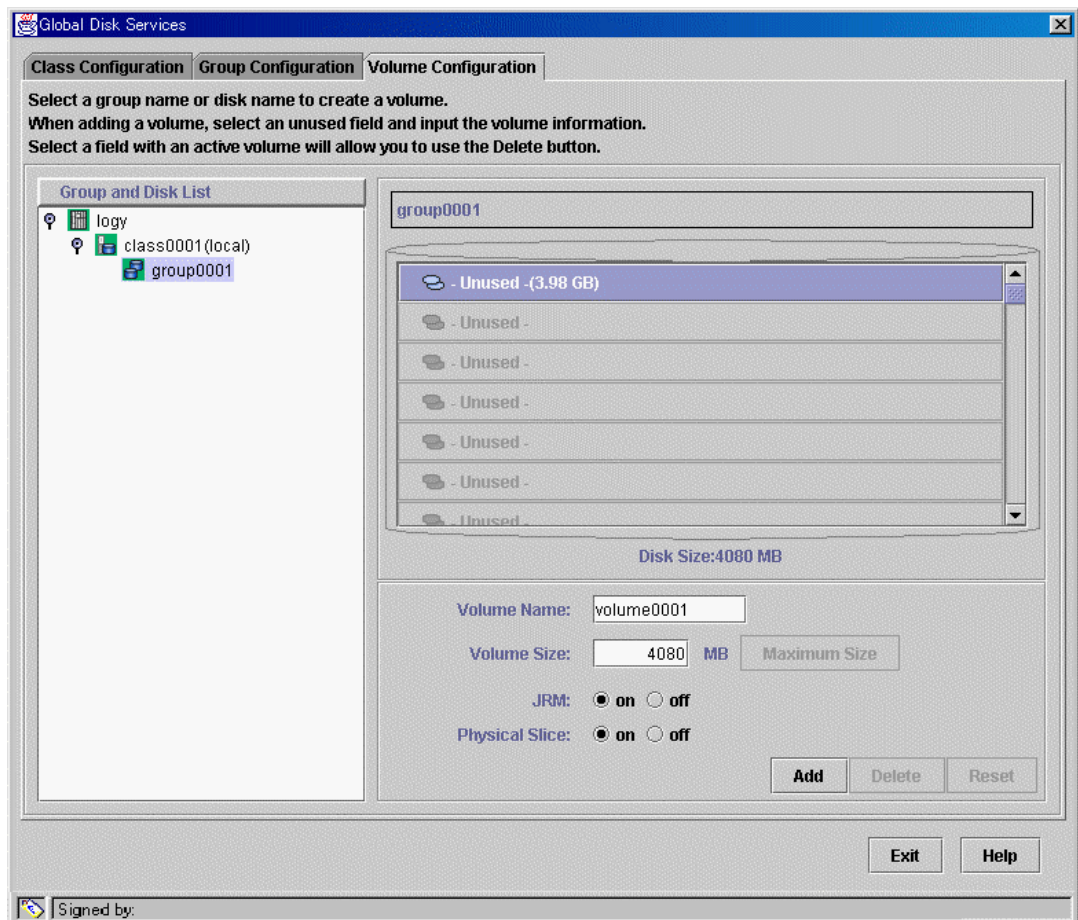
Sets the volume's physical slice attribute value.

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

/dev/sfdsk/class name/dsk/volume name



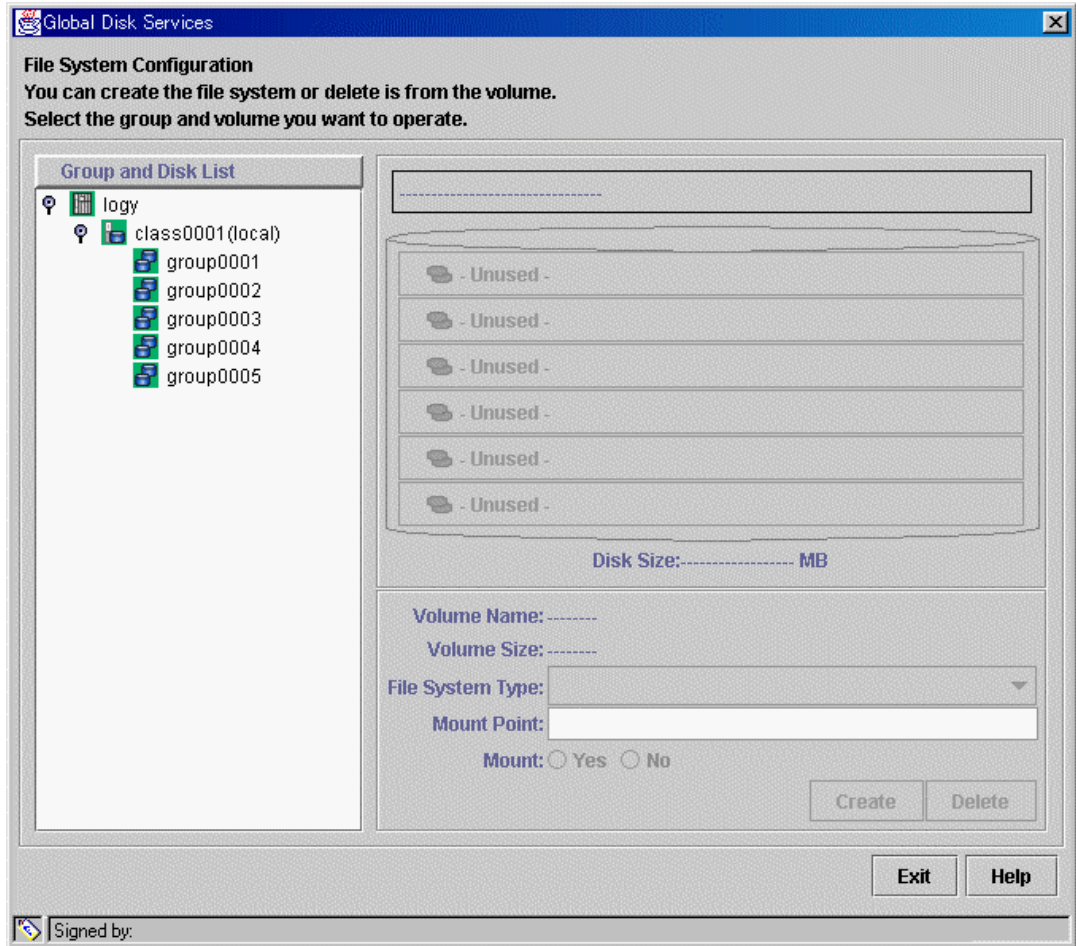
Note Shared Class Volume Operation

Volume created in a shared class cannot be used from other nodes immediately. If you wish to access from other nodes, you must activate the volume from the node you wish to access. After activating the volume, the node which created the volume, and the node which activated the volume will gain access. However, since operating from two nodes could affect data consistency, you must be careful when operating the volume.

5.2.4 File System Configuration

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

Figure: File System Configuration



1. Selecting a group/disk

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

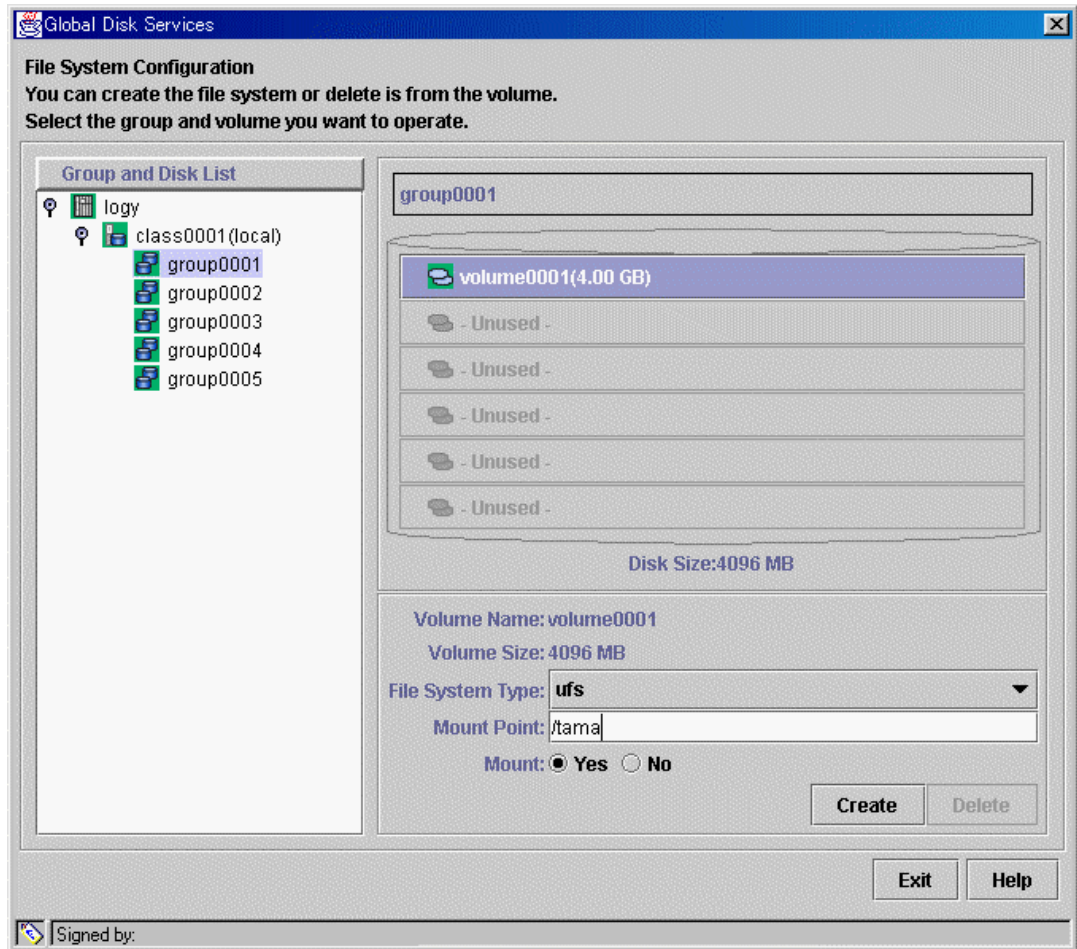
2. Selecting a volume

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. It is set to "ufs" by default.
When using GFS Local File System, select "sfxfs."

b) Setting "Mount Point"

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

c) Setting "Mount"

Select "Yes" if you want to automatically mount when you start the volume, and "No" if not.

Selecting "Yes" will mount the file system to the specified mount point immediately after its creation.

However, you cannot select "Yes" when the volume in which you will create the file system exists on a shared disk unit.



Note

When Selecting "No" for "Mount"

When you select "No" for "Mount," [mount at boot] field in the /etc/vfstab file will be set to "No" and the new mount information will be added. Change the mount information as necessary.

4. Creating the file system

After setting the attributes, press <Create>. A new file system will be created.

After creating all file systems, press <Exit>.



Note

When Using Cluster System

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 Reliant Monitor Services (RMS) Configuration and Administration Guide."

5.3 Operation in Use

GDS operates monitoring and maintenance from the main screen.
This section explains the following operation:

- confirming configuration and monitoring status
- backup
- in case of disk abnormality
- copying operation

5.3.1 Confirming Configuration and Monitoring Status

By selecting certain configuration unit from GDS configuration tree on the main screen, you can confirm the object configuration and monitor the status.

5.3.1.1 Confirming Object Configuration

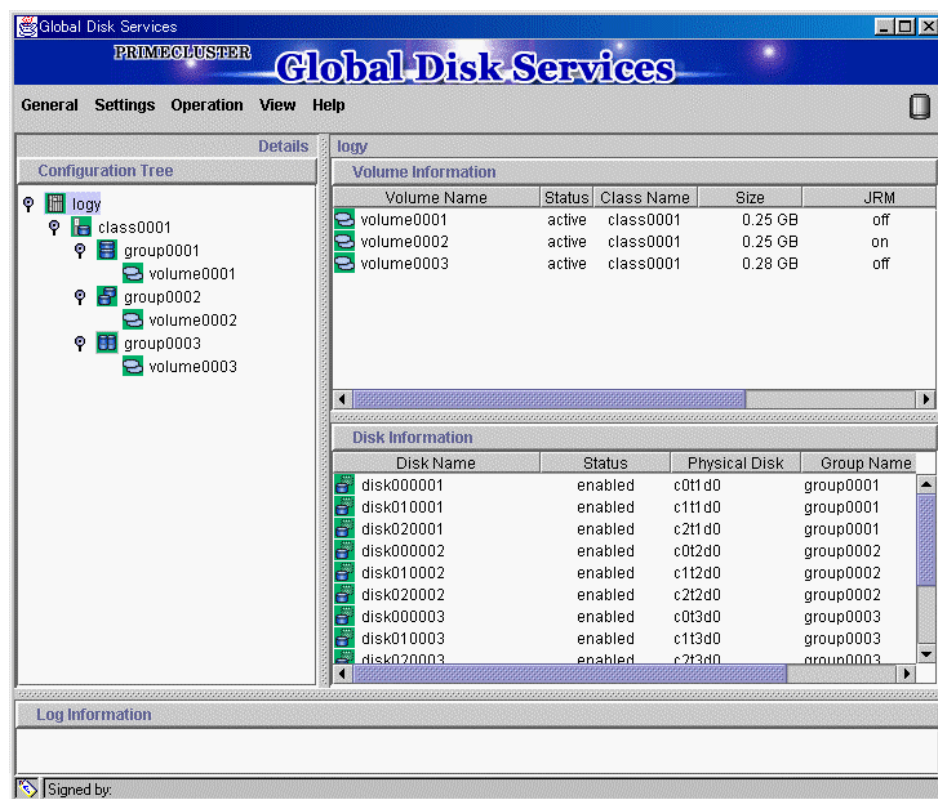
By selecting the following units on the GDS configuration tree, you can confirm the configuration of each object.

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

Object configuration within a node

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

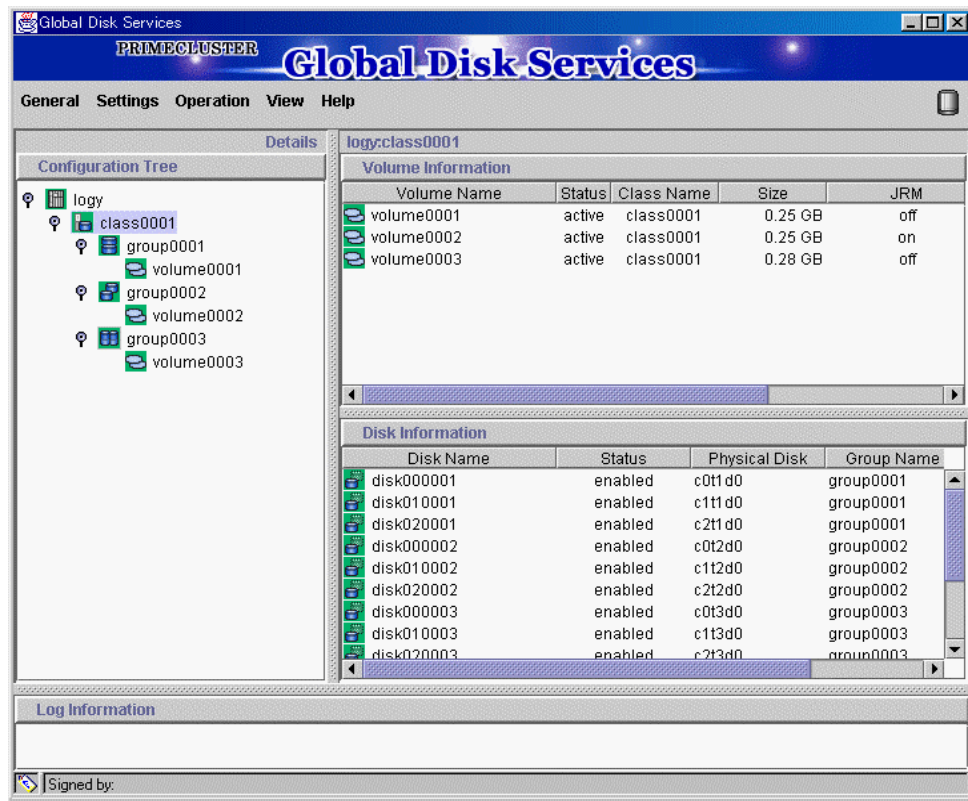
Figure: Main Screen (when selecting a node)



Object configuration within a class

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

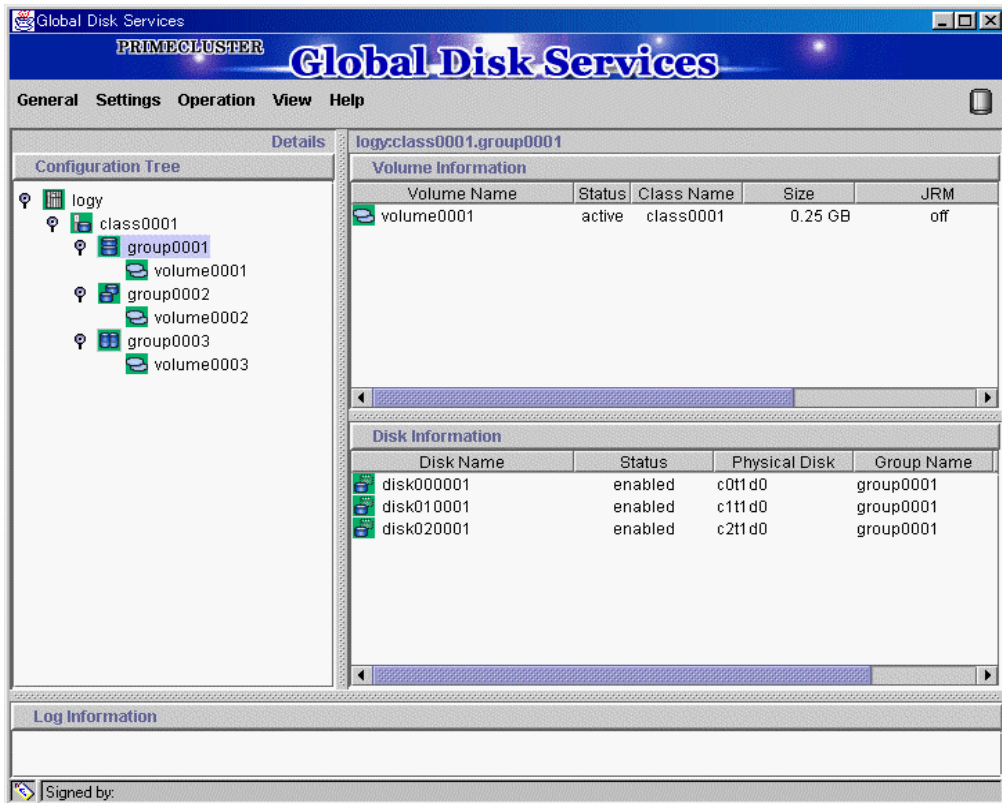
Figure: Main Screen (when selecting a class)



Object configuration within a group

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

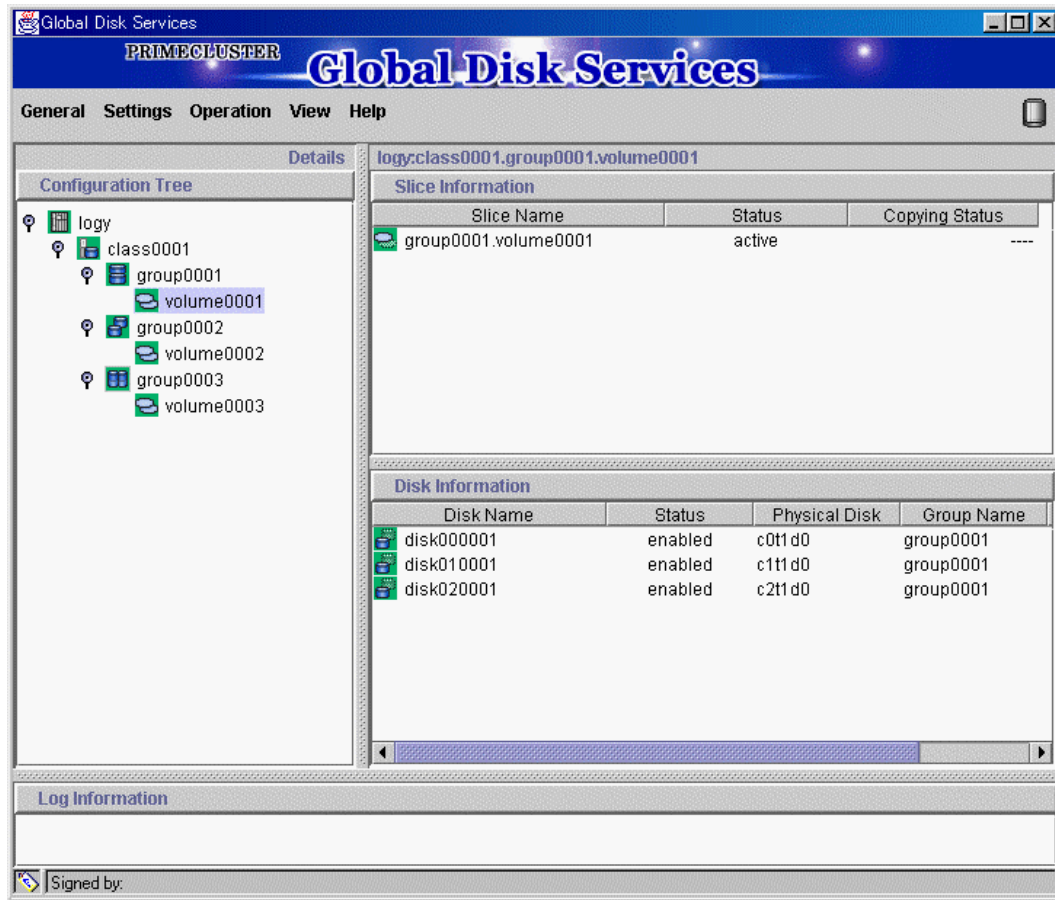
Figure: Main Screen (when selecting a group)



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 (when selecting a volume)



5.3.1.2 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] on the [General] menu. You may also use [Update] on the [View] menu to update the status immediately.

When an abnormality is detected with objects, 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, icon color and description in the status field will change as well.



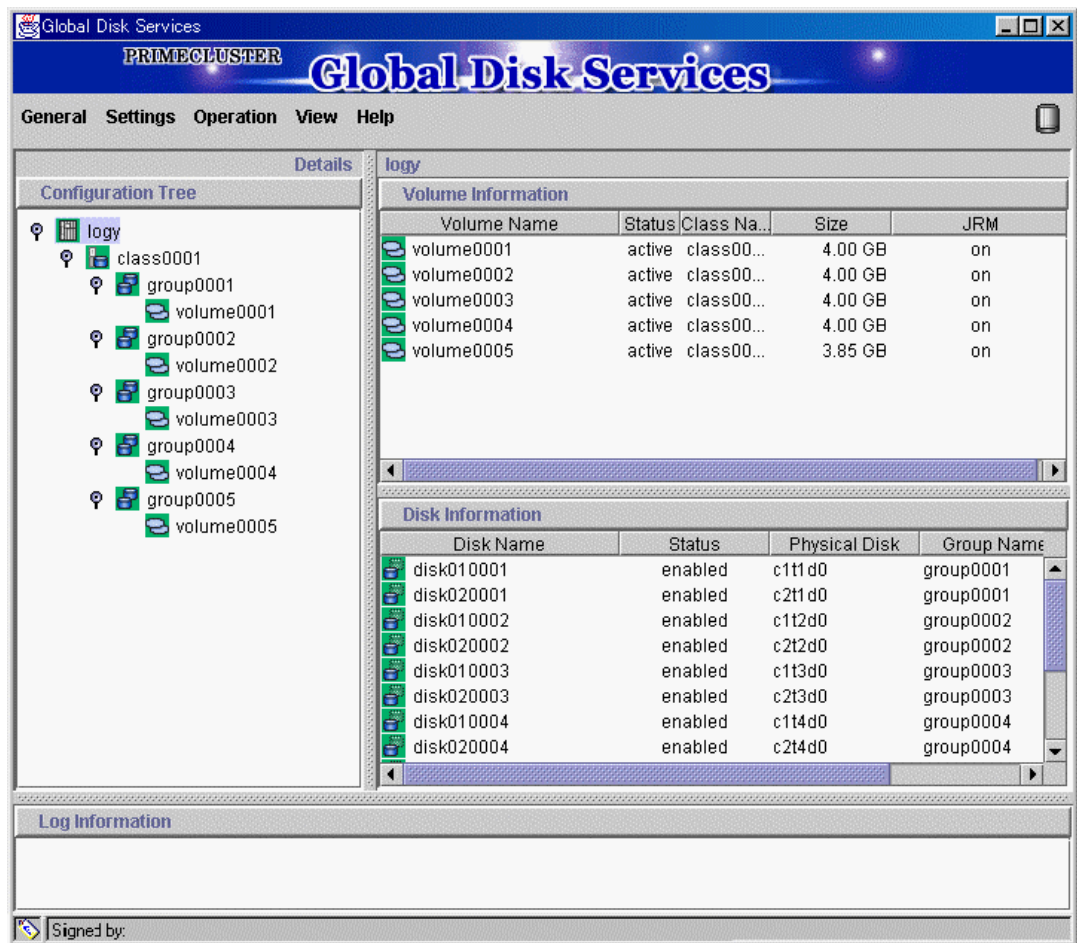
See

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

When a disk unit fails during operation, slice icon in which the abnormality is detected will turn red, and the disk icon, yellow.

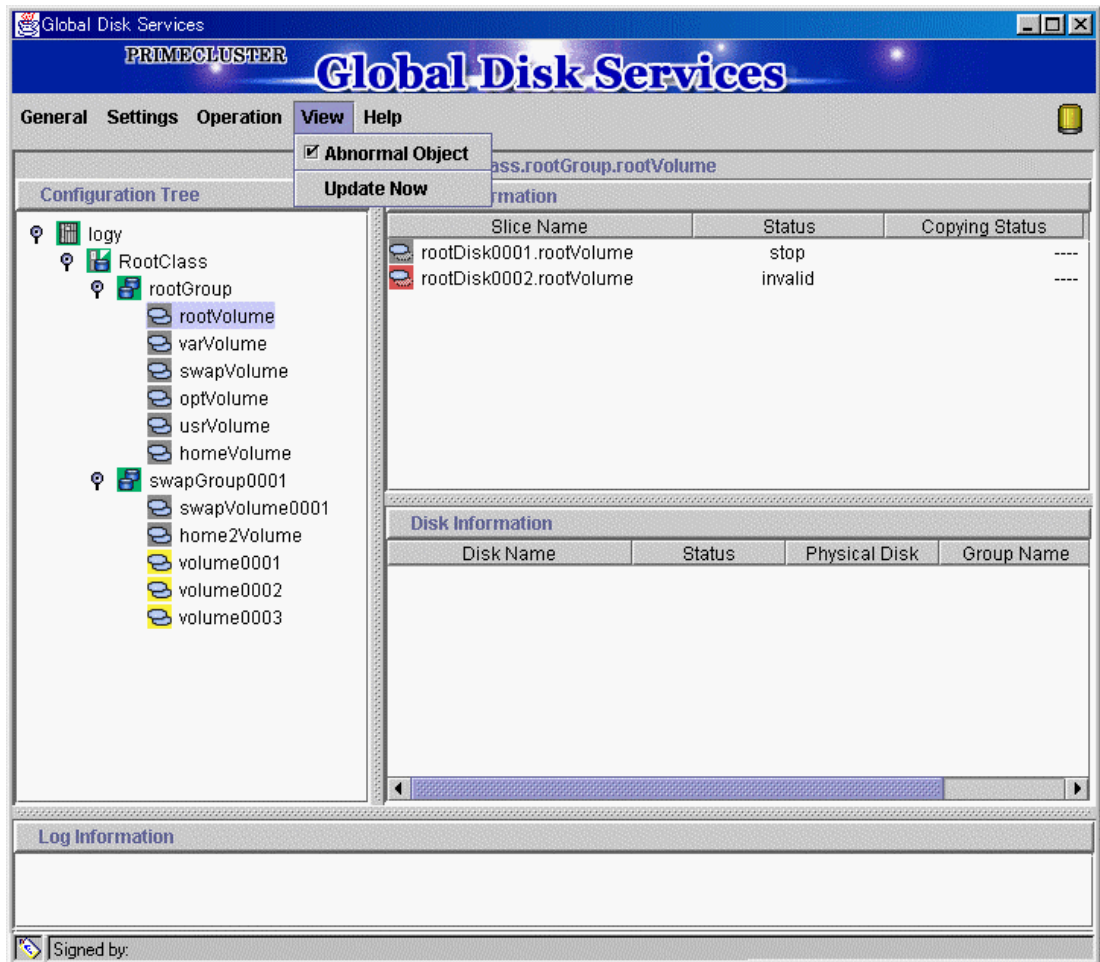
To recover, follow the procedures in "5.3.3 In Case of Disk Abnormality."

Figure: Object Status Monitoring in Main Screen



Clicking [Abnormal Object] on 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 then the abnormality is detected.

5.3.2 Backup

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

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 definite countermeasure against such problems is to create a "backup." To minimize the damages caused by troubles, periodical backup is strongly recommended.

GDS requires the following slice operation when creating a backup.

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



Note Requirement for Applying a Snapshot by Detaching a Slice

A slice can be detached only from a mirror volume with a physical slice. So if a disk is not connected directly to a mirror group, creation of a snapshot by detaching a slice is not available.



See

For procedures on system disk backup, see "F.1.5 System Disk Abnormality."



Information Restore

To restore a volume using the backup data, `/dev/sfdsk/class name/rdsk/volume name` which is the access path to the volume can be used. However, the restore using the access path to a volume is unavailable when it is the `/`(root), `/usr`, or `/var` file system that does not support unmounting during the system operation. For details on the restoring procedures in this case, see "(2) System can be booted, but the system disk data is invalid" and "(4) System cannot be booted (Boot disk data damage)" in "F.1.5 System Disk Abnormality."

Detach Slice

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.

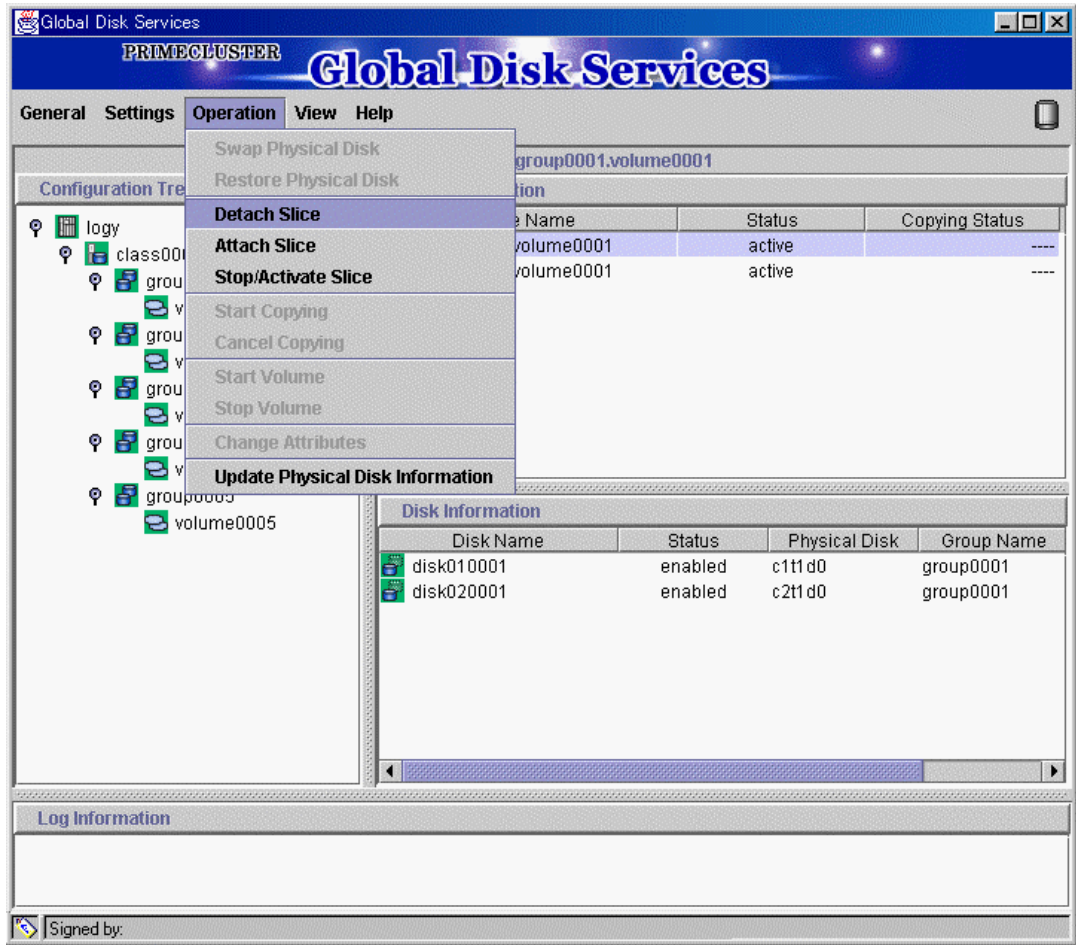
2. Selecting the slice to detach

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, click [Detach Slice].

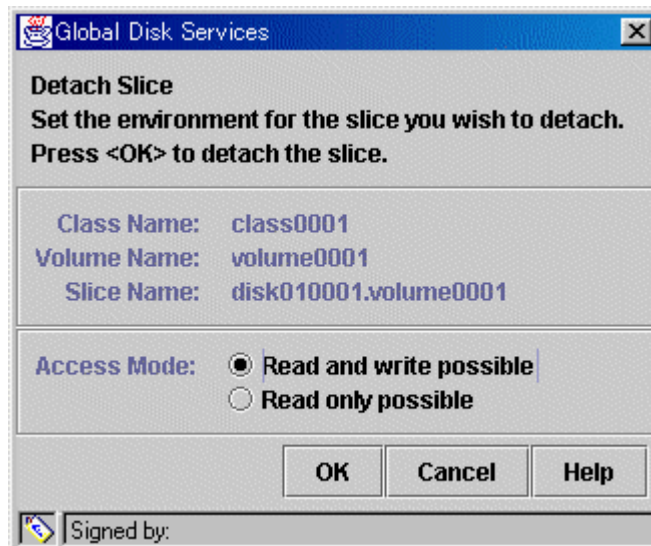
Figure: Detach Slice



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.

Default is set to [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.

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



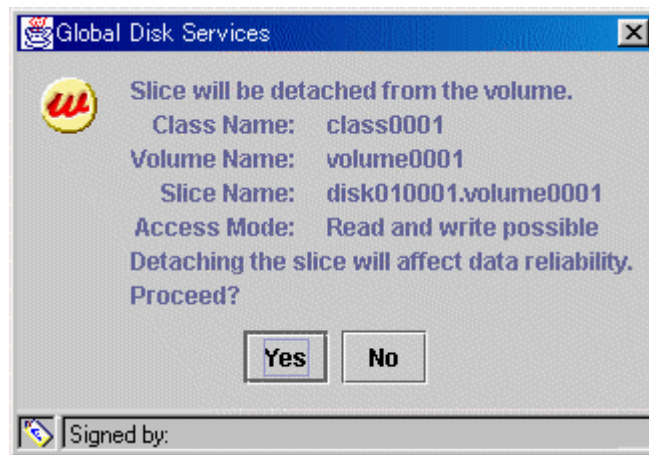
Note Slice status available for [Detach Slice]

You can only perform [Detach Slice] operation to slices that are in either "active" or "stop" status.

The confirmation screen shown below appears.

To proceed, press <Yes>. Pressing <No> cancels the [Detach Slice] process.

Figure: Confirming Slice Detach

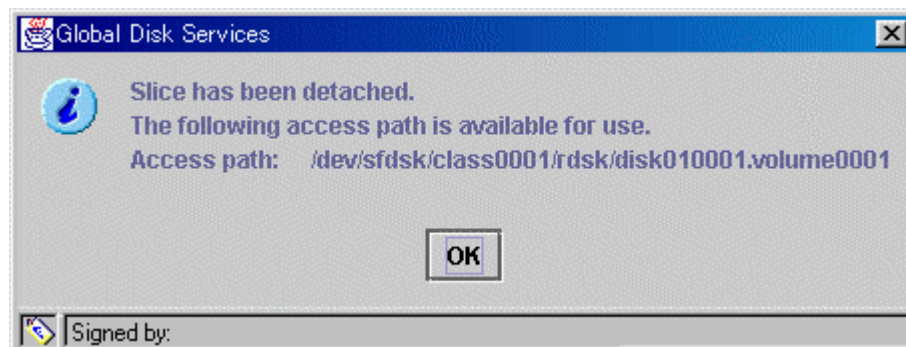


5. Backup Using Access Path

Press <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, slice that is temporarily detached will be attached to volume again. If the volume is activated, synchronization copying will begin. The procedures are explained below.

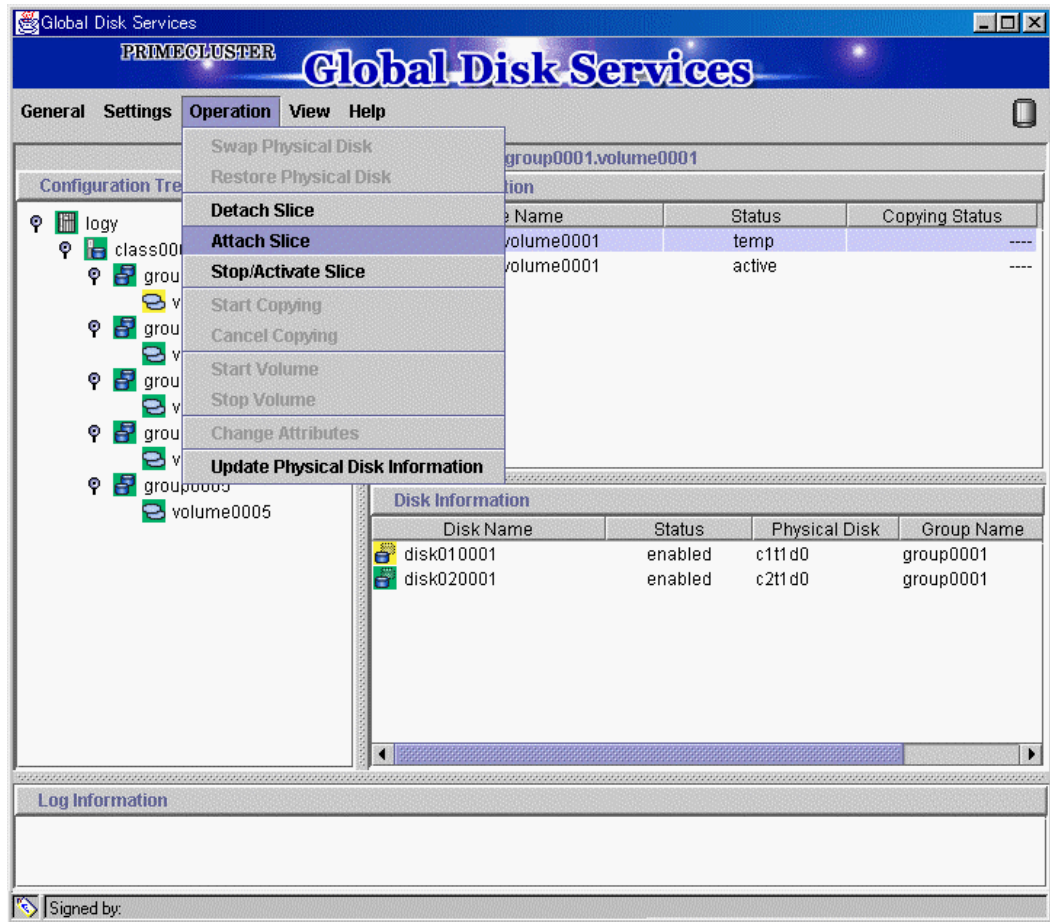
1. Selecting the slice to attach

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, click [Attach Slice].

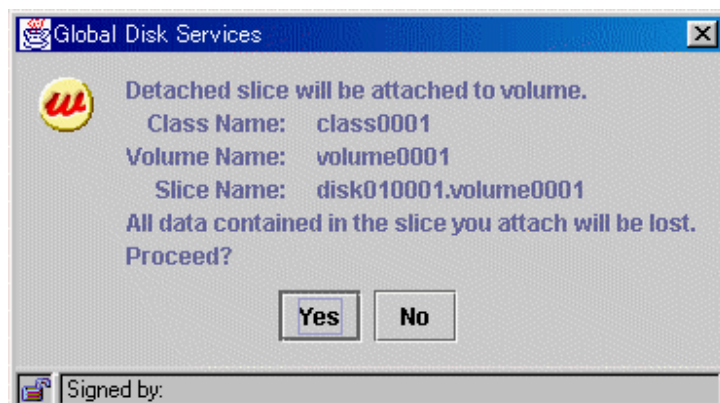
Figure: Attach Slice



The confirmation screen shown below appears.

To proceed, press <Yes>. Pressing <No> cancels the [Attach Slice] process.

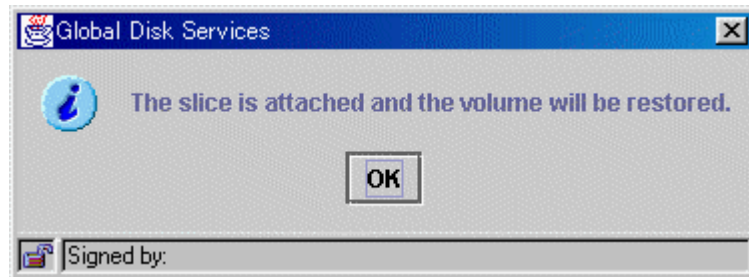
Figure: Confirming Attach Slice Completion



3. Notifying the completion of Attach Slice

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

Figure: Notifying Attach Slice Completion



Reliable Backup Procedures

Although the procedure above allows you 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, following the procedures below.

1. Register backup disk with class.
2. Connect backup disk to group.
3. After completing synchronization copying, stop operation.
4. Detach slice.
5. Resume operation.
6. Perform backup using access pat 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 Operation

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

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

When excluding the slice without stopping operation, 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, slice in "temp" status will temporarily become inaccessible.

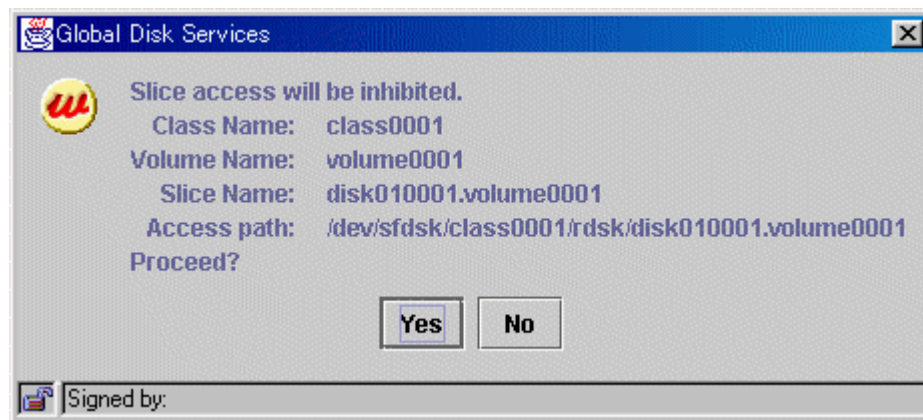
1. Selecting the slice to stop

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, click [Stop/Activate Slice].

Figure: Stop Slice



To proceed, press <Yes>. Pressing <No> cancels the [Stop Slice] 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, click [Stop/Activate Slice].

Figure: Activate Slice



To proceed, press <Yes>. Pressing <No> cancels the [Activate Slice] process.

5.3.3 In Case of Disk Abnormality

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 Physical disks that can not be swapped

A disk with the only active slice of the volume cannot be swapped. In order to swap such a disk, connect a new disk to the group. You will not be able to perform swapping under the following situations.

- A volume exists within a group, and only one disk is connected to the group.
- A volume exists in a "single" type disk that is not connected to a group.

A physical disk cannot be swapped with another either when the disk is connected to a group and the related highest level group is a stripe group or a concatenation group.

When swapping a physical disk with another, you should disconnect the disk after creating a backup copy of the data as necessary.

Swap Physical Disk

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

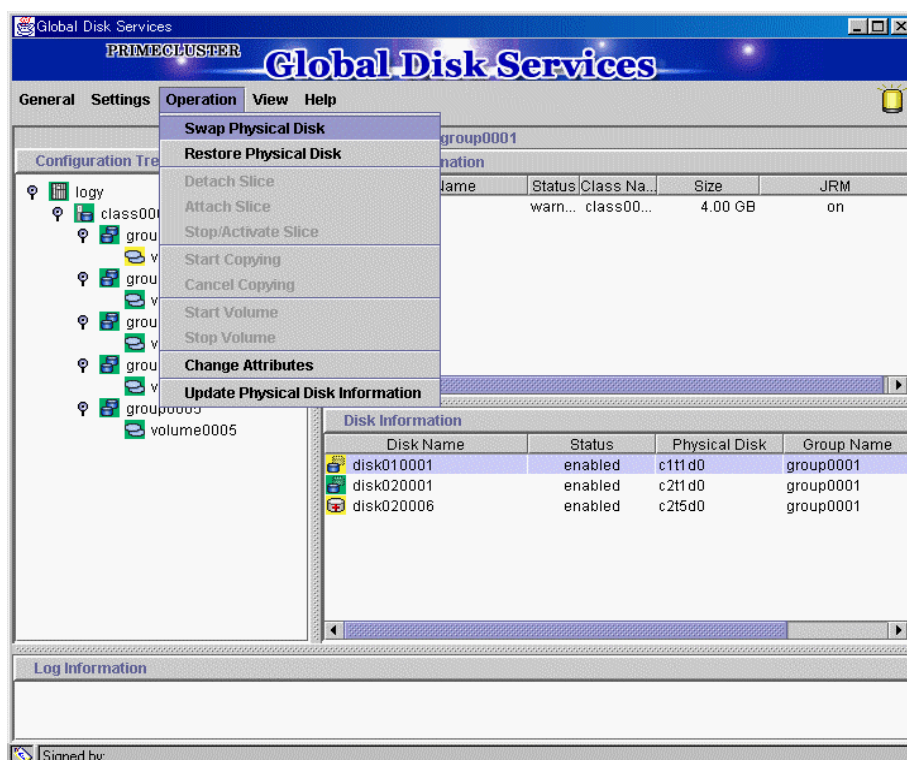
1. Displaying the status of physical disk

In the main screen, display the physical disk to be swapped. Click the icon and select the physical disk.

2. Selecting [Swap Physical Disk]

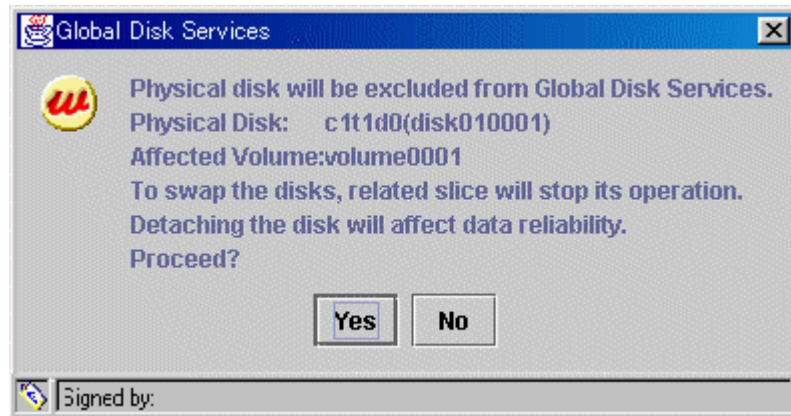
In the main screen [Operation] menu, select [Swap Physical Disk].

Figure: Swap Physical Disk



The confirmation screen shown below appears.
To proceed, click <Yes>. Clicking <No> cancels the [Swap Physical Disk] 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



Note To swap with a brand-new disk

It is necessary to label the new disk using the `format(1M)` command before performing the procedure described in "Restore Physical Disk."

Restore Physical Disk

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

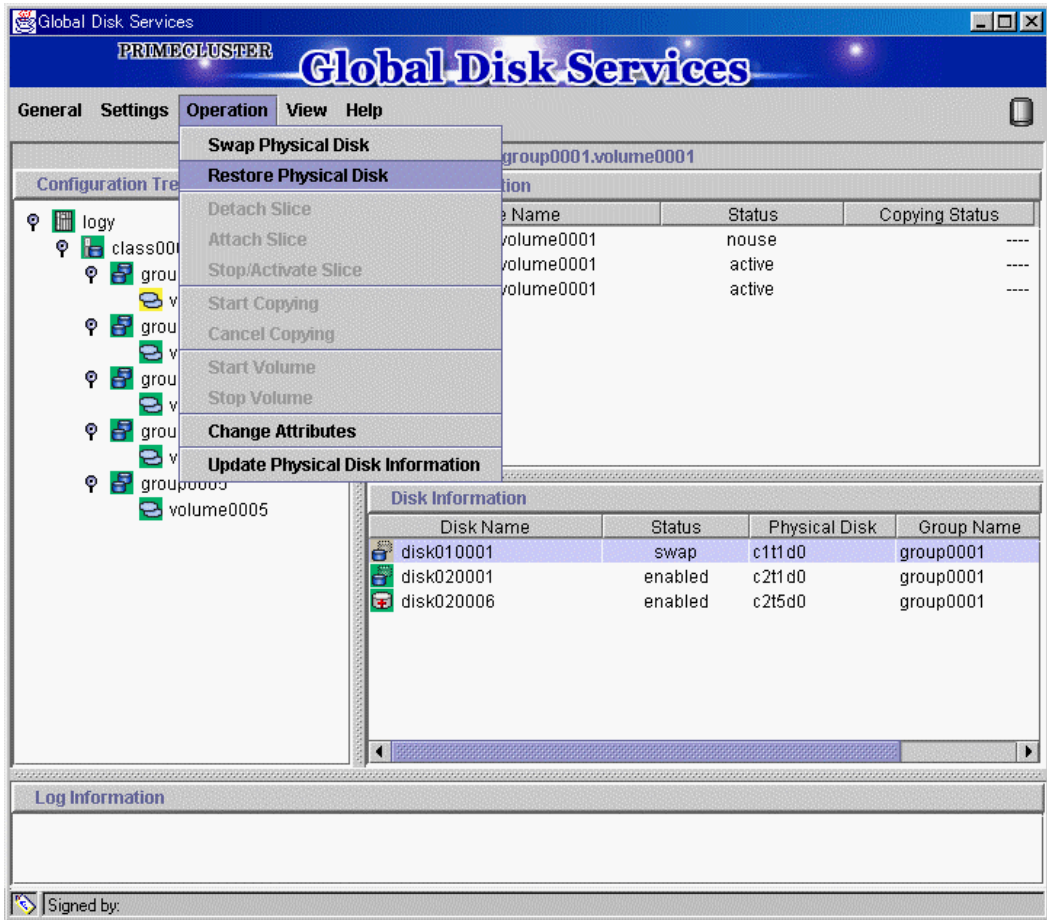
1. Selecting the physical disk to restore

Select the physical disk you want to restore.

2. Selecting [Restore Physical Disk] menu

In the main screen [Operation] menu, select [Restore Physical Disk].

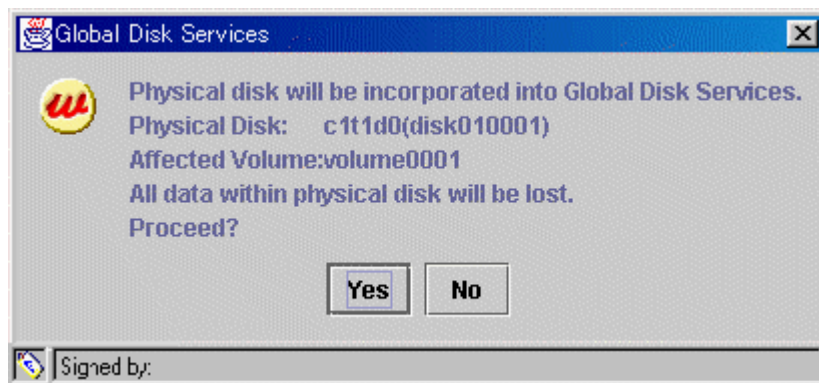
Figure: Restore Physical Disk



The confirmation screen shown below appears.

To proceed, click <Yes>. Clicking <No> cancels the [Restore Physical Disk] process.

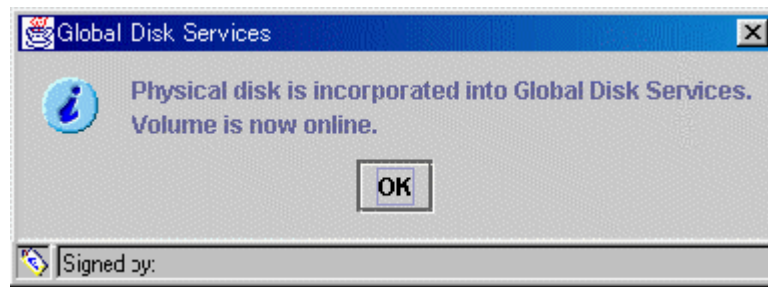
Figure: Confirming Restore Physical Disk



3. Notifying online status

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

Figure: Notifying Online Completion



5.3.4 Copying Operation

Copying operation allows you to control the synchronization copying of the mirror volume. 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.

Slice in "copy-stop" status will resume copying from the point where copying was interrupted.

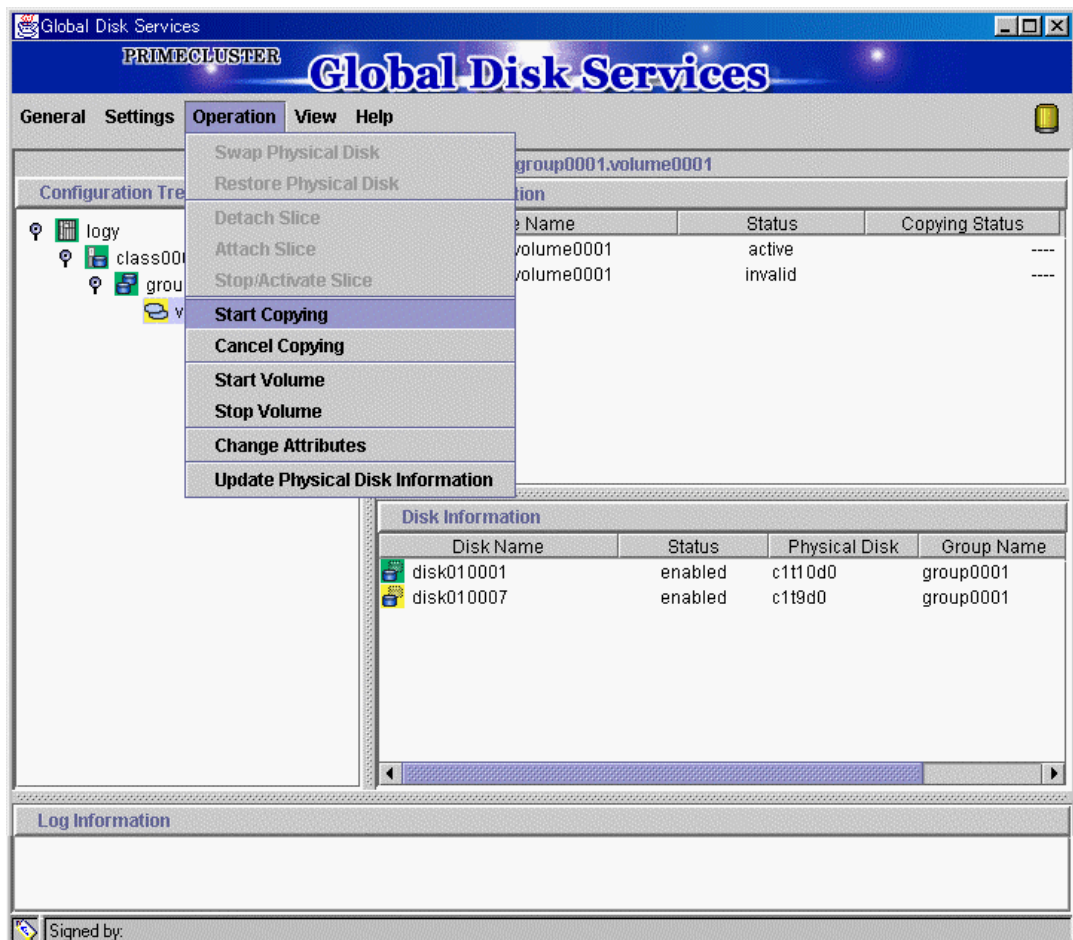
1. Select volume for synchronization copying

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

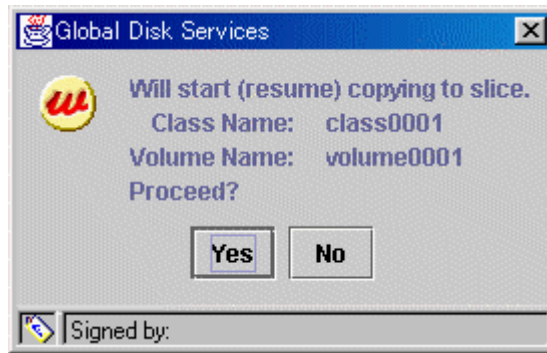
In the main screen [Operation] menu, click [Start Copying].

Figure: Start Copying



The confirmation screen shown below appears.
To proceed, press <Yes>. Pressing <No> cancels the [Start Copying] process.

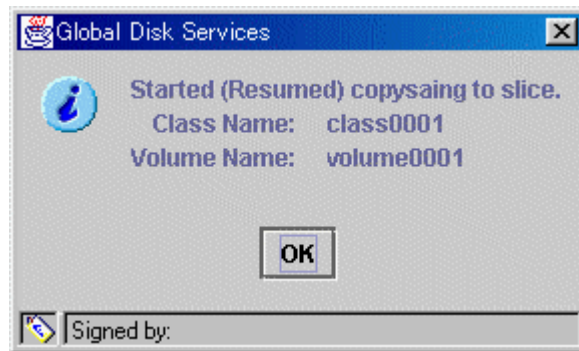
Figure: Confirming Start Copying



3. [Start Copying] completion screen

Press <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

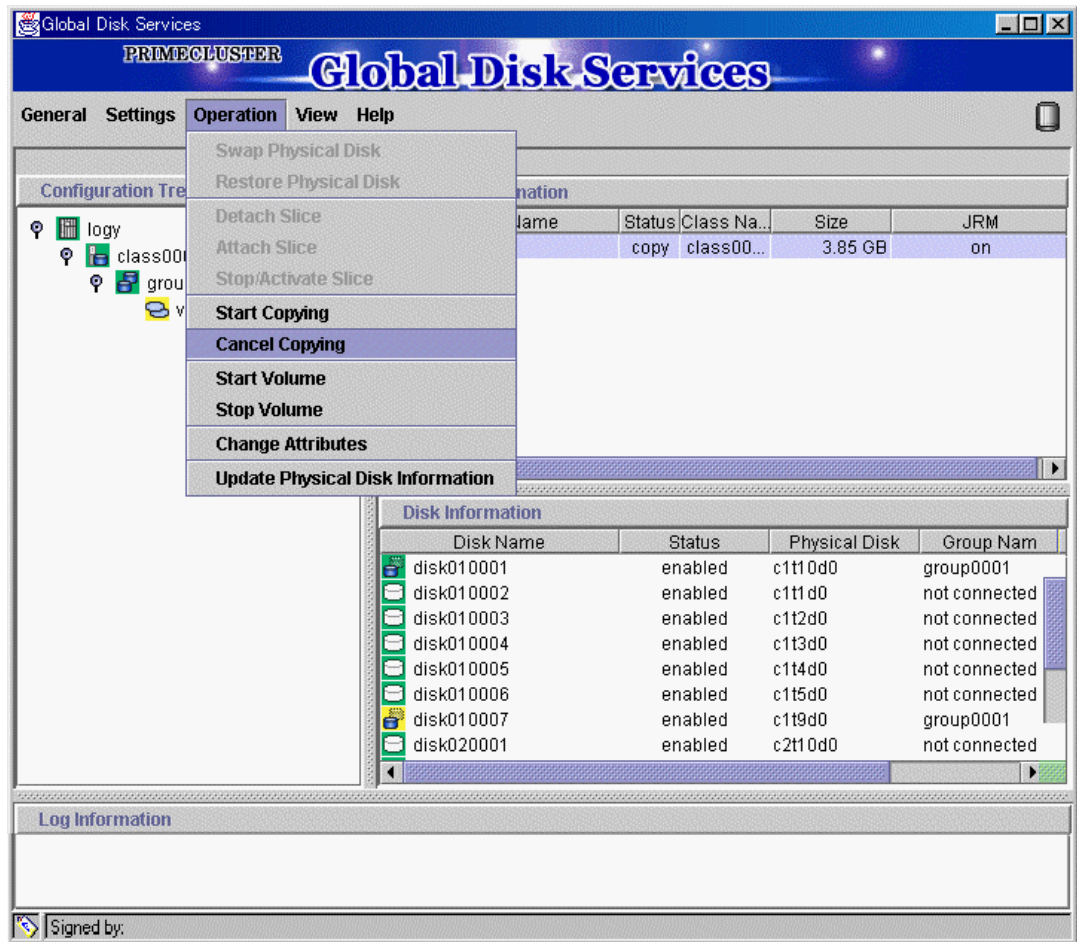
1. Select the volume to cancel copying

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, click [Cancel Copying].

Figure: Cancel Copying



The confirmation screen shown below appears.
To proceed, press <Yes>. Pressing <No> cancels the [Cancel Copying] process.

Figure: Confirming Cancel Copying



3. [Cancel Copying] completion screen

Press <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

Making changes refers to changes to configuration and attributes for class configuration, group configuration, and volume configuration.

This section explains the changing procedures for each configuration.

5.4.1 Class Configuration

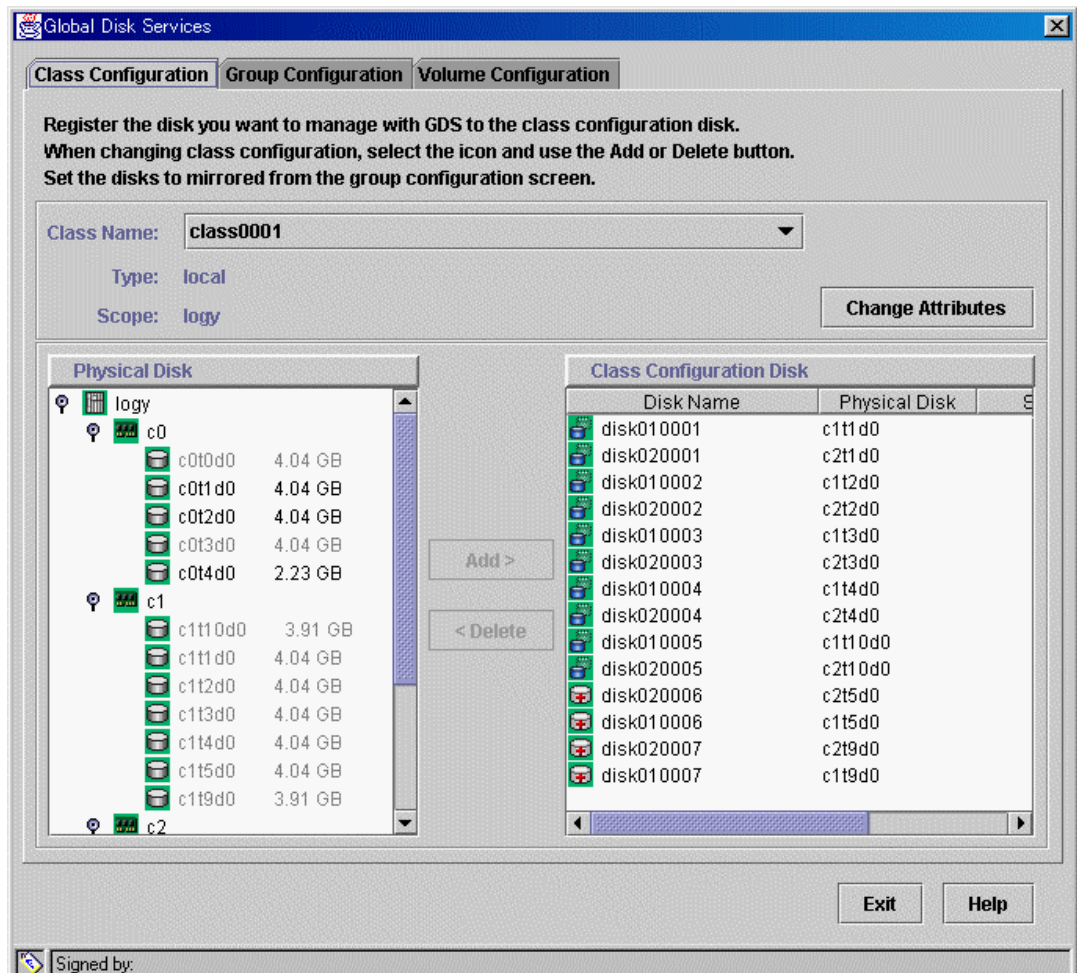
The procedures for changing class configuration are explained below.

Changing Class Configuration

1. Displaying Class Configuration screen

In the main screen [Settings] menu, click [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."

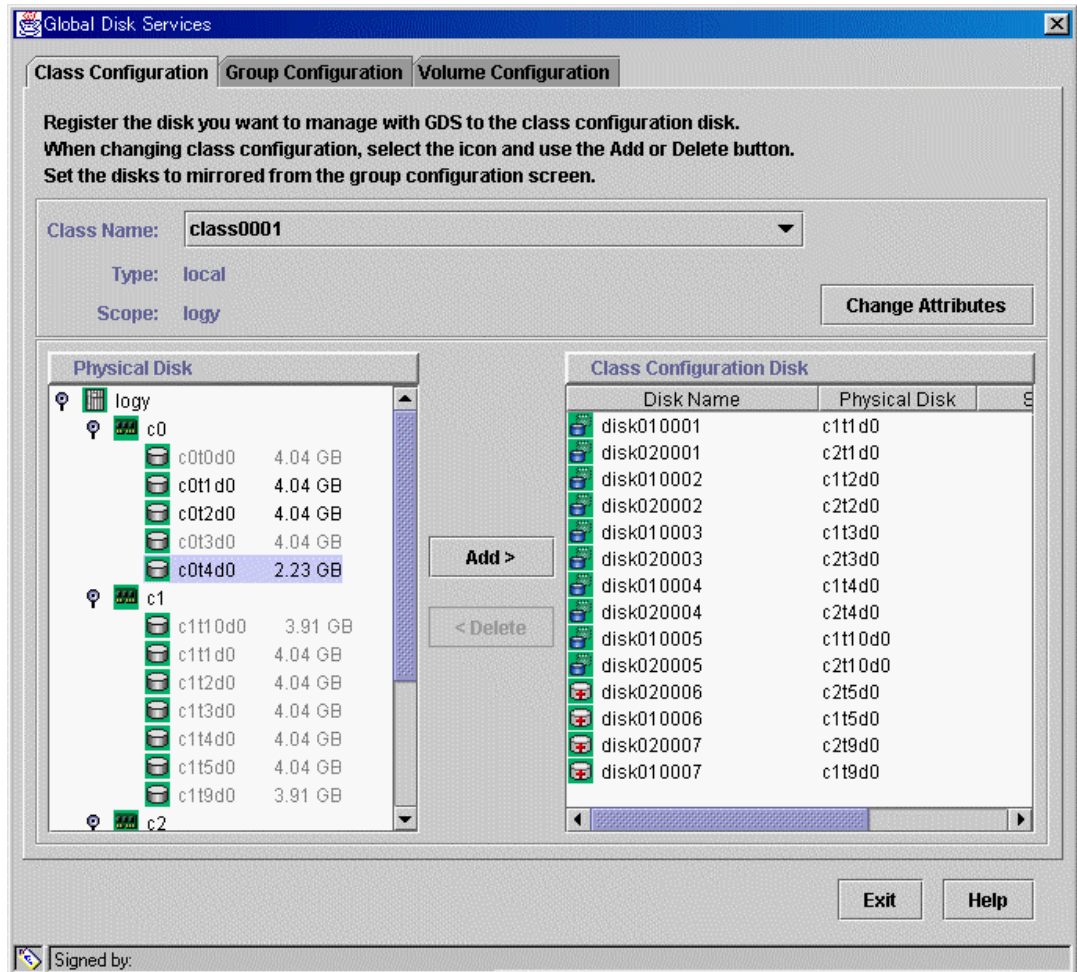
3. Changing configuration

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. Press <Add>.
3. When changing the default disk attributes, change the disk type by selecting the appropriate disk from the [Class Configuration Disk/Group] field.

Figure: Physical Disk Registration



b) Removing a disk

1. Select the disk you want to remove from the [Class Configuration Disk/Group] field.
2. Press <Delete>.

4. Completing the change

If you have no other changes to make, press <Exit>.

Changing Disk Class Attributes

When changing the disk class attributes, there are two procedures to choose from:

- using main screen Operation menu
- using Change Attributes button in the Class Configuration screen

Using main screen Operation menu

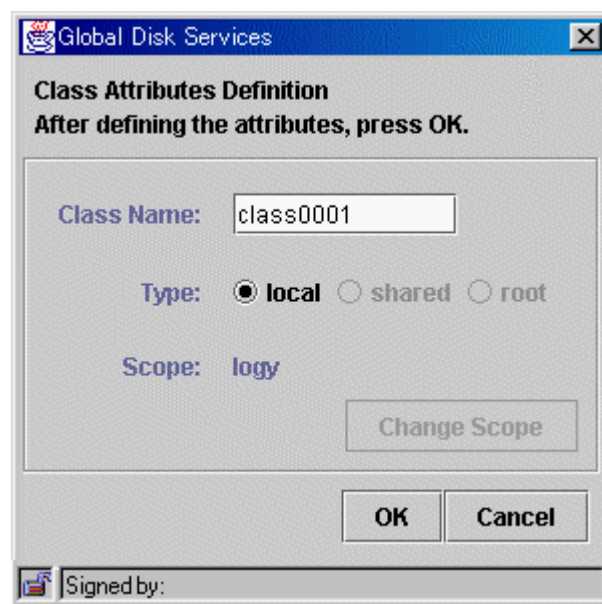
1. Selecting class to change

Display the class you want to change in the main screen Configuration Tree field. Click the icon and select the class you want to change.

2. Displaying Class Attributes Definition screen

In the main screen [Operation] menu, click [Change Attributes]. Class Attributes Definition screen appears.

Figure: Class Attributes Definition Screen



3. Changing attributes

- a) When using **single node**.
You can only change the class name.
- b) When using cluster system
You can change the following attributes:
 - Class Name
 - Type
 - Scope



See

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

4. Implementing changes

Press <OK> to implement the attributes change, or <Cancel> to cancel.

Using Change Attributes button in the Class Configuration screen

1. Displaying Class Attributes Definition screen

Press <Change Attributes> and Class Attributes Definition screen appears.

2. Changing attributes

a) When using single node.
You can only change the class name.

b) When using cluster system
You can change the following attributes:

- Class Name
- Type
- Scope



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

3. Implementing changes

Press <OK> to implement the attributes change, or <Cancel> to cancel.



Note

Changes in Special File Path Name

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.

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, click [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. Implementing changes

Press <OK> to implement the attributes change, or <Cancel> to cancel.

Using the Class Configuration Setting screen

1. Displaying the disk to change attributes

Select the disk of which you want to change attributes in the [Class Configuration Disk/Group] field.

2. 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."

3. Implementing changes

Press <OK> to implement the attributes change, or <Cancel> to cancel.

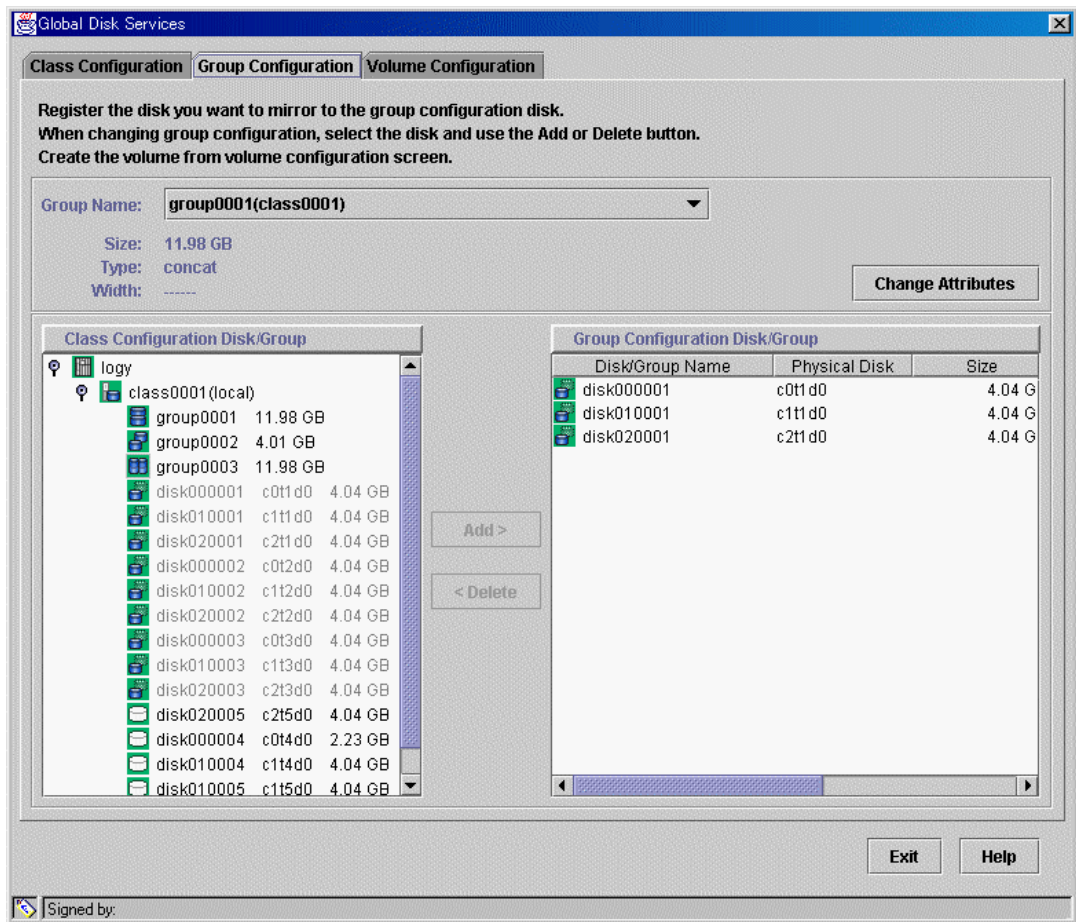
5.4.2 Group Configuration

The procedures for changing group are explained below.

1. Displaying Group Configuration screen

In the main screen [Settings] menu, click [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."

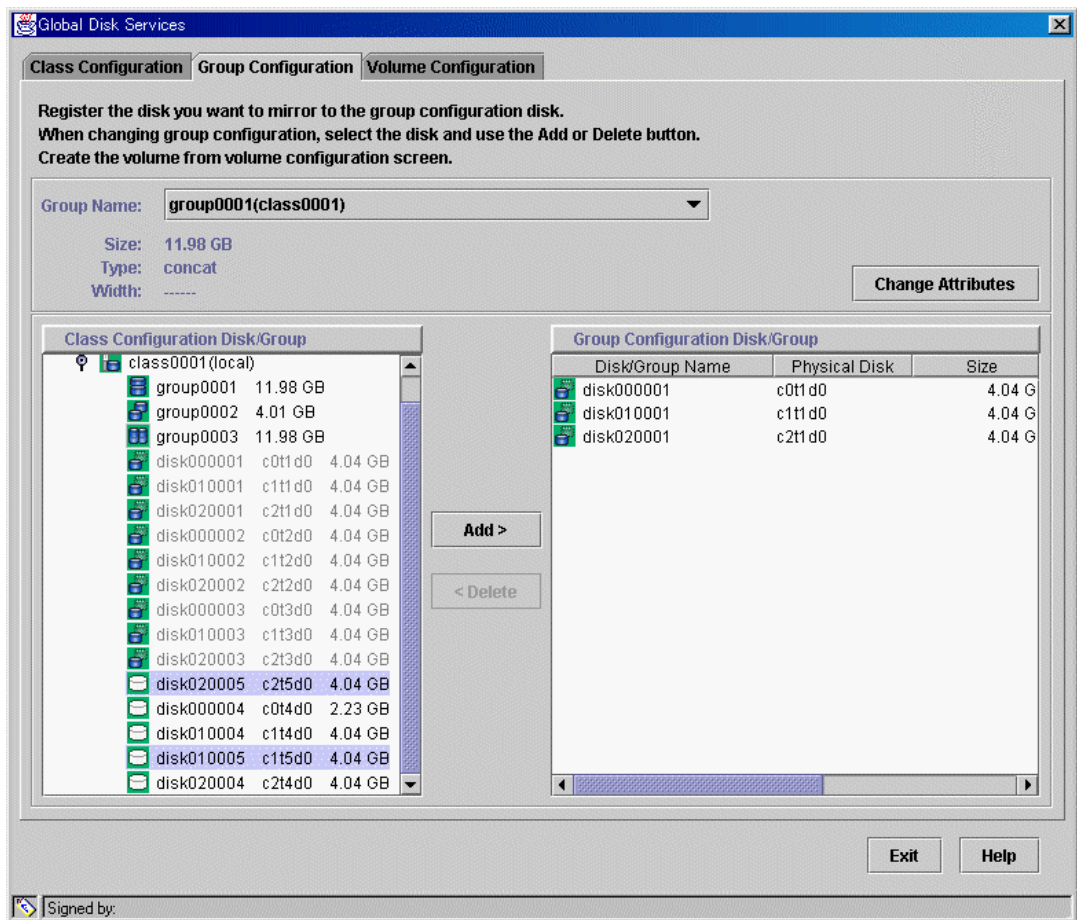
3. Changing configuration

Follow the procedures below to connect disk/lower level group to a group, or to disconnect disk/lower level group from a group.

a) Connecting disk/lower level group

1. Select the disk/lower level group you want to add to group from [Class Configuration Disk/Group] field .
2. Press <Add>.

Figure: Connecting disk/lower level group



b) Disconnecting disk/lower level group

1. Select the disk/lower level group you want to disconnect from [Group Configuration Disk/Group] field.
2. Press <Delete>.

4. Completing the change

If you have no other changes to make, press <Exit>.

Changing Group Attributes

You can only change "group name."

When changing the group attributes, there are two procedures to choose from:

- using main screen [Operation] menu
- using <Change Attributes> button in the [Group Configuration] screen

Using main screen Operation menu

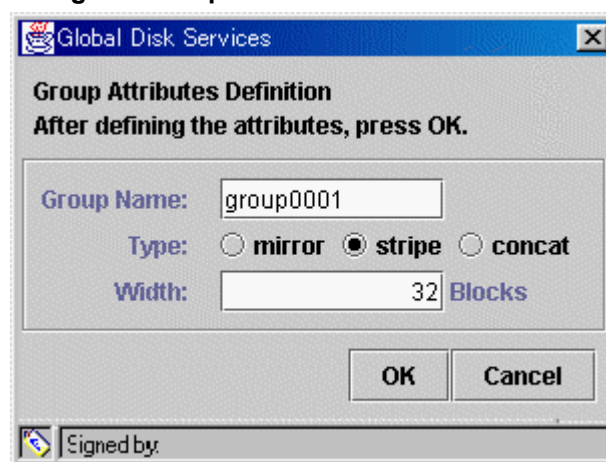
1. Selecting group to change

Display the group you want to change in the main screen Configuration Tree field. Click the icon and select the group you want to change.

2. Displaying Group Attributes Definition screen

In the main screen [Operation] menu, click [Change Attributes]. [Group Attributes Definition] screen appears.

Figure: Group Attributes Definition Screen



3. Changing group name

You can only change the group name.

Press <OK> to implement the change, or <Cancel> to cancel.



[See](#)

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

Using Change Attributes button in the Group Configuration screen

1. Displaying Group Attributes Definition screen

Press <Change Attributes>, and [Group Attributes Definition] screen appears.

2. Changing group name

Change the group name.

Press <OK> to implement 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

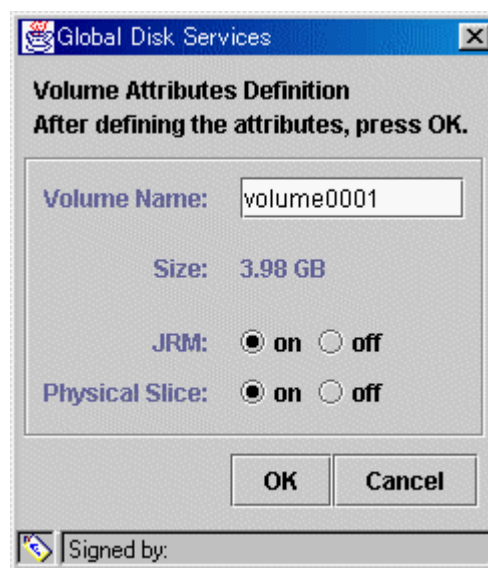
1. Selecting volume to change

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, click [Change Attributes]. [Volume Attributes Definition] screen appears.

Figure: Volume Attributes Definition Screen



3. Changing attributes

You can change the following attributes:

- Volume Name
- JRM (on/off)
- Physical Slice (on/off)



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

4. Implementing changes

Press <OK> to implement 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

This section explains procedures to remove file system, volume, group, and class.

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, click [File System Configuration]. File System Configuration screen appears.

2. Selecting group/disk

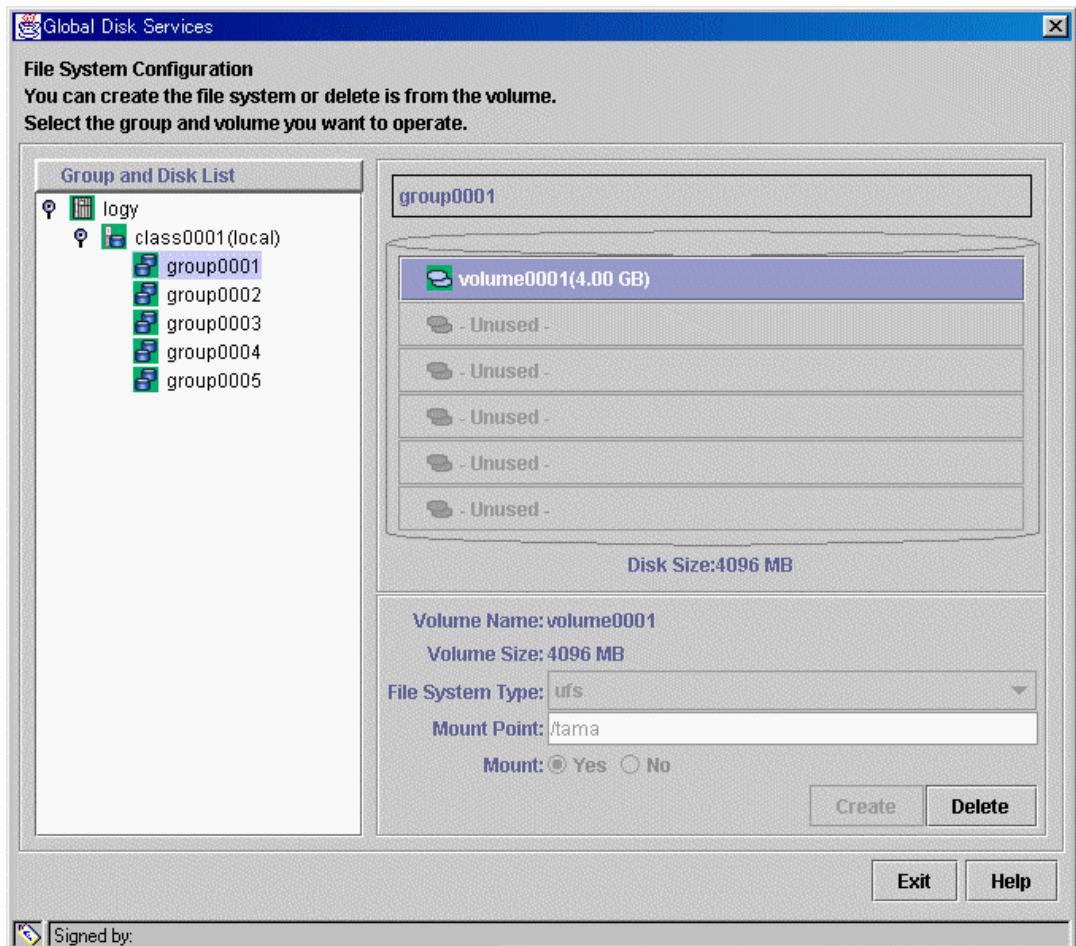
In the Group/Disk List, select a group or disk with which you want to perform an operation.

3. Removing file system

Select the volume from which you want to remove the file system, and press <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

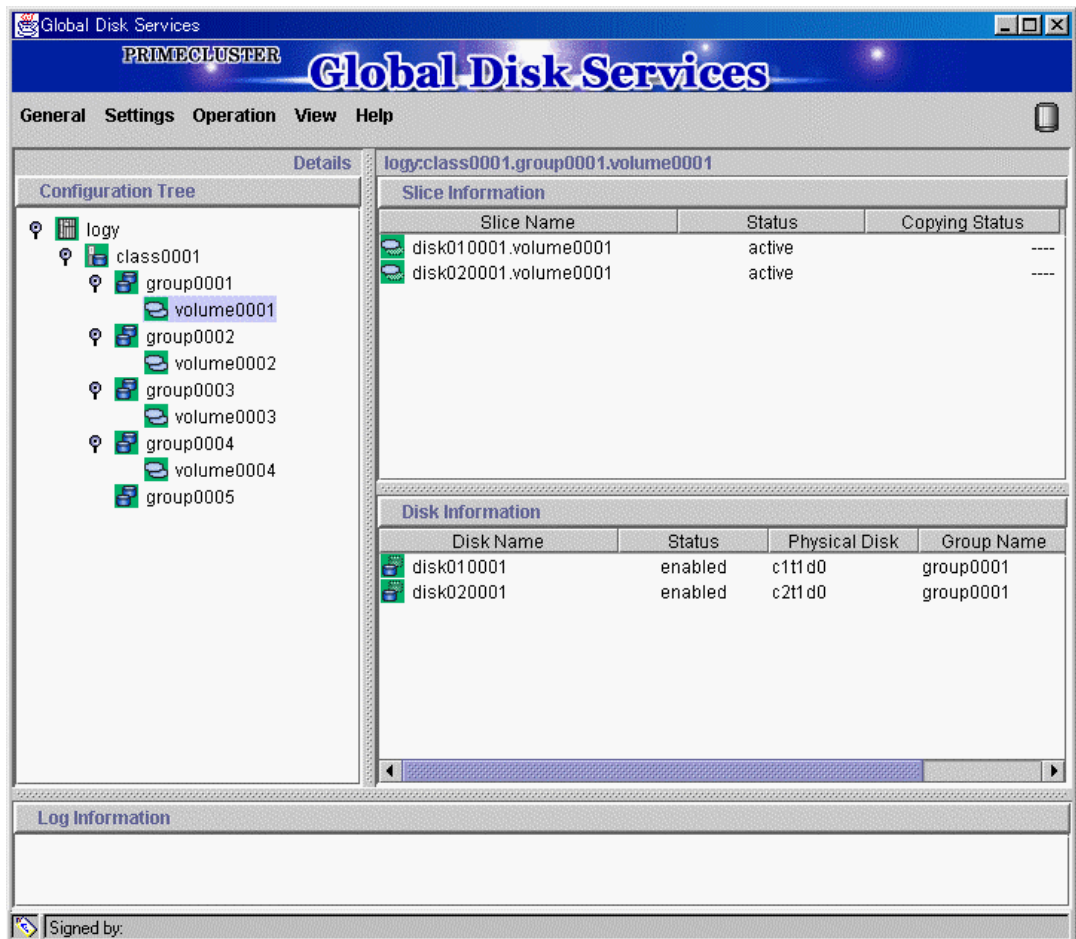
The procedures for removing a mirror volume are explained below.

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 click [Attach Slice]. The detached slice will be attached again.

For information on Attach Slice operation, see "5.3.2 Backup."

3. Removing volume

The procedures for removing a volume are explained below.

(1) Displaying the Volume Configuration screen

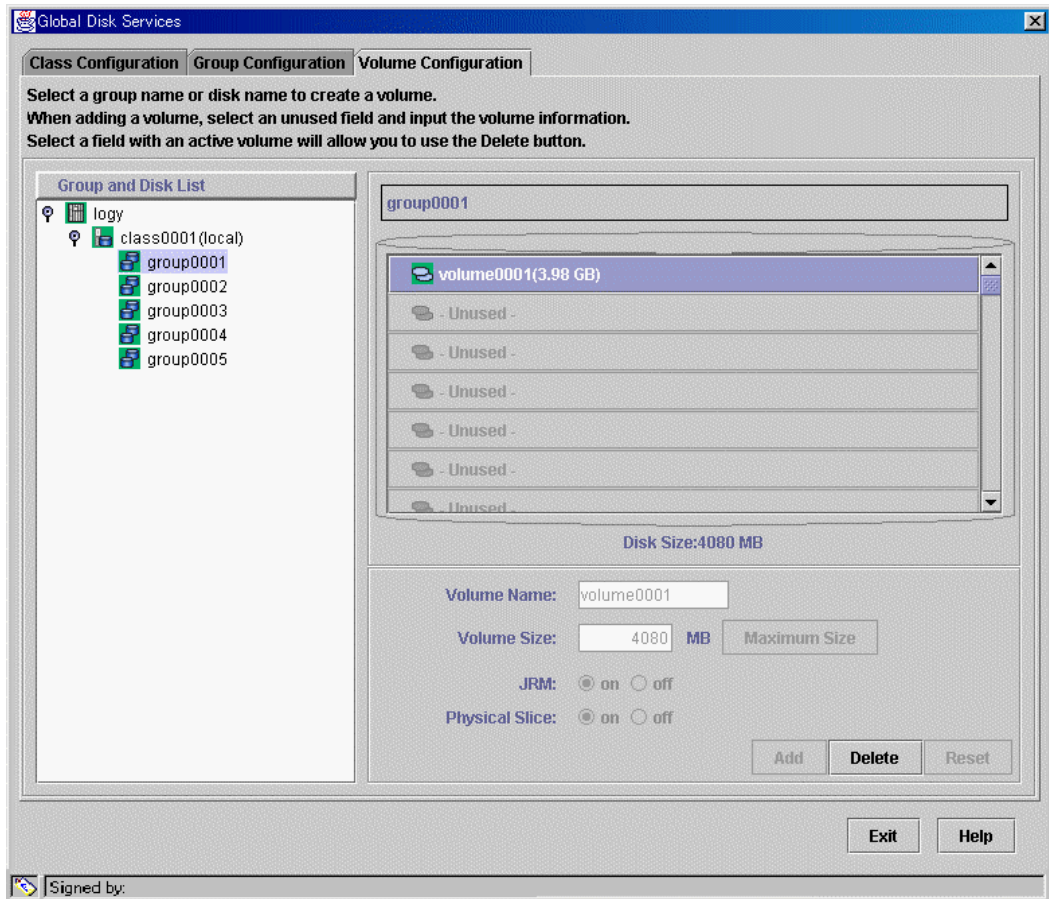
In the main screen [Settings] menu, click [Volume Configuration]. Volume Configuration screen appears.

(2) Removing a volume

Select the volume you want to remove, and press <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, click [Group Configuration]. [Group Configuration] screen appears.

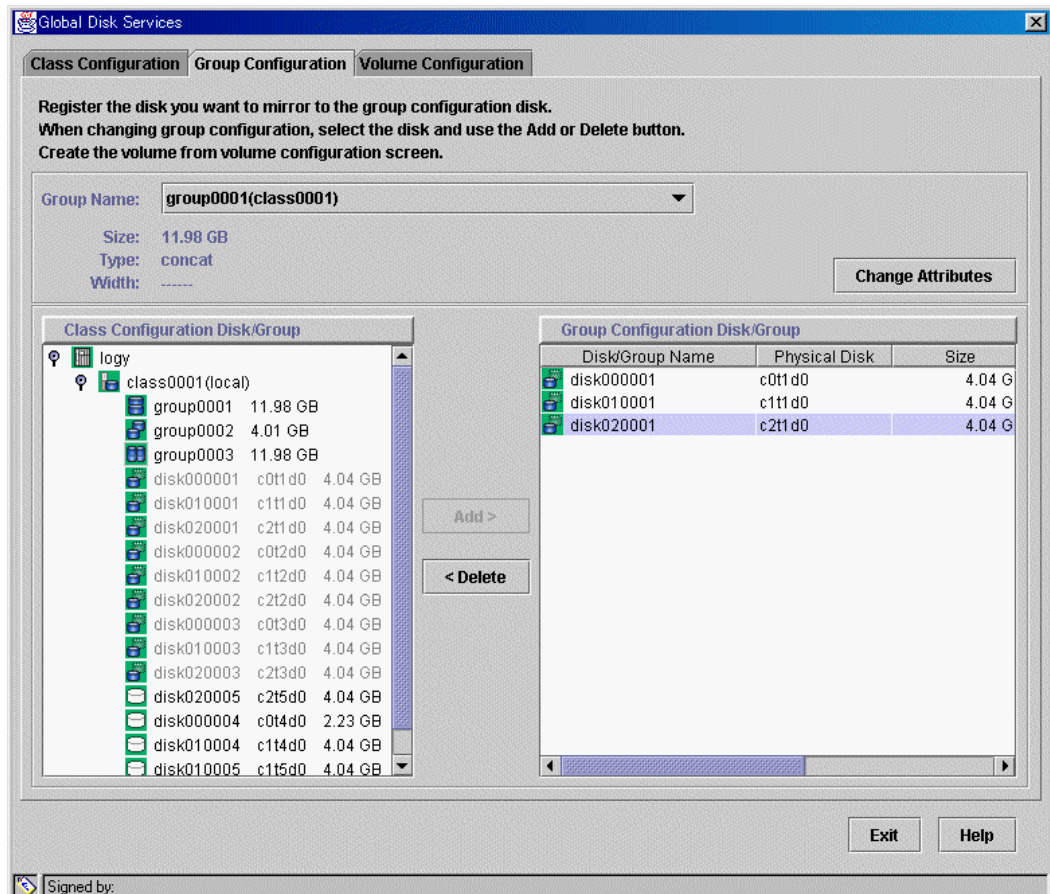
(2) Selecting group to remove

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 press <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.

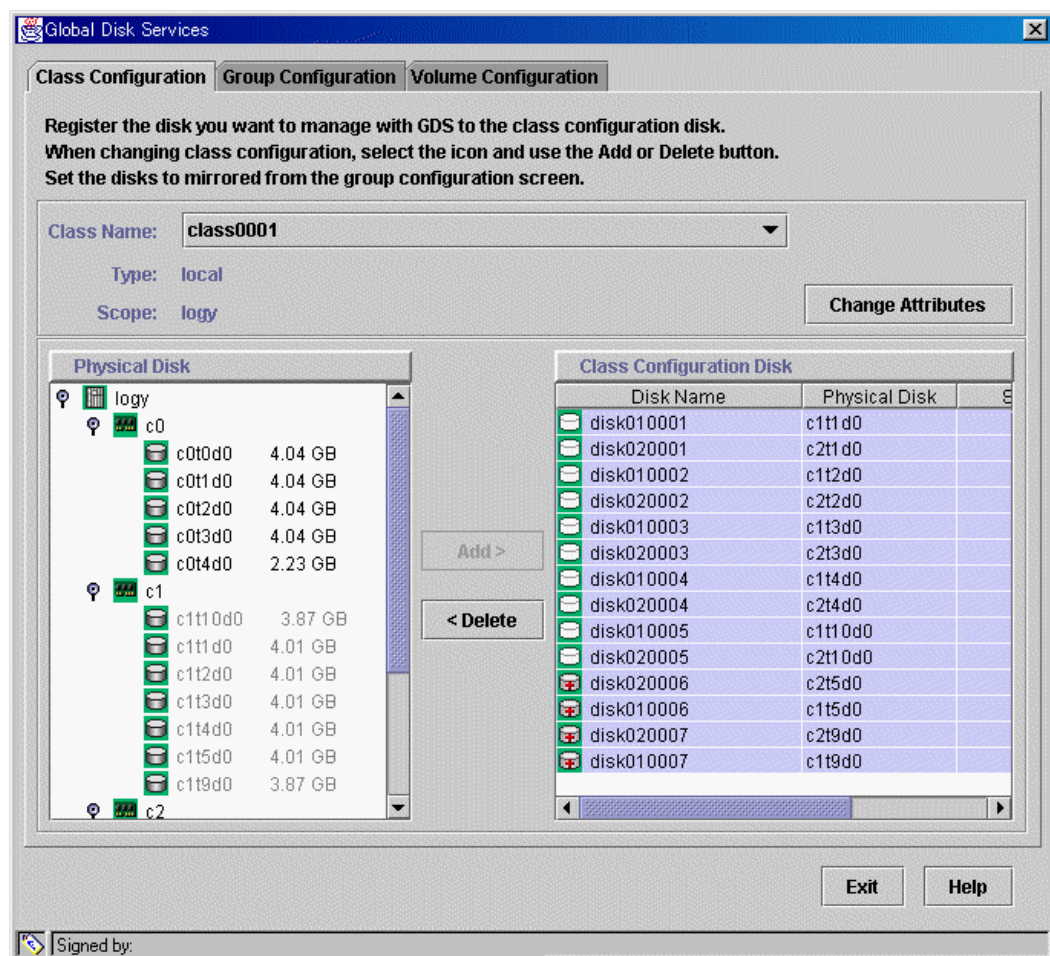
(2) Selecting class to remove

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

In this section, how to unmirror a system disk is explained.
On the [Settings] menu, select [Unmirror System Disk].



Note

When You Cannot Unmirror the System Disk

You cannot unmirror the system disk under the conditions listed below. Try unmirroring the system disk after restoring the disk from these conditions.

- A slice is temporarily excluded due to "Detach Slice" command.
- A disk is excluded due to "Swap Physical Disk" command.
- A disk is in disabled status.
- All disks in each group contain a slice that is not active.



Note

When Registering an Additional Disk with Root Class

If additional disk(s) have been registered with the root class and connected to an additional group by using commands or [Class Configuration] and [Group Configuration], ensure to unmirror the system disk after deleting the additional disk(s) from the root class.



Note

For Safe Unmirroring of System Disk

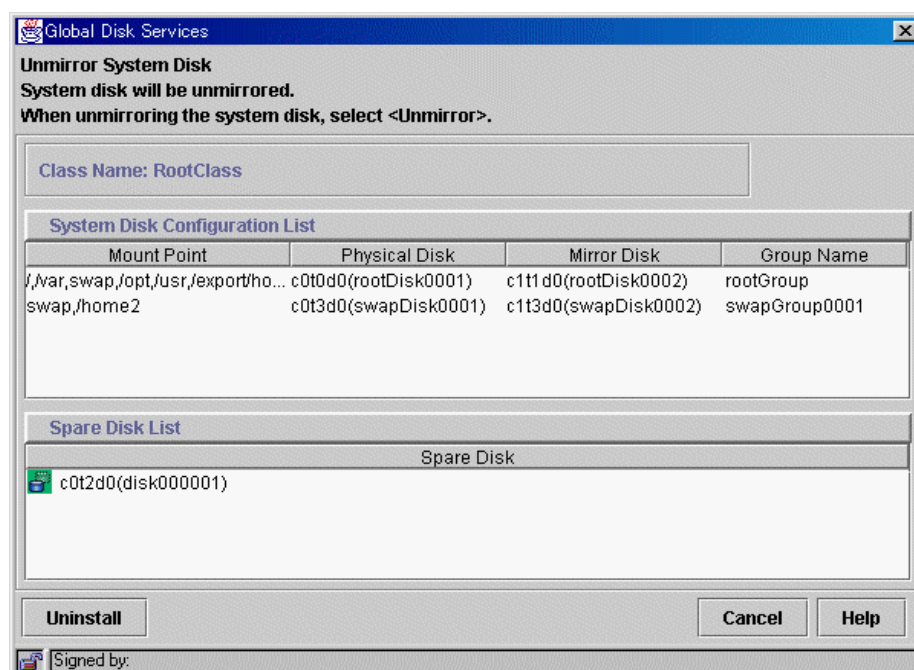
To safely unmirror the system disk, exit all active applications before proceeding to cancel system disk settings.

After unmirroring the system disk is complete, promptly reboot the system in multi-user mode.

1. Confirming system disk configuration

Mirroring status of system disk will be displayed.

Figure: Unmirror System Disk




When unmirroring is complete, the disk displayed in the "Mirror Disk" field will be excluded, and the disk displayed in the "Physical Disk" field will be used as the system disk. When canceling the mirroring of system disk, click <Uninstall>. When you do not want to unmirror, click <Cancel>.

2. Confirming the unmirroring of system disk

If you click <Uninstall> in the Unmirror System Disk screen, the screen below appears.

Figure: Unmirroring System Disk Confirmation Screen

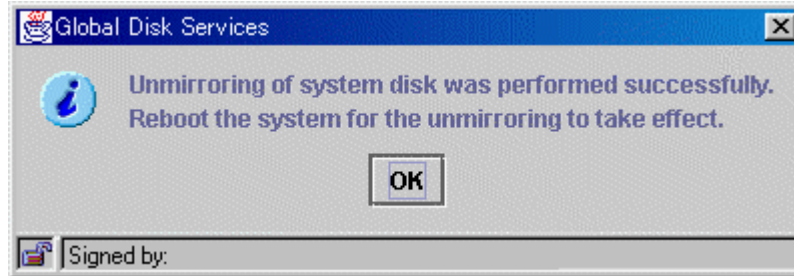


Click <Yes> to continue. Click <No> to cancel the unmirroring of system disk.

3. Unmirroring System Disk Completion Screen

If you click <OK> 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>. After unmirroring of system disk is complete, promptly reboot the system in multi-user mode.



Note

When does the unmirroring of system disk become valid?

Unmirroring of system disk will take place when the system is activated in multi-user mode.

Appendix A General Notes

A.1 Rules

A.1.1 Object Name

You can name the following objects:

- Class
- Disk
- Group
- Volume

When naming objects, you can use alphanumeric characters, -(minus), and _(underscore). However, you cannot start the name with a -(minus), or an _(underscore). The first letter should always be an alphanumeric character. You may use up to 32 characters.

Disk class name must be unique within the entire system (when using cluster, within the entire cluster system). As such, you can not create more than one disk class with the same name.

Since other object names are unique within the class, creating more than one object with the same name in a class will result in an error.

GDS assigns slice names by combining the name of disk, group and volume, to which the slice belongs. The slices are named according to the following rules.

When slice is a mirror slice, and it exists in a disk which is connected directly to the highest level mirror group:

disk_name.volume_name

When slice is a mirror slice, and it exists in a lower level group which is connected directly to the highest level mirror group:

lower_level_group_name.volume_name

When slice belongs to a stripe volume:

the_highest_level_stripe_group_name.volume_name

When slice belongs to a volume created within the highest level concatenation group:

the_highest_level_concatenation_group_name.volume_name

When slice is a single slice:

single_disk_name.volume_name

After all, you can designate each object uniquely in the entire system by the object name and the class name to which it belongs.

A.1.2 Number of Classes

If you specify a new class name when registering a physical disk, a new class is automatically created. Although there is no limit to the number of classes you can create, you should not create more classes than necessary.



Note

Separating classes beyond what is necessary could reduce reliability. Usually there should be one class, unless you are using a cluster system, or a large-scale class.

You should separate classes only under the following conditions:

- When the number of disks within a class is 1025 or more. (You can register a maximum of 100 disks with a root class).
- When there are both disks shared in a cluster system and disks used only on a single node.
- When more than one service is defined in the cluster.
- When creating a shared class with different scopes.

If you are just adding on a disk, you do not need to separate the class. Only one root class may be created within a node.

However, there are situations as listed below, when separating classes improves the manageability. Select the appropriate configuration suitable for your system requirement.

- Separating the system disk to root class, and local disk to local class improves manageability.
- When there are more than 100 disks connected to the system, separating classes according to their physical location and the type of stored data improves manageability.

A.1.3 Number of Disks

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

- To root class, you can register up to 100 disks.
- To local or shared class, you can register up to 1024 disks.

The number of disks you can connect to one group has following limitations:

- To mirror group, you can connect a total of 8 disks and lower level groups combined. In other words, a maximum of eight-way multiplex mirroring is supported. However, be aware that the spare disk that will be automatically connected when a disk failure occurs is also included in the count.
- To stripe group, you can connect a total of 64 disks and lower level groups combined. In other words, you can stripe across a maximum of 64 columns.
- To concatenation group, you can connect a maximum of 64 disks. In other words, you can concatenate up to 64 disks.

A.1.4 Number of Groups

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

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

A.1.5 Number of Volumes

Within a group or a single disk, the number of volumes you can create has following restrictions:

- You can create a maximum of 6 volumes with their physical slice attribute set to "on".
- You can create a total of 254 volumes with their physical slice attribute set to "on" or "off".
- You can not create a volume with its physical slice set to "on" in a stripe group or a concatenation group.
- You can not create a volume with its physical slice set to "off" in root class.

Within a class, the total number of volumes you can create has following restrictions:

- For the root class, you can create up to 600 volumes
- For a local class or a shared class, you can create up to 6144 volumes

A.1.6 Number of Spare Disks

There is no limit to the number of spare disks you can register with one class.

Although there is no general rule in deciding the number of spare disks, you can designate, assigning 10% of disks and lower level groups for spare disks is recommended. In other words, one spare disk for every 10 disks or lower level groups combined, would be a good rule of thumb.

Since high availability is critical for cluster system, preparing a spare disk for shared disks in a cluster system is strongly recommended.

Also when mirroring disk arrays with spare disk function, using the function provided by the disk array is recommended.

A.1.7 Number of Keep Disks

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

A.1.8 Creating Group Hierarchy

There are seven kinds of group hierarchy configuration including when the group is not hierarchised.

higher level group <-----> lower level group
mirror group (*1)
mirror group (*1) - stripe group (*3)
mirror group (*1) - stripe group (*3) - concatenation group (*7)
mirror group (*1) - concatenation group (*5)
stripe group (*2)
stripe group (*2) - concatenation group (*6)
concatenation group (*4)

Possible operation to change the configuration of hierarchical group is listed below.

(*1) The highest level mirror group

- Disks, lower level stripe groups and lower level concatenation groups can be connected to, or disconnected from the highest level mirror group. However, you cannot disconnect these objects if it may result in a change in the status of the volume.
- If there is no volume, the group itself can be removed.
- A volume can be created or removed.

(*2) The highest level stripe group

- If there is no volume, disks and lower level concatenation groups can be connected or disconnected.
- If there is no volume, it can be connected to a mirror group.
- If there is no volume, the group itself can be removed.
- If more than one disks or lower level concatenation groups are connected to the highest level stripe group, you can create and remove a volume.

(*3) The lower level stripe group

- It can be disconnected from the highest level mirror group. However, if disconnection may result in a change in the status of the volume, you cannot disconnect it.

(*4) The highest level concatenation group

- A disk can be connected.
- If there is no volume area on the disk that was connected last, that disk can be disconnected.
- If there is no volume, it can be connected to a mirror group or a stripe group.
- If there is no volume, the group itself can be removed.
- A volume can be created or removed.

(*5) Lower level concatenation group that is connected to the highest level mirror group

- A disk can be connected.
- If more than one disk is connected, and there is no volume area on the disk that was connected last, that disk can be disconnected.
- It can be disconnected from the highest level mirror group. However, if disconnection may result in a change in the status of the volume, you cannot disconnect it.

(*6) Lower level concatenation group that is connected to the highest level stripe group

- If there is no volume within the highest level stripe group, you can connect a disk.
- If more than one disk is connected, and there is no volume within the highest level group, the disk that was connected last can be disconnected.
- If there is no volume within the highest level stripe group, it can be disconnected from the highest level stripe group.

(*7) Lower level concatenation group that is connected to a lower level stripe group

- None.

A.1.9 Exception to Proxy Configuration

Users can relate a pair of volumes or a pair of groups as a master and a proxy. However, there are following exceptions:

- Objects in root class cannot be related as a master or a proxy.
- Objects belonging to different classes cannot be related as a master and a proxy.
- Objects belonging to the same group cannot be related as a master and a proxy.
- Objects of a stripe type cannot be related as a master and a proxy.
- Objects of a concatenation type cannot be related as a master and a proxy.
- Proxy volumes that related to another master cannot be related to a different proxy
- Master volumes that are already related to a proxy cannot be related to a different master as a proxy.

A.1.10 Number of Proxy Volumes

Although more than one proxy volume can be related to one master volume, the number should be limited to meet the following conditions.

- The total number of slices comprising the master volume, and the slices comprising the proxy volume that is related to the master volume, must not exceed thirty-two.

For example, if a single volume comprises both the master and proxy volume, a maximum of 31 proxy volumes may be related to the master volume.

A.1.11 Proxy Volume Size

The size of proxy volume must be equal to that of the master volume to which it will be related.

A.1.12 Proxy Group Size

The size of a proxy group must be larger than the last block number of volumes within the master group to which the proxy is related.

A.2 Important Points

A.2.1 Physical Disk Driver

GDS allows you to manage physical disks controlled by the following device drivers. In other words, physical disks controlled by other disk drivers cannot be registered with class.

sd, ssd, dad, hddv, mphd, mplb, emcpower, emcp

Also, physical disks controlled by the following device drivers cannot be registered with root class. They must be registered with **local class** or shared class.

hddv, mphd, mplb, emcpower, emcp

When the above requirements are not met, the `sdxdisk -M` command will terminate with an error and output the following error message.

```
SDX:sdxdisk: ERROR: physical_disk_name: disk driver driver_name not supported
```

```
or,
```

```
SDX:sdxdisk: ERROR: physical_disk_name: disk driver driver_name not supported for root class
```

A.2.2 Managing System Disks

You cannot configure the following file systems in volumes registered with local class or shared class.

Disk with these file systems must be managed in root class as a system disk.

`/(root), /usr, /var, swap area`

A.2.3 Restraining Access to Physical Special File

After you complete registering the disk with GDS, you cannot access the disk using physical special files such as those listed below. This protection measure is taken so that you do not access the physical special files by mistake and disrupt mirroring.

`/dev/(r)dsk/cCtTdDs[0-7]` (for normal hard disks)

`/dev/FJSVmphd/(r)dsk/mphd/ s[0-7]` (for MPHD disks)

`/dev/FJSVmplb/(r)dsk/mplb/ s[0-7]` (for MPLB disks)

C indicates the controller; *T*, the target ID; *D*, the disk number; *I*, the MPHD instance number or the MPLB instance number; [0-7], the slice number with 1 digit from 0 to 7.

For example, if you execute the `dd` command to write to the disk with physical special file, the following error message will be displayed.

```
# dd if=/dev/zero of=/dev/rdisk/c1t1d0s0
dd: /dev/rdisk/c1t1d0s0: open: Device busy
```

Disk list output by the `format(1M)` command does not display the disk registered with GDS.

This access protection is cancelled when disk is removed from class.

It will also be cancelled when object is operated for disk swap. The access protection function will turn back on if disk recovery operation is performed on the object.

However, when using EMC PowerPath 2.x.x, this function is not effective for `emcpower` devices. You must exercise care not to access the following special physical files, since access is not restrained.

`/dev/[r]dsk/emcpowerN{a-h}`

N indicates the `emcpower` device number; {a-h}, the slice number with one alphabet from a to h.

A.2.4 Booting from a CD-ROM Device

When you cannot boot a system, for example, you may recover the system by booting from a CD-ROM device.

However, there is the following concern when booting from a CD-ROM device.

- An operation mistake can be easily caused since circumstances when booting with this method and when booting from a boot disk may alter the correlation between the disk and the device special file (`/dev/[r]dsk/cCtTdDsS`).
- Only partial mirrored disks can be written in since access to the device special file of a physical slice is not restrained. Therefore, the mirroring state could collapse.

For these reasons, you should avoid booting from a CD-ROM device unless directed in this manual.

When mounting a file system by booting from a CD-ROM device in any procedure not described in this manual from necessity, it should be mounted as read-only.

A.2.5 Initializing Disk

When registering a physical disk with disk class, all data contained in the physical disk will be lost since GDS automatically reformats the disk (Keep disk is an exception). The same applies when swapping and restoring the physical disk.

When using the GDS Management View, a dialog box reminding you that all data will be lost appears. However, please be aware that there will be no reconfirmation when actually performing the command.

Also, when registering a disk containing data with class, you must specify the disk as "keep" type and register it with root class. Unless it is a system disk, confirm that the disk is in unused status (when used as a file system, it should be in unmounted status) before registering with root class.

When not specifying as "keep" type, the disk will be initialized at the time of registration. Backup data before registration, and restore data after creating the volume.

Be aware that when a physical special file name is set for `/etc/vfstab`, you will have to change it to a special file name for mirror volume.

```
/dev/sfdsk/class_name/rdsk/volume_name (character special file)  
/dev/sfdsk/class_name/dsk/volume_name (block special file)
```

A.2.6 Disk Size

Out of the physical disk area, the capacity available for creating volumes equals the physical disk size rounded down to the cylinder boundary, minus the private slice size. This size is called the available size of disk.

Private slice is an area that GDS reserves on each disk for storing data, such as configuration information and the JRM log.

When a physical disk is registered with a class, a private slice is reserved. Disks within the same class will all have the same private slice size. The size is determined by the size of the disk that is first registered with the class.

The size of a private slice can be estimated as below. The private slice size never exceeds the following estimated values.

- When the size of the disk registered first with the class is 10GB and below:
32MB
- When the size of the disk registered first with the class is over 10GB:
32MB + (size calculated by rounding-up 0.1% of the disk capacity to the cylinder boundary)

The size of a log for JRM is determined by the size of a volume. As a result, the log area for JRM may not be reserved because the private slice size becomes insufficient when a mirror volume is created registering a disk larger in size than the first registered disk with the class. In such a case, the volume is incapable of using JRM. Therefore you are recommended to first register a disk of the maximum size with the class.

A.2.7 Volume Size

The size of a volume is automatically adjusted conforming to the following conditions.

- When creating a volume into the group that a keep disk belongs to
A volume is created in the size calculated by rounding-up the size specified when creating the volume to a multiple of the keep disk's cylinder size.
(Example)
If the keep disk's cylinder size is 8 MB (= 16,384 blocks) and 20 MB (= 40,960 blocks) is specified to the size of a volume when it is created, the size is rounded to a multiple of 16,384 blocks, and a volume with 49,152 blocks (= 24 MB) is created.
- When creating a volume into a stripe group
A volume is created in the size calculated by rounding-up the size specified when creating the volume to a common multiple of (stripe width) * (stripe column number) and the cylinder size (32,768 blocks = 16 MB).
(Example)
If the stripe width is 32, the number of the stripe columns is 3 and 20 MB (= 40,960 blocks) is specified to the size of a volume when it is created, the size is rounded to a common multiple of 96 (= 32 * 3) blocks and 32,768 blocks, and a volume with 98,304 blocks (= 48 MB) is created.
- When creating other volumes
A volume is created in the size calculate by rounding-up the size specified when creating the volume to a multiple of the cylinder size (32,768 blocks = 16 MB).
(Example)
If 20 MB (= 40,960 blocks) is specified to the size of a volume when it is created, the size is rounded to a multiple of 32,768 blocks, and a volume with 65,536 blocks (= 32 MB) is created.

A.2.8 Spare Disk

Spare disk provides temporarily support to maintain mirroring status until the crashed disk is restored. If the spare disk that was connected in place of the failed disk also fails, it will not be replaced by another spare disk, even if there is another spare disk registered. Promptly locate the defective part and restore the disk.

After restoring the disk and completing synchronization copying, the spare disk is automatically disconnected, and will prepare for the next failure.

If a disk connected to a lower level group connected to a mirror group, or a disk connected to a group that is even lower in the hierarchy fails, spare disk will not replace the failed disk, but will replace the lower level group connected to the mirror group. For example, if a disk and a concatenation group is mirrored each other, and one of the disks connected to the concatenation group fails, the disk and a spare disk will be mirrored each other until the failed disk is restored.

A.2.9 System Disk Mirroring

You can mirror the system disk containing the operating root file system by registering the system disk with root class. When registering a system disk, you must specify "keep" for its disk type.

You cannot register system disk with local class and shared class.

When mirroring a system disk, GDS additionally sets the OpenBoot nvramrc parameter. It will also set use-nvramrc? parameter to "true" so that the command registered with NVRAMRC will automatically be executed when booting the system.

However, when a command is already registered with NVRAMRC and use-nvramrc? is set to "false", a following warning message will be displayed, and use-nvramrc? will not be changed to "true".

WARNING: cannot set use-nvramrc? parameter to true, NVRAMRC not empty

When use-nvramrc? is set to "false", removing the command registered with NVRAMRC is recommended.

When the warning message is displayed, follow the procedures below and set use-nvramrc? to "true."

1. Confirm the contents of NVRAMRC. When the setting is correct, change the use-nvramrc? parameter to "true" with eeprom(1M).
2. If the contents of NVRAMRC is invalid, remove the command registered with NVRAMRC, and reboot the system. use-nvramrc? will be automatically changed to "true".



Information

- How to confirm use-nvramrc? parameter:

```
# eeprom use-nvramrc?  
use-nvramrc?=false
```

- How to change and confirm use-nvramrc? parameter:

```
# eeprom use-nvramrc?=true  
# eeprom use-nvramrc?  
use-nvramrc?=true
```

- How to confirm the command registered with NVRAMRC:

```
# eeprom nvramrc  
(Registered command will be displayed.)
```

- How to remove and confirm the command registered to NVRAMRC:

```
# eeprom nvramrc=  
# eeprom nvramrc  
nvramrc: data not available.
```

A.2.10 Keep Disk

By specifying a disk as "keep" type and registering it with root class, you can retain the disk data while configuring mirroring. However, you cannot register the disk under the following conditions. In such case, change the disk configuration and try registering again.

- There is no unused slice number (size=0) in the disk.
- There is not sufficient free space area or swap area on the disk for private slice.

For details on free space area (or swap area), see "A.2.6 Disk Size".
You cannot register a keep disk with local class or shared class.

A.2.11 GFS Partition on a Keep Disk

When a keep disk registered with the root class contains a partition used by a GFS Local File System, you must make the following modifications on the settings for the relevant partition prior to rebooting the system after executing [System Disk Settings] in the GDS Management View or the `sdxroot -M` command.

- Changing the partition information settings
For details on the method of changing the settings, see "PRIMECLUSTER Global File Services Configuration and Administration Guide" or the manual page of the `sfxadm(1M)` command.
- Modifying the `/etc/vfstab` file
You must change the device special file name from that for the physical slice to that for the logical volume.

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

A.2.12 Creating a Snapshot by Detaching Slice

A slice can be detached only from a mirror volume with a physical slice. So if a disk is not connected directly to a mirror group, you cannot create a snapshot by detaching a slice.

A.2.13 The Difference between a Mirror Slice and a 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 slices are 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 results in that you cannot continue accessing a master volume where all slices comprising the master volume are abnormal even if proxy volumes are normal. The purpose of proxy volumes is to create snapshots (saved copies from the master volume at certain moment) for a different business running concurrently with the core business but not to improve the data redundancy of the master volume used in the core business operation.

While the function of GDS Snapshot creating snapshots by detaching slices is a by-product of mirroring, the function itself is the primary purpose of GDS Snapshot by proxy volumes.

Therefore, the application of proxy volumes provides more flexible disk configurations and business styles to snapshot management.

A.2.14 Just Resynchronization Mechanism (JRM)

There are three types of Just Resynchronization Mechanism (JRM): for volumes, for slices and for proxies.

JRM for Volumes

JRM for volumes speeds up the resynchronization process when booting the system after a system panic or the like. GDS records the changed portion in the private slice. The resynchronization copy performed at rebooting after an unexpected system failure copies the only portion that was written during the system down to realize high-speed resynchronization and minimize the load of copy processing.

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.

JRM for Proxies

JRM for proxies speeds up the resynchronization process when joining a parted proxy again to the master. GDS records the changes made on the master and proxy on the memory while the proxy is parted. The resynchronization copy performed when re-joining the parted proxy copies only the updated portions to realize high-speed resynchronization.

Unlike JRM for volumes or slices, JRM for proxies is always enabled and cannot be turned off.

However, if there is an occurrence of a system shutdown while the proxy is parted, just resynchronization is not performed when the proxy is later re-joined. Resynchronization is performed by copying the entire data, not only the updated portions.

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.

A.2.15 Object Operation When Using Proxy

Operation to change the class attributes related to master volume and proxy volume, operation to change group configuration or attributes, and slice operation cannot be performed. Possible object operations that are related to master and proxy are given below.

- Master volume can be activated and stopped.
- Proxy volume can be activated and stopped, if separated from the master.
- New volumes can be created within a group that is neither a master or proxy group, and existing volumes can be removed as long as they are not a master or a proxy volume.
- Hot spare function is valid for a group with a master volume. Spare disk cannot be connected to a group with a proxy volume.

The following operations can be performed, but the operation will result in an error if copying is in process, or there is an EC session, a BCV pair, or an SRDF pair, between the master and the proxy.

- Attributes of master and proxy volumes can be changed with the `sdxattr -V` command.
- The disk connected to a group with a master volume and a proxy volume can be prepared for swapping with the `sdxswap -O` command. After the swapping is complete, the disk will be ready for use with the `sdxswap -I` command.
- Synchronization copying of the master volume or the proxy volume separated from the master volume can be started, cancelled, interrupted and resumed, and the parameter of the synchronization copying changed. However, synchronization copying between volumes or its parameter cannot be changed.
- Master volume can be restored with the `sdxfix -V` command.

A.2.16 Snapshot by Cooperation with Advanced Copy of ETERNUS3000 or ETERNUS GR700 series

By cooperation with Advanced Copy function of ETERNUS3000 or GR700 series storage system, a proxy configuration realizes synchronization copying without imposing a load on the server accessed by the business application, or the SAN.

When Advanced Copy function is available, synchronization copying of master to proxy, or proxy to master is performed using that function. However, soft copy function (copy function of GDS driver running on the server) will be performed in the following cases.

- Soft copy function is specified to be used.
- The volume to which the data is copied is mirrored.
- Lower level groups are connected to the mirror group that contains the master volume or the proxy volume, but no disk is directly connected to the mirror group.

EC is used for synchronization copying from a master to a proxy after joining or rejoining them, copying for maintaining synchronization, and recording the portion updated while a master and a proxy are parted.

OPC is used for synchronization copying, instant snapshot processing, and restoring of a master using the proxy data.

When both EC and OPC is available as with GR740, EC has precedence over OPC.

Once an EC session is stopped, you can not use the EC function afterwards. EC session will be stopped in the following cases.

- Instant snapshot is created.
- Master data is restored using proxy.
- The master and the proxy are re-joined using the soft copy function which is specified to be used.

If you wish to use the EC function after performing the above operation, you must first cancel the master/proxy relation, and relate the two again.

The copying process being performed can be checked in the CPTYPE field using the `sdxinfo` command.

A.2.17 Instant Snapshot Created by OPC

Instant snapshot created by OPC is only available with the following storage systems that support the OPC function: ETERNUS3000, GR720, GR730 and GR740.

A.2.18 To Use EMC Symmetrix

When using an EMC's Symmetrix storage unit, EMC's PowerPath is essential. GDS is capable of managing EMC Symmetrix emcpower devices.

The following devices are exceptions to devices that can be managed by GDS.

- native devices configuring emcpower devices
- BCV (Business Continuance Volume) devices
- SRDF target (R2) devices
- GateKeeper devices
- CKD (Count Key Data) devices
- VCMDB (Volume Configuration Management Data Base) used by Volume Logix

After completing the configuration of these devices and EMC software, follow the procedure below and describe a list of devices excluded from disk management by GDS in the `/etc/opt/FJSVsdx/lib/exdevtab` file (referred to as the Excluded Device List). The Excluded Device List must include all disks that cannot be managed by GDS in addition to the devices as above.



Note Snapshot by Cooperation with EMC TimeFinder or EMC SRDF

When using EMC TimeFinder or EMC SRDF in a proxy configuration, the BCV and R2 devices connected with the proxy group may not be described in the Excluded Device List. However, native devices configuring such devices must be described in the list. For details see "A.2.19 Snapshot by Cooperation with EMC TimeFinder or EMC SRDF".



Note In the Case of a PRIMECLUSTER System

- The Excluded Device List must be created on all nodes configuring the cluster in the case of a PRIMECLUSTER System.
- These devices cannot be the targets of Automatic Resource Registration (Automatic Configuration), and they must also be described in the Excluded Device List for Automatic Resource Registration prior to conducting it. For details see "PRIMECLUSTER Cluster Foundation Configuration and Administration".
- When an R2 device of the SRDF pair is not described in the Excluded Device List, such a pair must be split prior to conducting Automatic Resource Registration.
- When the Excluded Device List is edited, Automatic Resource Registration must be executed and the node must be rebooted. For details on the procedure for Automatic Resource Registration, see "PRIMECLUSTER Cluster Foundation Configuration and Administration" or the manual page of the `clautoconfig` command.

[Procedure]

- 1) The `syminq` command provided by SYMCLI is available for checking respective BCV, R2, GK, and CKD devices. Execute the `syminq` command, and describe all devices indicated as BCV, R2, GK, and CKD(`cCtTdD`, `emcpowerN`) in the Excluded Device List.
- 2) The `syminq` command is unavailable for checking devices managed under the control of Volume Logix. Ask your EMC engineer or systems administrator who configured Volume Logix about devices that cannot be managed by GDS and describe them in the Excluded Device List when you are using Volume Logix.
- 3) Describe other devices that cannot be managed by GDS, if any, in the Excluded Device List in addition to the devices included in the list in step 1 and step 2.

You are recommended to suffix such tags as "PP", "BCV", "R2", "GK", "CKD" and "VL" to the device names for Excluded Device List management efficiency. A device name and a tag must be separated by one or more spaces.

The Excluded Device List should appear as follows.

```
# cat /etc/opt/FJSVsdx/lib/exdevtab
#device(full path) type
/dev/rdsd/emcpower7c BCV
/dev/rdsd/emcpower8c BCV
/dev/rdsd/emcpower9c GK
/dev/rdsd/emcpower10c GK
/dev/rdsd/emcpower11c CKD
/dev/rdsd/emcpower12c R2
/dev/rdsd/c2t0d10s2 PP
/dev/rdsd/c3t0d10s2 PP
...
/dev/rdsd/c4t0d31s2 PP
/dev/rdsd/c4t0d32s2 PP
#
```



Information **exdevtab.sh**

A sample script, `/etc/opt/FJSVsdx/bin/exdevtab.sh`, is provided in order to simply create an Excluded Device List, `/etc/opt/FJSVsdx/lib/exdevtab`.

When utilizing `exdevtab.sh`, use an editor and edit the following two parameters (symlinq command and powermt command paths) in `exdevtab.sh` for them to work in your environment.

```
SYMLINQ=/usr/symcli/bin/symlinq
POWERMT=/etc/powermt
```

Executing `exdevtab.sh` creates the descriptions of native devices of emcpower devices and respective BCV, GateKeeper, and CKD devices in the Excluded Device List. R2 devices, and devices under the control of Volume Logix, are not included. According to need, edit `exdevtab.sh` in advance or describe any additional disks that cannot be managed by GDS taking step 1) to step 3).



Note **Kernel Parameters**

- When an emcpower device is managed by GDS using EMC PowerPath 1.5.0.x, add the following line in the `/etc/system` file if the file is not including it.

```
forceload: drv/emcpower
```

- When an emcpower device is managed by GDS using EMC PowerPath 2.0.2 or after, add the following lines in the `/etc/system` file if the file is not including them.

```
set lwp_default_stksize=0x6000
set rpcmod:svc_default_stksize=0x6000
```

- When an emcpower device is managed by GDS using EMC PowerPath 2.1.1 or after, add the following lines in the `/etc/system` file if the file is not including them.

```
set lwp_default_stksize=0x6000
set rpcmod:svc_default_stksize=0x6000
set emcp:bPowerIgnoreFEXCL=1
```

- When using EMC PowerPath 3.x, no kernel parameter setting is required.



Note **Operation in the GDS Management View**

Although native devices (cCtTdD) configuring emcpower devices may be displayed in a list of physical disks on the Class Configuration screen of the GDS Management View, such devices should not be selected. Also emcpower devices (emcpowerN) may be displayed in the System Disk Settings screen, such devices should not be selected



Information Devices Described in exdevtab

A disk device described in the exdevtab file cannot be registered with a class.

If an attempt to register the disk device included in the exdevtab file is made, an error message "ERROR: *physical disk name*: no such device" is output.

A.2.19 Snapshot by Cooperation with EMC TimeFinder or EMC SRDF

The synchronization copying between a master and its proxy imposes almost no load on the servers and the SAN by using the EMC TimeFinder or EMC SRDF to use hard copy function of EMC Symmetrix.

TimeFinder and SRDF are available to copy data of an entire physical disk to another physical disk, but they are not available to copy data of a disk slice to the other disk area. Therefore, to utilize TimeFinder or SRDF with a proxy configuration, a pair of groups instead of a pair of volumes must be related as a master and a proxy. If a proxy volume instead of a proxy group is specified as a target of the operation to part, rejoin or restore, the operation results in an error, while utilizing TimeFinder or SRDF.

Configuring to meet the following requirements should be done before joining a master group and proxy group.

- To utilize TimeFinder,
 - 1) One of disks connected to the master group must be the standard device that has been registered with the device group on all of the nodes within the class scope.
 - 2) The disk connected to the proxy group must be a BCV device which is associated with the same device group as described in 1) on all nodes within the class scope.
- To utilize SRDF:
 - 1) One of disks connected to the master group must be the source (R1) device that has been registered with the device group on all of the nodes within the class scope.
 - 2) A disk connected to the proxy group must be the target (R2) device that is paired with the source (R1) device as above on all of the nodes within the class scope.

Moreover, ensure the following points about management of BCV devices, source (R1) devices and target (R2) devices, which are used with a proxy configuration.

- GDS configuration databases cannot be stored in BCV devices and target (R2) devices since the devices are overwritten by data in copy source disks. Therefore, GDS does not regard BCV devices and target (R2) devices as "disks that can be accessed normally" described in "[Explanation]" of "(1) Class becomes closed status during operation" in "F.1.4 Class Status Abnormality."
- The BCV and target (R2) devices connected to the proxy group may not be described in the Excluded Device List. However, native devices configuring such devices must be described in the list. For details on the Excluded Device List, see "A.2.18 To Use EMC Symmetrix."
- When the target (R2) device of the SRDF pair is not described in the Excluded Device List, such an SRDF pair must be split prior to conducting Automatic resource registration. For details on the procedure for Automatic resource registration, see "PRIMECLUSTER Cluster Foundation Configuration and Administration".
- When the target (R2) device of the SRDF pair is registered with a GDS class, such an SRDF pair must be split in advance.
- The BCV, source (R1), and target (R2) devices used in a proxy configuration should not be operated using the SYMCLI commands.
- If the master and proxy are parted forcibly while copying by TimeFinder or SRDF is in process, a message informing device abnormality is submitted to EMC's customer support center.

When the hard copy function is available, synchronization copying from a master to a proxy is performed using that function. However, the soft copy function (copy function of a GDS driver running on the server) will be used in the following cases.

- The soft copy function is specified to be used.
- A pair of volumes instead of a pair of groups is related as a master and a proxy.
- The configuration of the proxy group to which the data are copied is a mirroring configuration.
- A proxy volume with a different physical slice attribute from the master volume is created into a proxy group.
- Lower level groups are connected to the master group or the proxy group, but no disk is directly connected to the group.
- A disk in the same size as the disk to which the data are copied is not connected to the master group from which the data are copied.

TimeFinder and SRDF are used for synchronization copying from a master to a proxy after joining or rejoining them, copying for maintaining synchronization, and recording the portion updated while a master and a proxy are parted.

If both TimeFinder and SRDF are available, TimeFinder has precedence over SRDF.

Once the BCV pair or SRDF pair is cancelled, the TimeFinder function or SRDF function is not used. BCV and SRDF pairs will be cancelled in the following cases.

- Master data are restored using the proxy.
- The master and the proxy are re-joined using the soft copy function which is specified to be used.

If you wish to use the TimeFinder or SRDF function after performing the above operation, you must first cancel the master-proxy relation and relate them again.

The copying process being performed can be checked in the CPTYPE field using the `sdxinfo` command.

A.2.20 Ensuring Integrity of Snapshot Data

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

To ensure the integrity of your snapshot data, you must stop the application that is accessing the volume in advance. After creating the snapshot, start the application again. For example, when using the volume (master volume) as a file system such as Global File Services or `ufs`, unmount the volume with the `unmount(1M)` command before creating a snapshot, and mount it with the `mount(1M)` command afterwards. This way, you can ensure the integrity 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.

As an example, see "D.16 Online Backup Making Using Command".

A.2.21 Data Consistency at the time of Simultaneous Access

When the same block within a volume is accessed simultaneously from multiple nodes, data consistency is maintained by access exclusion control performed by the application that accesses the shared disk simultaneously.

A.2.22 Volume Access Mode

There are two types of volume access modes; "Default Access Mode" which is set by default as an access mode attribute, and "Current Access Mode" which is set for a volume that is activated. "Current Access Mode" is valid only while the volume is activated, and will become invalid when the volume is stopped. When the volume is restarted, it will start in "Default Access Mode," except for when the access mode is specified at the time of restart.

For example, if you wish to usually use the volume in "Read and write possible" mode, and temporarily switch to "Read only possible" mode, set the access mode attribute to "rw," and use the `sdxvolume -N` command with `-e mode=ro` option as necessary.

A.2.23 Configuration Operation in Cluster System

When using cluster system, objects cannot be operated under the following conditions.

- When the cluster control is not operating, you cannot perform any operations on shared objects. The executed command will terminate with an error. First, activate the cluster control.
- It is possible to perform operations on local objects when the cluster control is not operating. However, to avoid contradictions that may occur by changing configuration, activate the cluster control before performing any operation on local objects.

A.3 General Points

A.3.1 Spare Disk Size

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

A.3.2 Guidelines for Mirroring

Important points to keep in mind when mirroring data are listed below.

- Connecting disks and lower level groups with the same available sizes to mirror group is recommended.
The available size of mirror group (the capacity available as volume) is the same as the available size of the smallest disk or lower level group that is connected.
When connecting disks or lower level groups with different available sizes to a mirror group, you will only be able to use the capacity of the smallest disk or lower level group. For example, if a 4 GB disk and a 9 GB disk is connected to one mirror group, only 4 GB out of the 9 GB disk will be accessible.
- Mirroring disks with similar performance specifications, or groups with the same configuration (including the performance specifications of disks that is connected to the group) is recommended.
When mirroring disks with different performance specifications such as revolution speed, the read performance becomes unbalanced and the write performance will depend on the slower disk performance.
The same applies when mirroring disks and groups, or when mirroring groups with different configuration.

A.3.3 Guidelines for Striping

Important points to keep in mind when striping data are listed below.

- In order to improve I/O performance with striping, it is necessary to adjust the stripe width and the number of stripe columns depending on the way an application accesses the disk. If the striping configuration is not appropriate, you can not gain much performance improvement. And, depending on the way an application accesses the disk, the performance may not improve even after adjusting the stripe width or the number of stripe columns.
- Do not make the stripe widths too large.
Since the size of stripe groups and stripe volumes are rounded to the common multiple of stripe width times the number of stripe columns and cylinder size, setting the stripe width too large may result in using the disk area inefficiently, or not being able to create a volume with the intended size.

- Where possible, connect disks and lower level groups with the same available sizes to the same stripe group.
The available size of the stripe group (available capacity as volumes) equals the available size of the smallest disk or the lower level group connected to the stripe group multiplied by the number of stripe columns, and rounded down to the common multiple of stripe width times the number of stripe columns and cylinder size.
When connecting disks or lower level groups with different available sizes to a stripe group, the larger disk or lower level group will only be able to use the capacity of the smaller disk or lower level group. For example, if a 4 GB disk and a 9 GB disk are connected to one stripe group, the 9 GB disk will only be able to use approximately 4 GB. Which means, the available size of stripe group will be approximately 8 GB (4 GB x 2).
- Where possible, striping across disks with similar performance specifications is recommended.
When striping disks with different performance specifications such as revolution speed, the performance becomes unbalanced and will depend on the slower disk performance.
- Using striping in combination with mirroring is recommended.
In a striping configuration, the risk of losing data from a disk failure increases as more disks are involved compared to a usual disk configuration.
By mirroring stripe groups, both the I/O load balancing and data redundancy can be achieved at the same time.

A.3.4 Guidelines for Concatenation

Important points to keep in mind when concatenating disks are listed below.

- The available size of the concatenation group (available capacity as volumes) equals the total of the available size of disks connected to the concatenation group.
- Where possible, concatenating disks with similar performance specifications is recommended.
When concatenating disks with different performance specifications such as the revolution speed, the performance becomes unbalanced.
- Concatenation in combination with mirroring is recommended.
When concatenating disks, the risk of losing data from a disk failure increases as more disks are involved compared to a usual disk configuration.
By mirroring concatenation groups, large capacity and data redundancy can be achieved at the same time.

A.3.5 Guidelines for Combining Striping with Mirroring

Important points to keep in mind when combining striping with mirroring are listed below.

- Where possible, mirroring stripe groups with similar configuration is recommended.
See "A.3.2 Guidelines for Mirroring" and "A.3.3 Guidelines for Striping" also.

A.3.6 Object Status

In order to maintain a correct object state, GDS manages the status of slice, volume, and disk as described below.

Slice Status

ACTIVE

Data is valid, and slice is accessible.

Volume is in ACTIVE status. The total number of slices that are in 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 valid, and slice is inaccessible.

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

INVALID

Data is invalid, and slice is inaccessible.

When an I/O error is detected within a slice during the mirroring process, it becomes INVALID. Promptly identify the problem, and restore the slice by checking the disk driver message or the like.

After restoring the slice, execute the synchronization copying. Successful completion will make the data valid (slice will become ACTIVE or STOP status). If unsuccessful, it will become INVALID again.

COPY

Synchronization copying is being performed to maintain data uniformity.

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

Successful completion of synchronization copying will make the slice valid (ACTIVE or STOP), and if unsuccessful, invalid (INVALID).

TEMP

Slice is detached from the volume, and can temporarily be accessed in isolation. If the `sdxslice -M` command is successfully completed, the slice becomes TEMP status. To revert to the original status (ACTIVE or STOP), execute the `sdxslice -R` command.

TEMP-STOP

Temporarily detached from volume and cannot be accessed in isolation.

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

NOUSE

No normal operation possible.

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

Volume Status

ACTIVE

Valid data is accessible.

There are one or more slices in ACTIVE or COPY (in process of background copying) status within the volume.

STOP

Volume is inaccessible, but can be activated to change to ACTIVE status. However, in case of proxy volume it must first be separated from the master volume.

There are one or more slices in STOP or COPY (in process of background copying) status within the volume.

INVALID

There is no slice with valid data, and volume is inaccessible.

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

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

Disk Status

ENABLE

Disk is accessible.

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

DISABLE

Abnormality is detected, and disk is inaccessible.

The configuration information of disks is checked when booting the system. If some disk is determined as not managed with GDS, it becomes DISABLE. In such case, all slices related to the disk become NOUSE status.

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

SWAP

Disk is inaccessible, and disk can be swapped.

Successful completion of the `sdxswap -O` command will change the disk to SWAP status. In such case, all slices related to the disk become NOUSE status.

Appendix B Logging View

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

Appendix C Environment Setting

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

Appendix D Command Reference

This appendix explains the commands provided by GDS.

This appendix explains the format and facility of commands, specifiable options, and return values.

Commands available are as follows.

Command	Meaning
sdxcass	Class operations
sdxdisk	Disk operations
sdxgroup	Group operations
sdxvolume	Volume operations
sdxslice	Slice operations
sdxinfo	Display object configuration and status information
sdxattr	Change attribute values of an object
sdxswap	Swap disk
sdxfix	Restore a crashed object
sdxcopy	Synchronization copy operations
sdxroot	Define or cancel mirroring of root file system
sdxparam	Configuration parameter operations
sdxproxy	Proxy object operations

D.1 sdxclass - Class operations

SYNOPSIS

sdxclass -R -c *class*

DESCRIPTION

Use sdxclass to perform operations on the class object specified by *class*. You must be super-user to use this command.

PRIMARY OPTIONS

You can use the following option.

-R

Remove

Removes the class definition indicated by *class*. If *class* is a shared class, the definition is removed from all nodes.

A disk registered with *class* will be removed with the *class*. However, if there is a group or a volume, the class will not be removed.

To place a removed disk under GDS management again, you need to re-register the physical disk in the class.

For further details, refer to "D.2 sdxdisk – Disk operations."

SUB OPTIONS

SUB OPTIONS are as follows:

-c *class*

The *class* indicates the class name that is the target of the operation.

RETURNED VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a non-zero value is returned.

D.2 sdxdisk - Disk operations

SYNOPSIS

```
sdxdisk -C -c class -g group -d disk,...  
          [-v num=volume:jrm [,num=volume:jrm,...]]  
          [-a attribute=value [,attribute=value]] [-e delay=msec ]
```

```
sdxdisk -D -c class -g group -d disk
```

```
sdxdisk -M -c class [-a attribute=value [,attribute=value,...]]  
          -d device=disk [:type ] [,device=disk [:type ],...]
```

```
sdxdisk -R -c class -d disk
```

DESCRIPTION

Use sdxdisk to perform operations on the disk object specified by *disk*. You must be super-user to use this command.

PRIMARY OPTIONS

You can use either of the following options.

-C

Connect

Connects one or more disks (keep disk, single disk or undefined disk) indicated by *disk*,... to *group*, which can be a mirror group, stripe group or concatenation group. The *class* indicates the class name with which *disk* is registered. The *group* indicates the group name to which the disk is connected.

If there is no group with the name given by *group*, one is automatically created.

You cannot connect a spare disk to a group. You cannot connect a keep disk or a single disk to an existing group. Also, a keep disk and a single disk, or multiple keep disks, or multiple single disks cannot be connected to the same group simultaneously.

The disk attribute will be changed to match the attribute of the group (mirror, stripe or concatenation) it is connected to. Disks and lower level groups that are connected to the same group will be mirrored, striped or concatenated, depending on its type attribute.

Details on connecting a disk to a mirror group, a stripe group and a concatenating group are explained below.

- When connecting to a mirror group

Disks and lower level groups connected to the same mirror group will be mirrored each other. When only one disk or one lower level group is connected, the volume created within that mirror group will not be mirrored. When configuring a mirroring environment with "n"-way multiplexing, "n" numbers of disks or lower level groups must be connected. A maximum of eight-way multiplex mirroring is supported.

When one or more volumes already exist within the mirror group indicated by *group*, the slice configuration of disk or lower level group that is already connected to *group* will be automatically copied to the newly connected *disk*. Also, when there is an activated volume within *group*, volume data in addition to the slice configuration will be automatically copied upon returning from the sdxdisk command, therefore increasing the mirroring multiplicity.

By connecting a single disk with single volumes to a mirror group, single volumes will be changed to mirror volumes.

The available size of the mirror group (available capacity as volumes) will be the same as the available size of the smallest disk or the lower level group connected to the mirror

group. When a keep disk is connected, the available size of the group will be the same as the available size of the keep disk. If connecting *disk* results in decreasing the available size of *group*, a warning message will be sent to standard error output.

When a keep disk is connected, volume will be created for each physical slice within the keep disk. If the physical slice is not open, the created volume will be activated and synchronization copying will automatically begin after returning from the `sdxdisk` command.

- When connecting to a stripe group

Disks indicated by *disk*,... will be connected to *group* in the order they are listed. Disks and lower level groups connected to the same stripe group will each configure a stripe column, and will be striped in the order they were connected. When only one disk or one lower level group is connected, volume cannot be created within that stripe group. When striping "n" number of columns, "n" number of disks or lower level groups must be connected. A minimum of two columns and a maximum of 64 columns are supported.

When a stripe group indicated by *group* already exists, stripe columns will be added after the existing stripe columns in *group*, in the order they are indicated by *disk*,.... However, you cannot increase stripe columns by connecting disks to stripe groups with volumes, or to stripe groups connected to a higher level group.

The available size of the stripe group (available capacity as volumes) equals the available size of the smallest disk or the lower level group connected to the stripe group multiplied by the number of stripe columns, and rounded down to the common multiple of stripe width times stripe columns and cylinder size. If connecting *disk* decreases the available size of *group*, a warning message will be sent to standard error output.

You cannot connect a single disk to a stripe group.

- When connecting to a concatenation group

Disks connected to the same concatenation group indicated by *disk*,... will be concatenated in the order they are listed. A maximum of 64 disks can be concatenated.

The available size of the concatenation group (available capacity as volumes) equals the total of the available size of disks connected to the concatenation group.

The available size of the group can be increased by connecting disks to the existing concatenation group. When a concatenation group indicated by *group* already exists, disks will be added in the order they are indicated by *disk*,... after the disk that was last concatenated in *group*. However, you cannot add a disk to a concatenation group that is connected to the highest level stripe group which already has a volume. Also if the concatenation group indicated by *group* is connected to a stripe group connected to a mirror group, a disk cannot be added.

You cannot connect a single disk to a concatenation group.

-D

Disconnect

Disconnect the disk (including spare disk) specified by *disk* from *group* which is a group. The *class* indicates the class name with which the *disk* is registered, and *group* indicates the group name to which *disk* is connected.

The disconnected disk will return to its original type attributes (keep disk, single disk or undefined disk).

If only *disk* is connected to *group*, *group* will automatically be removed upon disconnecting *disk*. However, when *disk* is the only object connected to *group* and *group* is connected to a higher level group, disconnection will result in an error. In such case, disconnect group from the higher level group using the `sdxgroup -D` command, and then disconnect disk.

You cannot disconnect *disk* if the disconnection will result in a change in the status of any of the existing volume within *group*.

Conditions on when you cannot disconnect a disk from a mirror group, stripe group or a concatenation group are explained below.

- When disconnecting from a mirror group

For example, you cannot disconnect *disk* from a mirror group if one or more volumes exist within the mirror group indicated by *group*, and *disk* is the only object connected to *group*.

- When disconnecting from a stripe group

You cannot disconnect a disk from a stripe group with one or more existing volumes, or from a stripe group connected to a higher level group.

- When disconnecting from a concatenation group

The only disk you can disconnect from a concatenation group is the disk that was concatenated last.

Disk containing volume data cannot be disconnected from a concatenation group.

You cannot disconnect a disk from a concatenation group that is connected to the highest level stripe group with one or more existing volumes. Also, if the concatenation group indicated by *group* is connected to a stripe group connected to a mirror group, the disk cannot be disconnected.

-M

Make

Registers one or more physical disks, indicated by *device*, with *class*. The *class* gives the name of the destination class. Once physical disks have been registered, they can then be managed using GDS. A disk managed by GDS is called SDX disk. User will use the disk name specified by *disk* to perform operations on the disk.

If no class with the name specified by *class* already exists, then one is automatically created.

The "keep" type *device* may be registered with "root" type *class*. However, when registering a keep type *device*, you must register the same number or more of "undef" type *device* at the same time.



Note

Since the `sdxdisk` command initializes the registered physical disks (except for "keep" type device), when registering a physical disk containing data, you must first create data backup.

-R

Remove

Removes the disk indicated by *disk* from class indicated by *class*. The *class* indicates the class name with which the *disk* is registered.

Once the disk is removed, it can no longer be managed using GDS.

When the last disk is removed from a *class*, that class definition is automatically removed.

Disk cannot be removed when a volume exists within *disk*, or when *disk* is connected to a group.

SUB OPTIONS

SUB OPTIONS are as follows:

-a *attribute=value*[,*attribute=value*] (when using -C)

When using the -C option and defining a new group name with the -g option, a new group is automatically created. Using this option sets the *attribute* of the created *group* to *value*. The *attribute* indicates the attribute name, and *value* indicates the attribute value. Always separate *attribute* and *value* with an equal (=) sign. When indicating multiple attributes, specifiers should be combined using comma (,) as the delimiter.

If no group is created, indicating a different attribute value from the existing *group* will result in an error. You cannot change the attribute value of an existing *group*.

You can indicate the following combination to *attribute* and *value*.

If multiple attributes are indicated and an error results because of any part of them, the entire process is canceled.

type=mirror, type=stripe or type=concat (default is mirror)

Sets type attribute of *group*. When *class* is a root class, indicating "stripe" or "concat" will result in an error.

mirror

Sets type attribute to "mirror".

stripe

Sets type attribute to "stripe".

concat

Sets type attribute to "concatenation".

width=*blks* (default is 32)

Sets the stripe width of *group*. The *blks* indicates the stripe width in blocks (base 10). One block is 512 bytes. For *blks*, you can indicate an integer (one and larger) that is two raised to the power, which is equal to or smaller than the available size of the smallest disk indicated by *disk*,.... If *group* is not a stripe group, this option will result in an error.

-a *attribute=value*[,*attribute=value*] (when using -M)

When using the -M option and defining a new class name with the -c option, a class is automatically created. Using this option sets the created *class attribute* to *value*.

The *attribute* indicates the attribute name, and the *value* indicates the attribute value.

Always separate *attribute* and *value* with an equal (=) sign. Specifiers should be combined using comma (,) as the delimiter. If no class is created, this option is ignored.

To change the attributes of an existing class, use the `sdxattr -C` command.

You can use the following *attribute* and *value* pairs in any combination.

If multiple attributes are specified and an error results because of any part of them, the entire process is canceled.

type=root, type=local, or type=shared (default is local)

Sets the attributes of *class* type.

root

Sets the type attribute to "root".

Objects within "root" type *class* can only be used by current node.

Unlike local type class, system disk including root file system may be registered with "root" type *class*.

Only one root type class can be created within a system.

When a root type class already exists, indicating a new "root" type *class* will result in an error.

For the scope attribute, the current node's node identifier will be automatically set.

local

Sets the type attribute to "local".

Objects within "local" type *class* can only be used by current node.

For the scope attribute, the current node's node identifier will be automatically set.

shared

Sets the type attribute to "shared."

By combining this with the scope attribute, the objects in the *class* can be shared among multiple nodes including the current node.

scope=node [:node:...] (default is the current node)

Sets the node set which share the *class* whose type attribute is specified to be "shared".

In *node*, indicates a node identifier that is defined by PRIMECLUSTER.

hs=on or **hs=off** (default is on)

It sets the operation mode of the hot spare.

on

Hot spare feature is enabled.

off

Hot spare feature is disabled.

When hot spare mode is set to "off," be aware that the spare disk will be prevented from automatically connecting.

-c class

The *class* indicates the class name to which the disk is registered or is to be registered, where the disk is the target of the operation.

-d device=disk [:type] [,device=disk [:type] ,...] (when using **-M**)

The *device* indicates the name of the physical disk, the *disk*, the name of the disk, and *type*, the type attribute of the disk.

The *device* must always be followed by an equal sign (=), and if a type is given, it is delimited from *disk* by a colon (:).

To register multiple *devices*, combine multiple definitions with a comma (,) as the delimiter.

The physical disk name can be specified with one of the following formats:

cCtTdD (for normal hard disks)

mphdI (for MPHD disks)

mplbI (for MPLB disks)

emcpowerN (for emcpower disks)

Where *C* is the controller number, *T* is the target ID, *D* is the disk number, *I* is the MPHD instance number or the MPLB instance number, and *N* is the emcpower device number.

emcpower device number is either a 1 or 2 digit number using numbers from 0 to 9, and does not specify small letters corresponding to slice number.

The *type* can be indicated as one of the following. If no type is indicated, the default value of undef (undefined disk) is used. When *class* is not a root class, indicating "keep" will result in an error.

keep

Keep disk. Its format and data will remain even when registered to class, and connected to group.

single

Single disk. Single volumes can be created on it.

spare

Spare disk.

undef

Undefined disk.

When "spare" is indicated for *type*, and the available size of *device* is smaller than the available size of the smallest mirror group in *class*, a warning message notifying that the hot spare function may not be available will be sent to standard error output.

When *keep* is indicated by *type*, the lack of sufficient space area or swap area on *device*, or lack of unused physical slice will result in a *sdxdisk* command error.

The *sdxdisk* command on system disks indicated as *keep* by *type* will terminate normally, even if all or part of the physical slices are open.

However, if anything other than system disk is indicated as *keep*, an open physical slice will result in an error.

For example, if a physical slice is used as a file system, you must first unmount the file system to regain the unused status of physical slice before executing the *sdxdisk* command.

When "single" is specified for *type*, *device* will be registered as a single disk.

For a "single" disk, you can create single volumes on it using the *sdxvolume* command without connecting the disk to any group.

-d *disk* (when using **-D, -R**)

The *disk* indicates the disk name that is the target of the operation.

-d *disk*,... (when using **-C**)

The *disk* indicates the disk name that is the target of the operation. To indicate multiple disks, separate each disk name with a comma (,) as the delimiter.

-e *delay=msec* (when using **-C**)

When a disk is connected to a mirror group, data contained in the volume will be copied as needed.

This option delays the issuing of the input/output request to the disk at the time of copying by milliseconds indicated by *msec*, allowing to adjust the influence on the application accessing the volume. The value is set to 0 by default. Values from 0 to 1000 may be indicated for *msec*.

If *group* is not a mirror group, this option is ignored.

-g *group* (when using **-C,-D**)

The *group* indicates the group name to which the disk is connected, or is to be connected, where *disk* is the target of the operation.

-v *num=volume:jrm* [,*num=volume:jrm*,...] (when using **-C**)

When connecting a "keep" type disk, indicates the attribute value of the volume that is to be created.

If no "keep" type *disk* is indicated, this option is ignored.

The *num* must always be followed by an equal sign (=), and *volume* and *jrm* must be delimited by a colon (:). To indicate multiple volume attributes, combine multiple definitions with a comma (,) as the delimiter.

The *num* indicates the physical slice number (integer between 0 and 7 inclusive, excluding 2) of the "keep" type disk where the volume data is stored, *volume* gives the name of the volume to be created, and *jrm* the just resynchronization mode setting of the created volume, which can be either on or off.

When there is more than one physical slice larger than 0 in the "keep" type *disk*, you must specify volume attributes corresponding to all physical slices.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.

Otherwise, a non-zero value is returned.

D.3 sdxgroup - Group operations

SYNOPSIS

```
sdxgroup -C -c class -h hgroup -l lgroup,...  
[ -a attribute=value [, attribute=value] ] [-e delay=msec ]
```

```
sdxgroup -D -c class -h hgroup -l lgroup
```

```
sdxgroup -R -c class -g group
```

DESCRIPTION

Use sdxgroup to perform operations on the group object.
You must be super-user to use this command.

PRIMARY OPTIONS

You can use one of the following primary options.

-C

Connect

Connects one or more groups (stripe group or concatenation group) indicated by *lgroup*,... to group (mirror group or stripe group) indicated by *hgroup*. The *class* indicates the class name to which *lgroup* belongs, and *hgroup* indicates the higher level group to which *lgroup* is to be connected. When *class* is a root class, this option will result in an error.

When no group with the same name as *hgroup* exists, a group will be created automatically.

Groups indicated by *hgroup* are referred to as higher level group, and groups indicated by *lgroup* are referred to as lower level group.

Lower level groups and disks connected to the same higher level group are mirrored or striped according to the type attribute of the higher level group. Connecting a group to a higher level group does not change the type attribute of the lower level group.

You cannot connect groups when:

- *lgroup* is a mirror group
- *hgroup* is a concatenation group
- type attributes of *lgroup* and *hgroup* are the same

A group that already contains volumes cannot be connected to another group.
Connecting a group to a mirror group and a stripe group is explained below.

- When connecting to a mirror group

You can connect one or more groups (stripe group or concatenation group) indicated by *lgroup*,... to *hgroup* which is a mirror group.

Disks and lower level groups connected to the same mirror group will be mirrored each other. When only one disk or one lower level group is connected, volumes created within that mirror group will not be mirrored. When configuring a mirroring environment with "n"-way multiplexing, "n" numbers of disks or lower level groups must be connected. A maximum of eight-way multiplex mirroring is supported.

When one or more volumes already exist within the mirror group indicated by *hgroup*, the slice configuration of disk or lower level group that is already connected to *hgroup* will be automatically copied to the newly connected *lgroup*. Also, when there is an activated volume within *hgroup*, volume data in addition to the slice configuration will be automatically copied upon returning from the sdxgroup command, therefore increasing the mirroring multiplexity.

The available size of the mirror group (available capacity as volumes) will be the same as the available size of the smallest disk or the lower level group connected to the mirror group. If connecting *lgroup* decreases the available size of *hgroup*, a warning message

will be sent to standard error output.

- When connecting to a stripe group

You can connect one or more groups (concatenation group) indicated by *lgroup*,... to *hgroup* which is a stripe group. Groups indicated by *lgroup*,..., will be connected to *hgroup* in the order they are listed.

Disks and lower level groups connected to the same stripe group will each configure a stripe column, and will be striped in the order they are connected. When only one disk or one lower level group is connected, volume cannot be created within that stripe group. When striping "n" number of columns, "n" number of disks or lower level groups must be connected. A minimum of two columns and a maximum of 64 columns are supported.

When a stripe group indicated by *hgroup* already exists, stripe columns will be added after the stripe columns that already exist in *hgroup*, in the order they are indicated by *lgroup*,.... However, you cannot increase the stripe columns by connecting groups to stripe groups with volumes, or to stripe groups connected to a higher level group.

The available size of a stripe group (available capacity as volumes) equals the available size of the smallest disk or the lower level group connected to the stripe group multiplied by the number of stripe columns, and rounded down to the common multiple of stripe width times stripe columns and cylinder size. If connecting *lgroup* decreases the available size of *hgroup*, a warning message will be sent to standard error output.

-D

Disconnect

Disconnects group indicated by *lgroup* from the higher level group *hgroup*. The *class* indicates class name to which *lgroup* belongs, and *hgroup* indicates the higher level group name to which *lgroup* is connected.

When *lgroup* is the only object connected to *hgroup*, *hgroup* will automatically be removed upon disconnecting *lgroup*. However, when *lgroup* is the only object connected to *hgroup*, and *hgroup* is connected to a higher level group, disconnection will result in an error. In such case, disconnect *hgroup* from its higher level group, and then disconnect *lgroup*.

You cannot disconnect *lgroup* if the disconnection may result in a change in the status of any existing volume within *hgroup*.

Conditions on when you cannot disconnect a group from a mirror group or a stripe group are explained below.

- When disconnecting from a mirror group

For example, you cannot disconnect *lgroup* from a mirror group if one or more volumes exist within the mirror group indicated by *hgroup*, and *lgroup* is the only object connected to *hgroup*.

- When disconnecting from a stripe group

You cannot disconnect a lower level group from a stripe group with one or more existing volumes, or from a stripe group connected to a higher level group.

-R

Remove

Remove the group definition specified by *group*. The *class* indicates the class name to which *group* belongs.

Disks and lower level groups connected to *group* will be disconnected. The disconnected disk's attribute will return to its original setting (keep disk, single disk or undefined disk).

group cannot be removed when one or more volumes exist within *group*, or when *group* is connected to a higher level group.

SUB OPTIONS

SUB OPTIONS are as follows:

-a *attribute=value*[,*attribute=value*] (when using **-C)**

When using the **-C** option and defining a new group name with the **-h** option, a new group *hgroup* is automatically created. Using this option sets the *attribute* of the created *hgroup* to *value*.

The *attribute* indicates the attribute name, and *value* indicates the attribute value. Always separate *attribute* and *value* with an equal (=) sign. When indicating multiple attributes, specifies should be combined using comma (,) as the delimiter.

If no group is created, indicating a different attribute value from the existing *hgroup* will result in an error. You cannot change the attribute value of an existing *hgroup*.

You can indicate the following combination to *attribute* and *value*.

If multiple attributes are indicated and an error results because of any part of them, the entire process is canceled.

type=mirror or type=stripe (default is mirror)

Sets the type attribute of *hgroup*.

mirror

Sets type attribute to "mirror."

stripe

Sets type attribute to "stripe."

width=blks (default is 32)

Sets the stripe width of *hgroup*. The *blks* indicates the stripe width in block number (base 10). One block is 512 bytes. For *blks*, you can indicate an integer that is two raised to the power (one and larger), which is equal to or smaller than the available size of the smallest group indicated by *lgroup*,.... If *hgroup* is not a stripe group, this option will result in an error.

-c *class*

The *class* indicates the class name to which the group belongs, where group is the target of the operation.

-e delay=msec (when using **-C)**

When a group is connected to a mirror group, data contained in the volume will be copied as needed. This option delays the issuing of the input/output request to the disk by milliseconds indicated by *msec*, allowing to adjust the effect on the application accessing the volume. Default is 0. Values from 0 to 1000 may be indicated for *msec*. If *hgroup* is not a mirror group, this option is ignored.

-g *group* (when using **-R)**

The *group* indicates the group name that is the target of the operation.

-h *hgroup* (when using **-C,-D)**

The *hgroup* indicates the higher level group name to which the lower level group is connected or is to be connected, where the lower level group is the target of the operation.

-l *lgroup* (when using **-D)**

The *lgroup* indicates the lower level group name that is the target of the operation.

-l *lgroup*,... (when using **-C)**

The *lgroup* indicates the lower level group name that is the target of the operation. To connect multiple groups, separate each group name with a comma (,) as the delimiter.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.

Otherwise, a non-zero value is returned.

D.4 sdxvolume - Volume operations

SYNOPSIS

```
sdxvolume -F -c class [ -v volume,... ]
           [-e { allnodes|node=node [:node:... ] } ]

sdxvolume -M -c class { -g group | -d disk } -v volume -s size
           [ -a attribute=value [,attribute=value]] [-e delay=msec ]

sdxvolume -N -c class [ -v volume,... ]
           [-e [allnodes|node=node [:node:...]],delay=msec, mode=val, nosync, unlock]

sdxvolume -R -c class -v volume
```

DESCRIPTION

Use sdxvolume to perform operations on the volume object specified by *volume*. You must be super-user to use this command.

PRIMARY OPTIONS

You can use either of the following options.

-F

offline

Stops one or more volumes specified by *volume*. If the -v option is omitted, all volumes within the *class* are taken offline. Offline volumes cannot be accessed.

If synchronization copying of *volume* is in process, it cannot be stopped. You can stop the synchronization copying using the sdxcopy -C command.

The stopped volume will be activated when the node is rebooted. (Except for when the volume is locked.) Attempting this operation while a *volume* is in use results in an error.

-M

Make

Creates a volume indicated by *volume*, within the highest level group indicated by *group*, or within a single disk indicated by *disk*. The *size* indicates the number of blocks on *volume*, *class* indicates the class name associated with the *group* or *disk*.

Within the same *group* or *disk*, you can create up to six volumes with their physical slice attribute set to "on." A maximum of 254 volumes can be created in total, including the volumes with physical slice attribute set to "off." When -a pslice=*value* option is omitted, volumes with physical slice attribute set to "on," will be created. However, note that you cannot create a volume with physical slice attribute set to "on," if *group* is either stripe group, concatenation group, or a mirror group where its only directly-connected object is a lower level group. In such case, you must indicate the -a pslice=off option, and set the physical slice attribute to "off."

Once created, the volume is activated, and can be accessed using the following special files:

```
/dev/sfdsk/class/rdsk/volume
/dev/sfdsk/class/dsk/volume
```

If *group* is a mirror group, the system will automatically execute a synchronization copying upon returning from the sdxvolume command.

The features of volumes created when *group* is mirror group and stripe group is explained below.

- **When group is a mirror group**

To ensure data availability, GDS restricts the mirroring on a single piece of disk unit. In case of mirror groups, a mirror volume that consists of mirror-multiplexing equal to the number of connected disks or lower level groups is created (maximum of eight). When only one disk or one lower level group is connected, the volume created within that mirror group will not be mirrored.

If the last block number of the volume that is created within a mirror group is larger than the available size of any of the spare disks registered with *class*, a warning message is sent to standard error output informing you that the hot spare feature is disabled.

- **When group is a stripe group**

In a stripe group, stripe volumes with columns equal to the number of connected disk or lower level groups are created. When only one disk or lower level group is connected, volume cannot be created.

-N

oNline

Activates one or more volumes specified by *volume*, If the *-v* option is omitted, all volumes within *class* are activated.

Activated volumes can be accessed.

If there is a slice in the TEMP status on the *volume*, a warning message is sent to standard error output.

When *volume* is a mirror volume, the system will determine whether synchronization has been lost upon returning from the *sdxvolume* command, and automatically execute a synchronization copying as needed (except for when *-e nosync* is indicated).

-R

Remove

Removes the volume specified by *volume* and releases the disk area used in the group or the single disk.

If the specified volume is activated, this command results in an error.



Note

Be aware that any data stored on *volume* will be lost.

SUB OPTIONS

SUB OPTIONS are as follows:

-a *attribute=value*[,*attribute=value*] (when using **-M**)

Use this to set an *attribute* for the *volume*.

The *attribute* indicates the attribute name, and the *value* indicates the attribute value. Always separate *attribute* and *value* with an equal (=) sign.

When indicating multiple attributes, specifiers should be combined using comma (,) as the delimiter.

You can indicate the following combination to *attribute* and *value*.

If multiple attributes are specified and an error results because of any part of them, the entire process is canceled.

jrm=on or **jrm=off** (default is on)

Sets the volume's JRM mode.

When stripe group or concatenation group is indicated with *-g* option, this option will result in an error.

on

JRM is on.

off

JRM is off.

pslice=on or **pslice=off** (default is on)

Sets the physical slice attribute value of volume.

When using *-g* option to indicate a group in which a physical slice cannot be created (i.e. stripe group, concatenation group, or mirror group where the only object directly connected to mirror group is lower level group), you cannot set this option to "on."

When the type of *class* is root, you cannot set this value to be "off".

on

Physical slice attribute value is set to be "on".

Of the slices consisting the volume, the slice within the single disk, or slices within disks that are directly connected to the mirror group is registered to the VTOC, and physical slices are created.

off

Physical slice attribute value is set to be "off".

None of the slices consisting the volume is registered to the VTOC, and physical slice will not be created. When physical slice attribute is set to "off", the slice cannot be detached.

-c class

The *class* indicates the class name to which the volume that is the target of the operation belongs, or the class name in which the volume is to be created.

-d disk (when using -M)

The *disk* indicates the single disk name in which the single volume will be created.

-e allnodes (when using -F,-N)

Stops or activates the volume on all nodes included in the scope of *class*. Stopped nodes are ignored. *class* must be a shared class.

When neither this option nor `-e node=node[:node:...]` option is specified, volume is stopped or started only on the self-node.

-e delay=msec (when using -M,-N)

If synchronization is not maintained when creating or activating a mirror volume, synchronization copying will take place automatically (except for when `-e nosync` is indicated). This option delays the issuing of the input/output request to the disk at the time of copying by milliseconds indicated by *msec*, allowing to adjust the effect on the application accessing the volume. The value is set to 0 by default. Values from 0 to 1000 may be indicated for *msec*.

If mirror group is not indicated with *group*, this option is ignored.

-e mode=val (when using -N)

Specifies the access mode for one or more volumes that will be activated.

val indicates either of the following options.

rw

Sets access mode for read and write.

ro

Sets access mode for read only. Opening a read-only volume in write mode will result in an error.

Although volume will be activated in the access mode specified by *val*, the access mode attribute for volume will remain unchanged. Access mode indicated by *val* ("Current Access Mode") is valid only while the volume is activated, and will become invalid once the volume is stopped. When the volume is restarted, it will start in the mode set by access mode attribute ("Default Access Mode"), except for when the access mode is specified at the time of restart. In order to start a volume that is already activated on the current node in a different access mode, you must first stop the volume.

-e node=node[:node:...](when using -F,-N)

Stops or activates the volume on one or more specified nodes.

Stopped nodes are ignored. You must specify the node identifier of the node to stop or activate the volume to *node*. If a node not included in the scope of *class* is specified, the volume is not stopped or activated on any node. *class* must be a shared class.

If this option and the `-e allnodes` are both omitted, the volume is stopped or activated only on the current node.

-e nosync (when using **-N**)

Disables automatic synchronization copying after activating a mirror volume.
If mirror group is not indicated with *group*, this option is ignored.



Note

Volumes that are activated using this option will not be mirrored.
In order to configure a mirroring environment, you must perform synchronization copying with the `sdxcopy -B` command.

-e unlock (when using **-N**)

The volume will be activated regardless to whether it is locked. Lock mode will not be changed unless you change it with the `stxattr -V` command.

-g group (when using **-M**)

The *group* indicates the group name in which the volume will be created.

-s size (when using **-M**)

Specifies the size of the *volume* being created, in blocks (base 10).

One block is 512 bytes.

When *group* indicates stripe group, the size of *volume* created will be *size* rounded up to a common multiple of stripe width multiplied by stripe columns and cylinder size. In other cases, the size of *volume* created will be *size* rounded up to the integer multiple of cylinder size.

-v volume (when using **-M,-R**)

The *volume* indicates the volume name that is the target of operation.

-v volume,... (when using **-F,-N**)

The *volume,...* indicates one or more volume names that is the target of the operation. To indicate multiple volumes, separate each volume name with a comma (,) as the delimiter.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.

Otherwise, a non-zero value is returned.

D.5 sdxslice - Slice operations

SYNOPSIS

```
sdxslice -F -c class -s slice,...  
sdxslice -M -c class -d disk -v volume [-a attribute=value[,attribute=value]]  
sdxslice -N -c class -s slice,...  
sdxslice -R -c class {-d disk -v volume|-s slice } [-e delay=msec,waitsync]  
sdxslice -T -c class -s slice,...
```

DESCRIPTION

Use sdxslice to perform operations on slice objects.
You must be super-user to use this command.

PRIMARY OPTIONS

You can use either of the following options.

-F

oFline

Stops the slice or slices indicated by *slice*,.... Offline slices cannot be accessed. The *slice* indicates the mirror slice name detached from the mirror volume using -M option.

The *class* indicates the class name to which *slice* belongs.
Offline slices will be activated upon reboot.
Attempting this command while *slice* is in use results in an error.

-M

Make

Temporarily detaches one of the mirror slices used in a copy of a mirror volume indicated by *volume* with mirroring multiplicity of two and higher, which is a part of the disk indicated by *disk*. The *class* indicates the class name to which *volume* belongs.

Only when the physical slice attribute value of *volume* is "on", you can detach the slice.
When the physical slice attribute is set to "off", you must turn it to "on" using the sdxattr -V command before executing this command.

Once detached, a special file is placed on the system. The path name is given below.

/dev/sfdsk/class/[r]dsk/disk.volume

User will be able to access the slice with this special file. You can use this slice to create a data backup of volume.

When the *class* is a shared class, only the node that detached the slice can access the slice. Other nodes sharing the *class* cannot access it. If you need to access from other nodes, you can take over the access right with the -T option.

A slice can be detached even though the *volume* is active.
You must ensure the integrity of backup data at the file-system layer or database layer.
If you are handling the volume as the file-system, for instance, there will be situations where you must regain integrity using the fsck command.



Note

Be aware that as long as a slice is not attached using the -R option, the degree of mirror multiplexing stays reduced.

-N

oNline

Activates the slice or slices indicated by *slice*,.... Activated slices can be accessed.
The *slice* indicates the mirror slice name detached from the mirror volume using -M option.
The *class* indicates the class name to which *slice* belongs.

-R

Remove

Reassembles the slice as part of the volume, where the slice is indicated by *slice* or combination of *disk* and *volume*.

The *disk* and *volume* combination or *slice* indicates the mirror slice name disconnected from the mirror volume using -M option.

The *class* indicates the class name to which the slice belongs.

After (when using -e waitsync option, before) returning from the sdxslice command, the slice is automatically reassembled with the volume. If the volume is active at this time, a synchronization copy is executed.

Attempting this command while the slice is in use results in an error.

-T

Takeover

Takes over the slice or slices indicated by *slice* from another node. When the takeover is complete, the slice will stop on the original node and be activated on the current node, allowing operation to the slice on the current node.

When a slice is attached to the volume after executing this command, the entire block will be copied regardless to the setting of JRM mode.

This option is valid only for shared class.

The *slice* indicates the mirror slice name disconnected from the mirror volume using -M option.

The *class* indicates the class name to which *slice* belongs.

SUB OPTIONS

SUB OPTIONS are as follows:

-a attribute=value[,attribute=value] (when using -M)

Sets the *attribute* attribute of the slice to be *value*.

Both attribute values become invalid at the point when the slice is assembled with the volume.

The *attribute* indicates the attribute name, and the *value* indicates the attribute value.

Always separate *attribute* and *value* with an equal (=) sign.

Specifiers should be combined using comma (,) as the delimiter.

You can use the following *attribute* and *value* pairs in any combination.

If multiple attributes are specified and an error results because of any part of them, the entire process is canceled.

jrm=on or jrm=off (default is on)

Sets the slice's JRM mode.

on

JRM is on.

off

JRM is off.

mode=rw or mode=ro (rw when omitted)

Sets the slice access mode.

rw

Sets access mode for read and write.

ro

Sets access mode for read only. Opening a read-only slice in write mode will result in an error.

-c class

The *class* indicates the class name to which the slice belongs.

-d disk (when using -M,-R)

The *disk* indicates the disk name to which the slice belongs, where slice is the target of the operation.

-e delay=msec (when using **-R**)

This option delays the issuing of the input/output request to the disk at the time of synchronization copying of the slice detached from volume, by milliseconds indicated by *msec*.

Always separate delay and *msec* with an equal (=) sign.

This option allows you to adjust the influence on the application accessing the volume.

The value is set to 0 by default. Values from 0 to 1000 may be indicated for *msec*.

-e waitsync (when using **-R**)

When synchronization copying is executed, returns the command after the copying process is complete.

-s slice (when using **-R**)

The *slice* indicates slice that is the target of operation. Slice name should be indicated in *disk.volume* format.

-s slice,... (when using **-F,-N,-T**)

The *slice* indicates one or more slice names that is the target of the operation. To indicate multiple slices, separate each slice name with a comma (,) as the delimiter.

Slice name should be indicated in *disk.volume* format.

-v volume (when using **-M,-R**)

Specifies the name of volume comprising the slice that is the target of the operation.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.

Otherwise, a non-zero value is returned.

D.6 sdxinfo - Display object configuration and status information

SYNOPSIS

```
sdxinfo [ -ACDGSV ] [ -c class ] [ -o object ] [ -e long ]
```

DESCRIPTION

Use sdxinfo to display configuration and status information of the objects which the current node shares.

PRIMARY OPTIONS

Primary options indicate the type of objects to display.

If nothing is specified, only information on the pertinent object will be displayed.

Object types can be any combination of the following.

-A

All

Displays all specified objects, and information on all related objects.

This is equivalent to -CDGSV.

If any other options are combined with this, they are ignored.

-C

Class

Displays all specified objects, and class information of the related objects.

-D

Disk

Displays all specified objects, and disk information of the related objects.

-G

Group

Displays all specified objects, and group information of the related objects.

-S

Slice

Displays all specified objects, and slice information of the related objects.

-V

Volume

Displays all specified objects, and volume information of the related objects.

SUB OPTIONS

SUB OPTIONS are used to specify the names of objects to display.

If nothing is specified, the command is interpreted as if all objects at the current node had been specified.

-c *class*

The *class* indicates the class name whose information will be displayed.

If this option is omitted, this command is interpreted as if all classes had been specified.

In case of combining this option with the -o option, the objects related to the specified *object* within the *class* is displayed.

-e long

Displays detailed object information.

-o *object*

The *object* indicates the object name (class name, disk name, group name or volume name) whose information will be displayed.

If this option is omitted, this command is interpreted as if all object names had been specified.

In case of combining this option with the -c option, the objects related to the specified *object* within the *class* is displayed.

DISPLAYED INFORMATION

Information displayed in response to the `sdxinfo` command can be interpreted as follows.

Values displayed for class information are as follows:

OBJ

Display class as an object classification.

NAME

Display class name.

TYPE

The type attribute value can be any of the following:

root

root class

local

local class

shared

shared class

SCOPE

Node identifier is displayed as the scope attribute value. If the class is shared by multiple nodes within a PRIMECLUSTER system, node identifiers are separated by colons (:).

HS

When -e long option is used, displays hot spare mode, which can be either of the following.

on

Hot spare feature is enabled.

off

Hot spare feature is disabled.

When hot spare mode is set to "off," be aware that the spare disk will be prevented from automatically connecting.

SPARE

Displays the number of spare disks that are not connected to the group.

Values displayed for disk information are as follows:

OBJ

Displays disk as an object classification.

NAME

Displays disk name.

TYPE

The type attribute value can be any of the following:

mirror

Mirror. It is connected to a mirror group.

stripe

Stripe. It is connected to a stripe group.

concat

Concatenation. It is connected to a concatenation group.

keep

Keep. Its format and data will remain even when registered with class, and connected to group.

single

Single. Single volumes can be created on it.

spare

Spare. Also the same when the spare disk is connected to the group.

undef

Undefined. Its usage is not determined yet.

CLASS

Displays the class name to which the disk belongs.

GROUP

Displays the group name to which the disk is connected. If the disk is not connected to any group, an asterisk (*) is displayed.

DEVNAM

Displays the physical disk name in one of the following formats:

cCtTdD (for normal hard disks)

mphd/ (for MPHD disks)

mplb/ (for MPLB disks)

emcpowerN (for emcpower disks)

where *C* is the controller number, *T* is the target ID, *D* is the disk number, *I* is the MPHD instance number or the MPLB instance number, and *N* is the emcpower device number.

DEVBLKS

Displays the size of the physical disk. The size is given in blocks (sectors.)

FREEBLKS

When -e long option is used, displays the number of free allocatable blocks (sectors) for a new volume. If the disk is not a single disk, an asterisk (*) is displayed.

DEVCONNECT

Displays a list of node identifiers of the nodes to which the physical disk is connected, separated using colons (:) as delimiters.

STATUS

Disk status can be any of the following:

ENABLE

Available for work.

DISABLE

Not available for work.

SWAP

Not available for work, but available for disk exchanging.

E

When -e long option is used, displays error status of disk, which can be any of the following.

0

I/O error has not occurred on both the current node and the shared node.

1

I/O error has occurred on either the current node or the shared node.

Values displayed for group information are as follows:

OBJ

Display group as an object classification.

NAME

Display group name.

CLASS

Displays the class name to which the group belongs.

DISKS

Displays the name of disks or lower level groups that are connected to the group, separated using colons (:) as delimiters. In case of a stripe group, names are listed in the order they are striped. Likewise, in case of a concatenation group, names are listed in the order they are concatenated.

BLKS

Displays size of the group, which is the total of available size of group (available capacity as volumes) plus one private slice size. Size is given in blocks (sectors).

FREEBLKS

The number of free allocatable blocks for a new volume.
If the group is a lower level group, an asterisk (*) is displayed.

SPARE

Displays the number of spare disks that can be connected to the group. If the group is a stripe group or a concatenation group, an asterisk (*) is displayed.

MASTER

When -e long option is used, it displays the group name of master group. When it is not a proxy group, an asterisk (*) is displayed.

TYPE

When -e long option is used, displays type attribute value, which can be any of the following.

mirror

Mirror group.

stripe

Stripe group.

concat

Concatenation group.

WIDTH

When -e long option is used, displays stripe width in blocks (sectors). If the group is not a stripe group, an asterisk (*) is displayed.

Values displayed for volume information are as follows:

OBJ

Displays volume as an object classification.

NAME

Displays the name of the volume.
If it is an area that cannot be allocated (area for private slice) or that can be allocated but have not (unallocated area), an asterisk (*) is displayed.

TYPE

When -e long option is used, displays type attribute, which can be any of the following.

mirror

Mirror. Belongs to a mirror group.

stripe

Stripe. Belongs to a stripe group.

concat

Concatenation. Belongs to a concatenation group.

single

Single. Belongs to a single disk.

CLASS

Displays the class name to which the volume belongs.

GROUP

Displays the highest level group name to which the volume belongs. When it belongs to a single disk, an asterisk (*) is displayed.

DISK

When -e long option is used, displays the name of the single disk to which the volume belongs. When it belongs to a group, an asterisk (*) is displayed.

MASTER

When -e long option is used, it displays the volume name of master volume. When it is not a proxy volume, an asterisk (*) is displayed.

PROXY

When -e long option is used, it displays the proxy volume status in one of the two ways as given below. When it is not a proxy volume, an asterisk (*) is displayed.

Join

The volume is being joined to a master volume.

Part

The volume is being parted from a master volume.

SKIP

Displays the skip-resynchronization mode setting, which can be either of the following. If the volume is neither a mirror volume nor a single volume, an asterisk (*) is displayed.

on

Skip resynchronization.

off

Execute resynchronization.



Note

Note that the interface for setting or changing this option is not available.

JRM

Displays the just resynchronization mode setting, which can be either on or off. If the volume is neither a mirror volume nor a single volume, an asterisk (*) is displayed.

on

JRM is on.

off

JRM is off.

MODE

When -e long option is used, it displays the access mode attribute value (default access mode) for the current node, which can be either of the following. If it is either an area for private slice or an unallocated area, an asterisk (*) is displayed.

rw

Read and write mode.

ro

Read only mode.

CMODE

When -e long option is used, it displays the present access mode of the activated volume from the current node. If the volume is not activated, an asterisk (*) is displayed.

rw

Read and write mode.

ro

Read only mode.

LOCK

When -e long option is used, displays the lock mode of current node, which can be either of the following. If it is either a private area or an unallocated area, an asterisk (*) is displayed.

on

The volume is locked from activating thereafter.

off

The volume is not locked from activating thereafter.

1STBLK

Displays the block (sector) number of the first block.

The block number is the logical block number, which is the offset in the group to which the volume belongs, and not the physical block number indicating the offset on the physical disk. However, when the volume belongs to a single disk, the block number will match the physical block number on the single disk. Also, when it belongs to a mirror group to which a disk is directly connected, the block number will match the physical block number on the disk.

LASTBLK

Displays the block (sector) number of the last block.

The block number is the logical block number, which is the offset in the group to which the volume belongs, and not the physical block number indicating the offset on the physical disk. However, when the volume belongs to a single disk, the block number will match the physical block number on the single disk. Also, when it belongs to a mirror group to which a disk is directly connected, the block number will match the physical block number on the disk.

BLOCKS

Displays the size in blocks (sectors).

STATUS

Displays the volume status of the current node, which can be any of the following.

ACTIVE

Ready for work.

STOP

Stopped.

INVALID

Stopped, and cannot be activated due to problem with data.

FREE

Not yet allocated as a volume.

PRIVATE

An area reserved for GDS control, so cannot be allocated as a volume.

PSLICE

When -e long option is used, displays the physical slice attribute value, which can be either of the following. If it is either a private area or an unallocated area, an asterisk (*) is displayed.

on

Physical slice attribute of the volume is set to "on." Of the slices consisting the volume, the slice within the single disk, or each slice within each disk where the disk is directly connected to the mirror group is registered to the VTOC, and a physical slice is created. If lower level group is the only object directly connected to mirror group, the volume will not have a physical slice, regardless to this attribute set to "on." Also, when the volume belongs to either a stripe group or a concatenation group, this attribute value will never be set to "on."

off

Physical slice attribute of the volume is set to "off." The volume has no physical slices, and none of the slices in the volume is registered to the VTOC.

SNUM

When -e long option is used, displays the slice number of slices in the volume. If no physical slice which constitute volume, an asterisk (*) is displayed.

Values displayed for slice information are as follows:**OBJ**

Display slice as an object classification.

NAME

When -e long option is used, display the name of slice.

When the slice is not a mirror slice that is temporarily detached from the mirror volume using the sdxslice -M command, an asterisk (*) is displayed.

CLASS

Displays the class name to which the slice belongs.

GROUP

Displays the highest level group name to which the slice belongs. If it is a single slice, an asterisk (*) is displayed.

DISK

Displays the name of the disk or the lower level group (i.e. the group to which this slice belongs, among the groups that are directly connected to the relevant highest level group) to which the slice belongs. If the highest level group is either a stripe group or a concatenation group, an asterisk (*) is displayed.

VOLUME

Displays the volume name to which the slice belongs.

JRM

When -e long option is used, displays the just resynchronization mode setting, which can be either on or off.

When the slice is not a mirror slice that is temporarily detached from the mirror volume using the sdxslice -M command, an asterisk (*) is displayed.

on

JRM is on.

off

JRM is off.

MODE

When -e long option is used, displays the access mode, which can be either of the following. When the slice is not a mirror slice that is temporarily detached from the mirror volume using the sdxslice -M command, an asterisk (*) is displayed.

rw

Read and write mode.

ro

Read only mode.

STATUS

Displays the slice status on the current node, which can be any of the following.

ACTIVE

Ready for work.

STOP

Stopped.

INVALID

Due to problem with data, has been temporarily detached from the volume.

COPY

Specifies a copy is underway, to maintain data uniformity.

TEMP

Temporarily detached from volume. Slice is operating in isolation.

TEMP-STOP

Temporarily detached from volume. Slice is stopping in isolation.

NOUSE

Stopped, with no operations possible.

COPY

When -e long option is used, displays either of the following copying process statuses. When not in COPY status, an asterisk (*) is displayed.

run

Copying is underway.

bg

Copying is in process in the background, but you can access valid data.

intr

Copying is interrupted. Executing the sdxcopy -I command interrupts the copy.

wait

Since many copying processes are in progress, it is waiting to be scheduled.

CURBLKS

When -e long option is used, displays the number of blocks (sectors) that has been copied.

When CURBLKS and the later described COPYBLKS match, all copying has completed.

When not in COPY status, and when copying is performed using TimeFinder or SRDF, an asterisk (*) is displayed.

COPYBLKS

When -e long option is used, displays the number of blocks (sectors) that needs to be copied. Usually this is the same size as the volume size it is registered with, but when just resynchronization copying is in process, the number of blocks that actually needs to be copied will be displayed.

When not in COPY status, and when copying is performed using the EMC TimeFinder or SRDF, an asterisk (*) is displayed.

DLY

When -e long option is used, displays the copy delay time in milliseconds.

When not in COPY status, an asterisk (*) is displayed.

CPTYPE

When -e long option is used, one of the following values is displayed as copy function which is used for copying process between the master and the proxy. When copying is not in process, or if the volume to which the slice belongs is not a target volume of copying process between a master volume and a proxy volume, an asterisk (*) is displayed.

soft

Copying, using the soft copy function provided by the sfdsk driver, is in process.

EC

Copying, using the EC function provided by GR700 series is in process.

OPC

Copying, using the OPC function provided by GR700 series or ETERNUS3000, is in process. When master and proxy are joined, and original volume is activated, part of the copying process may be performed using the soft copy function.

REC

Copying, using the REC function provided by GR700 series, is in process.

ROPC

Copying, using the ROPC function provided by GR700 series, is in process. When master and proxy are joined, and original volume is activated, part of the copying process may be performed using the soft copy function.

TF

Copying, using the EMC TimeFinder, is in process.

SRDF

Copying, using the EMC SRDF, is in process.

CPSOURCE

When -e long option is used, the volume name of source proxy volume which is used to restore data is displayed. When the volume to which the slice belongs is not a target master volume of restoring process, an asterisk (*) is displayed.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.
Otherwise, a non-zero value is returned.

USAGE EXAMPLES

Displays all the objects within the current node.

```
sdxinfo
```

Displays information on all the disks registered with the class called "Class1."

```
sdxinfo -A -c Class1
```

Use this to check whether an object called "foo" is currently being used.

```
sdxinfo -o foo
```

CAUTION

Additional information may be displayed in accordance with new functions provided.

D.7 sdxattr - Set objects attributes

SYNOPSIS

```
sdxattr -C -c class -a attribute=value[,attribute=value,...]
sdxattr -D -c class -d disk -a attribute=value[,attribute=value]
sdxattr -G -c class -g group -a attribute=value
sdxattr -S -c class -s slice -a attribute=value[,attribute=value]
sdxattr -V -c class -v volume -a attribute=value[,attribute=value,...]
```

DESCRIPTION

Sdxattr sets attribute values of objects on the current node.
You must be super-user to use this command.

PRIMARY OPTIONS

Primary options are used to specify the category of the object whose attributes are to be set.

-C

Class

Set the attributes of the class specified by *class*.

-D

Disk

Set the attributes of the disk specified by *disk*.

The *class* indicates the class name with which *disk* is registered.

If *disk* is connected to a group, or if there is a volume within *disk*, it will result in an error, and the change in attribute will not be executed.

-G

Group

Set the attributes of the group specified by *group*.

The *class* indicates the class name to which *group* belongs.

When *group* is connected to another group, this option will result in an error and you cannot change the attribute. Also, this option will result in an error if there are one or more activated volumes within the group. Stop all volumes before executing this command.

-S

Slice

Set the attributes of the slice specified by *slice*.

The *class* indicates the class name to which the *slice* belongs.

-V

Volume

Set the attributes of the volume specified by *volume*.

The *class* indicates the class name to which the *volume* belongs.

SUB OPTIONS

SUB OPTIONS are as follows:

-a *attribute=value[,attribute=value,...]* (when using **-C**)

Sets the *attribute* attribute of the *class* to be *value*.

The *attribute* indicates the attribute name, and the *value* indicates the attribute value.

Always separate *attribute* and *value* with an equal (=) sign.

Specifiers should be combined using comma (,) as the delimiter.

You can use the following *attribute* and *value* pairs in any combination.

When attempting to set multiple attributes, if any of them result in an error, no attributes are set.

type=local or type=shared

Sets the class type attribute.

When there is an activated volume within *class*, you can change the type attribute from "local" to "shared," but not vice versa. In order to change it from "shared" to "local," you must execute the command after first stopping all volumes.

You cannot change it from "root" to "local" or "shared", or vice versa.

local

Sets the type attribute to "local."

Objects within "local" type class can only be used by current node.

For the scope attribute, current node's node identifier will be automatically set.

shared

Sets the type attribute to "shared."

By combining this with the scope attribute, objects in the class can be shared among multiple nodes including the current node.

scope=node:node:...

For a "shared" type class, changes the nodes which share the class.

When there is an activated volume within *class*, you can add new nodes, but you cannot remove a node that has already been included in the scope. In order to remove a node, you must execute the command after first stopping all volumes.

If the *node* settings are not all complete, this will result in an error.

The *node* indicates a node identifier that is defined by PRIMECLUSTER.

For *scope*, you can indicate up to sixteen nodes.

hs=on or hs=off

It sets the operation mode of the hot spare.

You can make changes regardless to whether there is an activated volume within *class*.

on

Hot spare feature is enabled.

off

Hot spare feature is disabled.

The spare disk will be prevented from automatically connecting.

name=classname

Sets the class name to *classname*.

When there is an activated volume within *class*, it results in an error. First, you must stop all volumes.

When changing a class name through this operation, the paths of special files for volumes are also changed, so you must update the files in which the paths are described, such as /etc/vfstab.

/dev/sfdsk/classname/[r]dsk/volume_name

-a *attribute=value*[,*attribute=value*] (when using -D)

Sets the *attribute* attribute of the *disk* to be *value*.

The *attribute* indicates the attribute name, and the *value* indicates the attribute value.

Always separate *attribute* and *value* with an equal (=) sign.

Specifiers should be combined using comma (,) as the delimiter.

You can use the following *attribute* and *value* pairs in any combination.

When attempting to set multiple attributes, if any of them result in an error, no attributes are set.

type=keep*, *type=single*, *type=spare* or *type=undef

Set the disk type attribute.

keep

Sets the type attribute to "keep."

Hereinafter, this disk will be regarded as a keep disk, and its format and data will remain when connected to a group.

You can not change a single disk to keep disk.

single

Sets the type attribute to "single."

Single volume may be created within *disk* thereafter. You cannot change a keep disk to a single disk.

spare

Sets the type attribute to "spare."

The disk will be used as a spare disk thereafter.

When the available size of *disk* is smaller than the available size of the smallest mirror group within *class*, a warning message notifying that the hot spare function may not be available will be sent to standard error output.

undef

Sets the type attribute to "undef."

Hereinafter, this disk will be regarded as an undefined disk, which use is not yet determined.

name=diskname

Sets the name of a disk to *diskname*.

-a *attribute=value* (when using -G)

Sets the *attribute* attribute of the *group* to be *value*.

The *attribute* indicates the attribute name, and the *value* indicates the attribute value.

Always separate *attribute* and *value* with an equal (=) sign.

You can indicate the following combination to *attribute* and *value*.

name=groupname

Sets the name of the group to *groupname*.

-a *attribute=value* (when using -S)

Sets the *attribute* attribute of the detached *slice* to be *value*.

Both attribute values become invalid at the point when the slice is assembled with the volume using the -R option.

The *attribute* indicates the attribute name, and the *value* indicates the attribute value.

Always separate *attribute* and *value* with an equal (=) sign.

Specifiers should be combined using comma (,) as the delimiter.

You can use the following *attribute* and *value* pairs in any combination.

When attempting to set multiple attributes, if any of them result in an error, no attributes are set.

jrm=off

Turns the slice's just resynchronization mechanism mode to off.

It could be turned off regardless to the slice status.

To set the *jrm* to "on," attach the slice to the volume and then detach it again.

mode=rw or mode=ro

Changes the access mode of current node for slice.

It will result in an error when *slice* is activated. Execute after stopping it.

rw

Sets access mode for read and write.

ro

Sets access mode for read only.

Opening a read-only volume in write mode will result in an error.

-a attribute=value[,attribute=value,...] (when using **-V**)

Sets the *attribute* attribute of the *volume* to be *value*.

The *attribute* indicates the attribute name, and the *value* indicates the attribute value.

Always separate *attribute* and *value* with an equal (=) sign.

Specifiers should be combined using comma (,) as the delimiter.

You can use the following *attribute* and *value* pairs in any combination.

When attempting to set multiple attributes, if any of them result in an error, no attributes are set.

jrm=on or jrm=off

Turn the JRM mode on or off.

You can make changes regardless to whether *volume* is activated or not.

When *volume* belongs to either a stripe group or a concatenation group, this option will result in an error.

on

JRM is on.

off

JRM is off.

lock=on or lock=off

Changes the lock mode of current node for *volume*.

You can make changes regardless to whether *volume* is activated or not.

When *class* is a root class, this option will result in an error.

on

The volume is locked from activating thereafter.

off

The volume is not locked from activating thereafter.

mode=rw or mode=ro

Changes the access mode of current node for *volume*.

When *volume* is activated, it results in an error. First, you must stop the *volume*.

rw

Sets access mode for read and write.

ro

Sets access mode for read only.

Opening a read-only volume in write mode will result in an error.

name=volumentname

Sets the name of the volume to *volumentname*.

It will result in an error when *volume* is activated. Execute after stopping *volume*.

pslice=on or pslice=off

Turns the physical slice attribute value to be "on" or "off", respectively.

If *volume* is activated, to avoid an error, stop *volume* before indicating this option. This option will also result in an error if there is a detached slice within *volume*. In such case, attach the slice before indicating this option.

on

The physical slice attribute value of volume is set to be "on".

Among the slices consisting *volume*, any slice on a single disk, and any slices on disks that are directly connected to a mirror group will be registered to the VTOC.

You cannot change this option to "on" when *volume* belongs to a group that cannot create a physical slice (stripe group, concatenation group, or a mirror group whose

only directly-connecting group is a lower level group), or when there is a maximum number (six) of volumes with its physical slice attribute set to "on" within the same group or single disk.

off

The physical slice attribute value of volume is set to be "off".
You cannot change to "off" when *class* is "root" type.

-c class

The *class* indicates the class name to which the object belongs, or is to be changed, where object is the target of the change.

-d disk

The *disk* indicates the disk name that is the target of the change.

-g group

The *group* indicates the group name that is the target of the change.

-s slice

The *slice* indicates the slice name that is the target of the change.
Slice name should be indicated in *disk.volume* format.

-v volume

The *volume* indicates the volume name that is the target of the change.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.
Otherwise, a non-zero value is returned.

D.8 sdxswap - Swap disk

SYNOPSIS

```
sdxswap -I -c class -d disk [ -e delay=msec,nowaitsync ]  
sdxswap -O -c class -d disk
```

DESCRIPTION

Use the `sdxswap` command to set the status of exchangeable disks registered with GDS, or to return swapped disks to their original state after swapping.

You must be super-user to use this command.

This command is mostly used when swapping crashed disks.

PRIMARY OPTIONS

You can use either of the following options.

-I

swapIn

Returns the disk (specified by *disk*) to a usable state, and restores its original status and configuration. You must execute the command after a crashed disk has been swapped.

The *disk* indicates the disk name that is now available for swapping using the `-O` option. The *class* indicates the class name with which the *disk* is registered. If *disk* was connected to a group, slice configuration or volume contents is copied as needed, and returns from the `sdxswap` command once the copying is complete (when using `-e nowaitsync` option, before the copying process).

If a spare disk is substituted in place of *disk*, then once the redundancy of all related volumes has been restored, the spare disk is disconnected.

-O

swapOut

Disconnects all slices within the disk (specified by *disk*), and sets the disk to exchangeable status. This command must be executed before swapping out a crashed disk.

The *class* indicates the class name with which the *disk* is registered.

When *disk* is connected to a group, and the highest level group related to *disk* is not a mirror group, this option will result in an error.

Also, when there are one or more single volumes within *disk*, or when there are one or more volumes within the highest level group related to *disk* and detaching slices within *disk* will change the volume's configuration or status, this option will result in an error.

For example, even if the highest level group is a mirror group, if there is only one mirror slice consisting the mirror volume and if its data is within *disk*, this option will result in an error.

If *disk* is already disabled, a warning message will be sent to standard error output after reconfirming the status.

SUB OPTIONS

SUB OPTIONS are as follows:

-c *class*

The *class* indicates the class name to which *disk* belongs, where *disk* is the target of the operation.

-d *disk*

The *disk* indicates the disk name that is the target of the operation.

-e delay=*msec* (when using **-l**)

When restoring the disk, data contained in the volume is copied as needed.

This option delays the issuing of the input/output request to the disk at the time of copying, by milliseconds indicated by *msec*.

This option allows you to adjust the influence on the application accessing the volume.

The value is set to 0 by default.

Values from 0 to 1000 may be indicated for *msec*.

-e nowaitsync (when using **-l**)

Returns the command before the copying is complete.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.

Otherwise, a non-zero value is returned.

D.9 sdxfix - Restore a failed object

SYNOPSIS

```
sdxfix -D -c class -d disk
sdxfix -V -c class { -g group | -d disk } -v volume
```

DESCRIPTION

sdxfix attempts to restore the state of the disk that detected an I/O error or the state of an inaccessible volume due to a failure. Data consistency and integrity might not be assured as a result of recovering the disk or volume. In such a case, you need to restore integrity from backup data or check by executing the "fsck" command. The "sdxfix" command can be executed with super-user access privileges only.

PRIMARY OPTIONS

You can use one of the following options.

-D

Disk

Restores the state of disk *disk* that detected an I/O error.

Reads all *disk* area, and if there's no problem, then clears the error status.

If there's a volume in the highest-level group to which *disk* belongs, or *disk* has a single volume, the volume should be stopped or inactive (STOP or INVALID) on all nodes.

-V

Volume

Restores an arbitrary slice with invalid data specified in the combination of *disk* and *volume* or *group* and *volume* to the STOP state. Then, the state of the volume with invalid data specified by *volume* is restored to STOP.

The *volume* should be stopped or inactive (STOP or INVALID) on all nodes. The slice state specified in the combination of *disk* and *volume*, or *group* and *volume* should be INVALID.

Reads the entire slice specified in the combination of *disk* and *volume*, or *group* and *volume*, and if there's no problem, changes the state to STOP, and then, changes the state of stopped slice to INVALID.

SUB OPTIONS

You can use the following sub-options.

-c class

Specify a name of the class to which the object belongs.

-d disk (when using -D)

Specify a name of the disk.

-d disk (when using -V)

When *volume* is a mirror volume, specify a name of the disk that is connected to the mirror group to which the *volume* belongs. This disk should have the INVALID mirror slice that needs to be restored to STOP.

Specify a single disk name when *volume* is a single volume.

-g group (when using -V)

When *volume* is a mirror volume, specify a name of lower-level group that is connected to the mirror group to which the *volume* belongs. This group should have the INVALID mirror slice that needs to be restored to STOP.

Specify a name of the highest-level stripe group when *volume* is a stripe volume.

Specify a name of the highest-level concatenation group when *volume* belongs to the highest-level concatenation group.

-v *volume*(when using **-V**)
Specify a name of the volume.

RETURNED VALUE

When it is normally terminated, "0" is returned. Otherwise, a non-zero value is returned.

D.10 sdxcopy - Synchronization copying operation

SYNOPSIS

```
sdxcopy -B -c class -v volume,... [ -e delay=msec,nowaitsync ]  
sdxcopy -C -c class -v volume,...  
sdxcopy -I -c class -v volume,...  
sdxcopy -P -c class -v volume,... -e delay=msec
```

DESCRIPTION

Use sdxcopy to execute a synchronization copying of the volume object specified by *volume*. You must be super-user to use this command.

PRIMARY OPTIONS

You can use either of the following options.

-B

Begin

Assembles and executes a synchronization copying of disconnected slices on the volume or volumes indicated by *volume*,..., returning once the synchronization copying is complete. (when using -e nowaitsync option, before the copying process.) *class* indicates the class name to which the *volume* belongs.

Slices on the *volume* that have a status of INVALID are assembled, and then a synchronization copying is executed.

Slices with a status of TEMP* or NOUSE are not assembled.

If there are any slices currently involved in a synchronization copy on the *volume*, this command will terminate with an error.

This command resumes copying from the point where it was interrupted with the -I option.

Synchronization copying will be executed while the volume is activated or stopped.

-C

Cancel

Cancels synchronization copying in process or interrupted on the volume or volumes indicated by *volume*,.... After the cancel is completed, the system returns from the command. The *class* indicates the class name to which *volume* belongs.

-B option executes the copying process again.

-I

Interrupt

Interrupts the synchronization copying currently underway on a volume or volumes specified by *volume*,....

The command returns after the interrupt is completed.

The *class* indicates the class name to which *volume* belongs.

-B option executes the copying process from the point where it was interrupted.

-P

Parameter

Changes parameter related to the synchronization copying in process or interrupted on the volume or volumes indicated by *volume*,....

The *class* indicates the class name to which *volume* belongs.

Copying in process will resume after the parameter is changed.

The current status of synchronization copying that is either in process or interrupted can be checked by executing the sdxinfo -S command.

SUB OPTIONS

SUB OPTIONS are as follows:

-c *class*

The *class* indicates the class name to which *volume* belongs.

-e *delay=msec* (when using **-B,-P**)

Delays the issuing of the input/output request to the disk at the time of copying by milliseconds indicated by *msec*. This option allows you to adjust the influence on the application accessing the volume.

The value is set to 0 by default.

If the copying process is either completed or canceled, the delay time will return to default (0).

Values from 0 to 1000 may be indicated for *msec*.

-e *nowaitsync* (when using **-B**)

Returns the command before the copying is complete.

-v *volume,...*

The *volume* indicates the volume name that is the target of the operation.

To indicate multiple *volumes*, separate each volume name with a comma (,) as the delimiter.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.

Otherwise, a non-zero value is returned.

D.11 sdxroot - Define or cancel mirroring of root file system

SYNOPSIS

```
sdxroot -M -c class -d disk [ ,disk,... ]  
sdxroot -R -c class -d disk [ ,disk,... ]
```

DESCRIPTION

Use sdxroot to define or cancel mirroring of a system disk with a root file system. You must have super user authority to execute the sdxroot command.

System disk is a physical disk installed with an active Solaris(TM) Operating Environment. Specifically, it refers to the entire disk including a slice which is active as one of the following file systems (or swap area).

/, /usr, /var, or swap

PRIMARY OPTIONS

You can use either of the following options.

-M

Make

Confirms that all preparation to mirror system disk or disks indicated by *disk* is complete, and completes the remaining mirroring definition by updating /etc/vfstab and /etc/system files.

After the sdxroot command is returned, promptly reboot the system to start mirroring.

When there is more than one system disk, for example, when the disk installed with /(root), /usr, and /var, and the disk assigned as swap area differ, you must first prepare each disk for mirroring.

Unless registering them to a class and connecting them to groups are complete, it will result in an error.

-R

Remove

Confirms that all preparation required to cancel mirroring of one or more system disks specified by *disk* is complete, and completes the remaining cancellation of mirroring definition by updating /etc/vfstab and /etc/system files.

After the sdxroot command is returned, promptly reboot the system to cancel mirroring of the system disk.

In order to completely cancel the system disk management by GDS, you must remove the volume, group, disk, and class contained in system disk.

When there is more than one system disk, for example, when the disk installed with /(root), /usr, and /var, and the disk assigned as swap area differ, you must first prepare each disk for unmirroring. Unless other disks are disconnected so that only one system disk is placed in the group, it will result in an error.

SUB OPTIONS

SUB OPTIONS are as follows:

-c *class*

Indicates the class name with which the *disk* is associated.

-d *disk* [,*disk*,...]

Indicates one or more disk names that is the target of operation.

To connect multiple disks, separate each disk name with a comma (,) as the delimiter.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.

Otherwise, a non-zero value is returned.

USAGE EXAMPLES

How to mirror and unmirror when the disk installed with /(root), /usr, and /var, and the disk assigned as swap area differ, are explained below.

Mirroring the System Disk

1. Stop all active applications.

To safely define mirroring, you must stop all active applications. The mirroring definition will be valid upon restarting your system.

When absolute safety is required, system disk should be backed up in advance.

2. Register the system disks with root class.

For this example, the disk installed with /(root), /usr, and /var will be c0t0d0, and the disk assigned as swap area will be c0t0d1.

```
# sdxdisk -M -c System -a type=root -d c0t0d0=Root1:keep,  
c1t0d0=Root2:undef,c0t0d1=Swap1:keep,c1t0d1=Swap2:undef
```

3. Connect system disks to groups.

```
# sdxdisk -C -c System -g Group1 -d Root1,Root2  
-v 0=root:on,1=usr:on,3=var:on,4=home:on,5=opt:on  
# sdxdisk -C -c System -g Group2 -d Swap1,Swap2 -v 0=swap:on
```

4. Confirm that the mirroring definition is complete.

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

5. Reboot the system.

```
# shutdown -y -g0 -i6
```

6. Confirm that mirroring is configured.

Use mount command and the sdxinfo command to check that the system disks are mirrored properly.

Canceling Mirroring of System Disk

1. Stop all active applications.

To safely cancel mirroring, you must stop all active applications.

The cancellation of mirroring will be valid upon restarting your system.

When absolute safety is required, system disk should be backed up in advance.

2. Disconnect all disks from disk group, except for the disk you will use after the cancellation.

```
# sdxdisk -D -c System -g Group1 -d Root2  
# sdxdisk -D -c System -g Group2 -d Swap2
```

3. Confirm that cancellation is complete.

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

4. Reboot the system.

```
# shutdown -y -g0 -i6
```

5. Confirm that the mirroring is canceled.

Use mount command and the sdxinfo command to check that the mirroring of system disks is canceled properly.

6.Cancel the system disk management.

```
# sdxvolume -F -c System -v root
# sdxvolume -F -c System -v usr
# sdxvolume -F -c System -v var
# sdxvolume -F -c System -v home
# sdxvolume -F -c System -v opt
# sdxvolume -F -c System -v swap
# sdxvolume -R -c System -v root
# sdxvolume -R -c System -v usr
# sdxvolume -R -c System -v var
# sdxvolume -R -c System -v home
# sdxvolume -R -c System -v opt
# sdxvolume -R -c System -v swap
# sdxgroup -R -c System -g Group1
# sdxgroup -R -c System -g Group2
# sdxdisk -R -c System -d Root1
# sdxdisk -R -c System -d Root2
# sdxdisk -R -c System -d Swap1
# sdxdisk -R -c System -d Swap2
```

D.12 sdxparam - Configuration parameter operations

SYNOPSIS

```
sdxparam -G [ -p param,... ]  
sdxparam -S -p param=val [ ,param=val,... ] [ -e default ]
```

DESCRIPTION

Use sdxparam to perform operations on GDS configuration parameter.
You must be super-user to use this command.

PRIMARY OPTIONS

You can use either of the following options.

-G

Get

Displays the current value of configuration parameter or parameters indicated by *param*.
When using cluster system, parameter value of the current node will be displayed. If the -p option is omitted, all configuration parameters are displayed.

-S

Set

Sets the value indicated by *val* to the configuration parameter or parameters indicated by *param*.
When using cluster system, this option sets the parameter value of the current node. The new value becomes valid upon returning from the command, and rebooting the system will not change the value.

SUB OPTIONS

SUB OPTIONS are as follows:

-e default (when using **-S**)

Resets all configuration parameter values to default.
When indicated at the same time as -p option, this option is ignored.

-p *param*,... (when using **-G**)

Displays the configuration parameter *param* value.

-p *param=val*[,*param=val*,...] (when using **-S**)

Sets *val* to configuration parameter *param*.

You can indicate the following combination to *param* and *val*.

copy_concurrency=*num*

Sets the maximum number of synchronization copying you can execute simultaneously to *num*.

The value is set to 8 by default.

Values from 1 to 1024 may be indicated for *num*.

copy_delay=*msec*

This option delays the synchronization copying by milliseconds indicated by *msec*, when the copying is initiated by an event other than hot spare.

The value is set to 0 by default.

Values from 0 to 1000 may be indicated for *msec*.

spare_copy_delay=*msec*

This option delays the synchronization copying by milliseconds indicated by *msec*, when the copying is initiated by hot spare.

The value is set to 50 by default. Values from 0 to 1000 may be indicated for *msec*.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.

Otherwise, a non-zero value is returned.

CAUTION

The default values and the range of values you can use for configuration parameters may change in the future.

D.13 sdxproxy - Proxy object operations

SYNOPSIS

sdxproxy Break -c *class* -p *proxy* [-e *force*]

sdxproxy Join -c *class* -m *master* -p *proxy*
[-a *mvol=pvol:jrm[:pslice][,mvol=pvol:jrm[:pslice],...]*]
[-e *delay=msec,softcopy,waitsync*]

sdxproxy Part -c *class* -p *proxy*,... [-e *instant,mode=val,unlock*]

sdxproxy Rejoin -c *class* -p *proxy*,... [-e *delay=msec,softcopy,waitsync*]

sdxproxy RejoinRestore -c *class* -p *proxy*,...
[-e *delay=msec,instant,nowaitsync,softcopy*]

sdxproxy Swap -c *class* -p *proxy*

DESCRIPTION

Use sdxproxy to perform operations on proxy objects. You must have super user authority to execute the sdxproxy command.

PRIMARY OPTIONS

You can use one of the following options.

Break

Cancels master-proxy relationship between a pair of volumes or groups. You can cancel the relationship when the objects are either joined or parted.

proxy indicates a proxy volume or a proxy group to be cancelled. You cannot indicate a proxy volume within a proxy group for the cancellation.

Even after the relationship has been cancelled, both the proxy volume and the proxy volume within a proxy group can be used as ordinary volumes retaining their volume data and attributes.

You can also cancel the relationship when master volume and proxy volume are in use. However, if the joined master volume is in use, the data integrity of the proxy volume should be ensured by either the file system layer or database layer whichever is managing the data, after the relationship is cancelled. For instance, if you are using the master volume as a file system, use the lockfs(1M) command to write-lock the file system, before canceling the relationship.

The command will result in an error when:

- copying is in process from the master volume to the proxy volume (except for when using -e force)
- copying is in process from the proxy volume to the master volume
- the master volume is in an INVALID status

Join

A pair of volumes or a pair of groups are related and joined as master and proxy. The volumes to be joined should be mirror or single type, and the groups to be joined should be mirror type.

When joining a pair of volumes, synchronization copying of master to proxy volume is performed after returning from the command (when using -e waitsync, before returning from the command). The size of the master and proxy volume must be the same. Also, they must be volumes belonging to different groups.

When joining a pair of groups, proxy volumes with the same offset and size as the master volumes within master group will be created in the proxy group. Synchronization copying from

master to proxy volume is also performed after returning from the command (when using `-e waitsync`, before returning from the command). The access mode of the proxy volume that will be created in the proxy group is set to `ro` (read-only). The size of a proxy group must be larger than the last block number of volumes within the master group to which the proxy is related. When there is no volume in a master group or there is already a volume in a proxy group, joining cannot be performed resulting in an error.

When hard copy function is available, the synchronization copying from master to proxy will be performed using that function (except for when using `-e softcopy`).

Proxy volume that is joined cannot be accessed or activated. In order to access the proxy volume, part the volume from master using the `Part` option, or break the relationship with master using the `Break` option.

You can create multiple snapshots by joining another proxy to a master, which is already joined with a proxy. However, the total number of slices which belong to the master volume and slices which belong to the proxy volumes that are related to the master volume, must not exceed thirty-two.

Master volumes that are already related to a proxy cannot be joined to another master as a proxy. Also, a proxy that is already related to a master cannot be joined with another proxy.

The command will result in an error when:

- the proxy volume is activated
- there is a slice that is temporarily detached or in copying process among any of the slices comprising the proxy volume or the master volume
- copying is in process between the master volume and the other proxy volume
- the master volume is in an `INVALID` status

Part

Separates a proxy or proxies in the joined status from the master. The master-proxy relationship will be maintained after parted. The parted proxy volume will be the snapshot containing the copy of the master volume data at the time of parting. By using the parted proxy volume, you can for instance, create a backup of the master volume at the time it was parted, or use it for other purposes.

proxy indicates a proxy volume or a proxy group in joined status. When proxy group is indicated, all proxy volumes in the group will be parted. You can also indicate a proxy volume in joined status in a proxy group.

After the parting is complete, the proxy volume will be activated as an independent volume and become accessible using the following special file.

`/dev/sfdsk/class/[r]dsk/volume_name`

When the proxy volume belongs to a "shared" type class, it will be activated on all nodes defined in the class scope.

You can part master and proxy volumes even if the master volume is activated, but the data integrity of the parted proxy volume must be ensured by the file system layer or the database layer that is managing the data. For instance, if you are using the master volume as a file system, use the `lockfs(1M)` command to write-lock the file system, before parting.

The command will result in an error when:

- copying is in process from the master volume to the proxy volume (except for when using `-e instant`)
- copying is in process from the proxy volume to the master volume

Rejoin

Rejoins one or more parted proxy with the master.

proxy indicates parted proxy volume, or proxy group. When proxy group is indicated, all proxy volumes in the group will be rejoined. You can also indicate the parted proxy volume in a proxy group.

Synchronization copying of master to proxy volume is performed after returning from the command (when using `-e waitsync`, before returning from the command). When hard copy function is available, the synchronization copying from master to proxy will be performed using that function (except for when using `-e softcopy`).

When more than one proxy volume related to the same master volume is indicated simultaneously, this command will result in an error.

The command will result in an error when:

- the proxy volume is activated
- there is a slice that is in copying process among any of the slices comprising the proxy volume or the master volume
- copying is in process between the master volume and the other proxy volume
- the master volume is in an INVALID status

RejoinRestore

Rejoins a proxy or proxies in the parted status with the master and restores the master volume data using the proxy volume data. Master volume data are restored by synchronization copying from the proxy volume to the master volume. When executing the command using this option, the master volume data are overwritten with the proxy volume data.

proxy indicates a proxy volume or a proxy group in parted status. When a proxy group is indicated, all proxy volumes in the group will be rejoined and data of the related master volumes will be restored. You can also indicate a proxy volume in parted status in a proxy group.

Synchronization copying of master to proxy volume is performed before returning from the command (when using *-e waitsync*, after returning from the command). When hard copy function is available, the synchronization copying from proxy to master will be performed using that function (except for when using *-e softcopy*).

When more than one proxy volume related to the same master volume is indicated simultaneously, this command will result in an error.

The command will result in an error when:

- the master volume or the proxy volume is activated
- there is a slice that is in copying process among any of the slices comprising the proxy volume or the master volume
- copying is in process between the master volume and the other proxy volume
- the master volume is in an INVALID status

Swap

Swaps the master's slices with the proxy's slices.

proxy indicates a proxy volume or a proxy group in the joined status. You cannot indicate a proxy volume within a proxy group for swapping.

You can swap the slices when master is in use.

The command will result in an error when:

- there is a slice that is in copying process among any of the slices comprising the proxy volume or the master volume
- copying is in process between the master volume and the other proxy volume
- the master volume is in an INVALID status

SUB OPTIONS

You can indicate the following SUB OPTIONS.

-a *mvol=pvol:jrm [:pslice] [,mvol=pvol:jrm [:pslice] ,...]* (when using **Join**)

Indicates the proxy volume attributes.

mvol must always be followed by an equal (=) sign, and *pvol*, *jrm* and *pslice* are delimited by a colon (:). When indicating attributes for more than one proxy volume, combine specifiers with a comma (,) as the delimiter.

When joining a pair of groups, *mvol* indicates the volume name of the master volume within the master group. *pvol* indicates the volume name of the proxy volume, that will be created in the proxy group to correspond with the master volume indicated by *mvol*. *jrm* indicates the just resynchronization mechanism mode for the proxy volume, which could either be "on" or "off." *pslice* indicates the physical slice attribute, which could either be "on" or "off." You must indicate attributes for each proxy volumes created for every master volume within master group. If *:pslice* is omitted, the physical slice attribute of the proxy volume will be equivalent

to that of the corresponding master volume.

When joining a pair of volumes, *mvol* indicates the volume name of the master volume, *pvol* the volume name of the proxy volume, *jrm* the just resynchronization mechanism mode for the proxy volume, which could either be "on" or "off", and *pslice* the physical slice attribute, which could either be "on" or "off." *mvol* and *pvol* should match with *master* and *proxy* respectively. When not using this option, the proxy volume attributes will be the same as before it was joined.

-c class

class indicates the class name to which the master object or proxy object that will be the target of operation belongs

-e delay=*msec* (when using **Join, Rejoin, RejoinRestore)**

Delays the issuing of the input/output request to the disk at the time of synchronization copying between master volume and proxy volume by milliseconds, indicated by *msec*.

This option allows users to adjust the influence on the application accessing the master volume.

When using the hard copy function, this option will be ignored.

When this option is omitted, the delay will be set to 0.

Values from 0 to 1000 may be indicated for *msec*.

-e force (when using **Break)**

Forcibly breaks master-proxy relationship even when copying is in process between master and proxy.

When using this option, copying process will be cancelled and the status of the proxy volume will become INVALID.

-e instant (when using **Part)**

The proxy volume will be parted and a virtual snapshot of the master volume will be created even if copying is in process from master to proxy. After returning from the command, the parted proxy volume will become accessible before the copying is complete, and will serve as a snapshot containing data of the master volume at the time of parted. When copying is in process from proxy to master, or if hard copy function is not available, the command will result in an error.

-e instant (when using **RejoinRestore)**

Restoration will be completed instantly, and the command will be returned. After returning from the command, synchronization copying from proxy volume to master volume will automatically begin. Although the copying is still in process, restoration will appear to be complete. After returning from the command, the master volume can be activated and accessed, before the copying is complete. Master volume data will appear to have been overwritten by the proxy volume data at the time of executing the command. When used with `-e nowaitsync` option, the `-e nowaitsync` is ignored.

-e mode=*val* (when using **Part)**

Indicates the access mode of the proxy volume which will be activated. *val* indicates either of the following options.

rw

Sets access mode for read and write.

ro

Sets access mode for read only.

Opening a read-only volume in write mode will result in an error.

Although proxy volume will be activated in the access mode specified by *val*, the proxy volume's access mode attribute will remain unchanged. Access mode indicated by *val* is valid only while the proxy volume is activated ("Current Access Mode"), and will become invalid once the proxy volume is stopped. When the proxy volume is restarted, it will start in the mode set by access mode attribute ("Default Access Mode"), except for when the access mode is specified at the time of restart.

When this option is omitted, proxy volume will be activated in the access mode according to the access mode attribute set on each node.

-e nowaitsync (when using **RejoinRestore**)

Returns from the command before the synchronization copying of the proxy volume to master volume is complete. After returning from the command, master volume can not be activated until the copying is complete. In order to activate the master volume before the copying is complete, use the **-e instant** option. When used with **-e instant** option, the **-e nowaitsync** is ignored.

-e softcopy (when using **Join**, **Rejoin** or **RejoinRestore**)

Hard copy function will not be used for synchronization copying between master volume and proxy volume.

-e unlock (when using **Part**)

The proxy volume will be activated regardless to whether it is locked. Lock mode will not be changed unless you change it with the **sdxattr -V** command.

-e waitsync (when using **Join** or **Rejoin**)

When performing synchronization copying, returns from the command after the copying is complete.

-m master (when using **Join**)

master indicates a master volume or a master group to be joined.

master indicates the volume name of the master volume, or the group name of the master group.

-p proxy (when using **Break**, **Join** or **Swap**)

proxy indicates a proxy volume or a proxy group that is the target of the operation.

proxy indicates the volume name of the proxy volume, or the group name of the proxy group.

-p proxy,... (when using **Part**, **Rejoin** or **RejoinRestore**)

proxy indicates one or more proxy volumes or proxy groups that will be the target of the operation. When indicating more than one volume or group, they must belong in the same class.

proxy indicates the volume name of the proxy volume or the group name of the proxy group that will be the target of the operation

When indicating more than one volume name or group name, combine them with a comma (,) as the delimiter.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.

Otherwise, a non-zero value is returned.

D.14 Volume Creation Using Command

In this section, operation outline of volume creation is explained. Please use as a reference when configuring the environment. For details, see the Command Reference.



Information

For procedures to mirror a system disk, see the examples given in "D.11 sdxroot – Define or cancel mirroring of root file system."



Note In PRIMECLUSTER Systems

In order to define the configuration of GDS objects such as classes and volumes, PRIMECLUSTER resources must be registered in advance. For details on registering resources, see "PRIMECLUSTER Cluster Foundation Configuration and Administration".

(1) Creating a mirror volume

The procedure which carries out mirroring of physical disks c1t0d0 and c2t0d0, and creates a volume as an example is shown.

1) Registering disks to class

Register the physical disks with a class. When the specified class does not exist, it will be created automatically.

Example: Registering physical disks c1t0d0 and c2t0d0 to local class "Local", and assigning the names "Disk1" and "Disk2" respectively.

```
# sdxdisk -M -a type=local -c Local -d c1t0d0=Disk1,c2t0d0=Disk2
```



Information

When registering disks to shared class, you must specify the "scope" attribute with the -a option. An example is given below.

```
# sdxdisk -M -c Shared -a type=shared,scope=node1:node2  
-d c1t0d0=Disk1,c2t0d0=Disk2
```

2) Connecting the disks to a mirror group

Connect the disks to a mirror group. When the specified mirror group does not exist, it will be created automatically.

Example: Connecting "Disk1" and "Disk2" to mirror group "Group1".

```
# sdxdisk -C -c Local -g Group1 -d Disk1,Disk2
```

3) Creating a mirror volume

Create a volume within the mirror group.

Example: Creating a volume of 1,000 blocks within mirror group "Group1", and assigning a volume name "Volume1".

```
# sdxvolume -M -c Local -g Group1 -v Volume1 -s 1000
```

After returning from the command, synchronization copying will automatically begin.

4) Confirming the completion of procedure

Confirm that the synchronization copying is complete.

Example: Confirming synchronization copying of volume "Volume1" is complete.

OBJ	CLASS	GROUP	DISK	VOLUME	STATUS
slice	Local	Group1	Disk1	Volume1	ACTIVE
slice	Local	Group1	Disk2	Volume1	COPY

If the all the displayed slices' STATUS fields are "ACTIVE," synchronization copying is complete.

If the synchronization copying is still in progress, "COPY" will be displayed in the STATUS field.

Using the -e long option, you can check the progress of the synchronization copying.

For details, see "D.6 sdxinfo – Display object configuration and status information."

(2) Creating a single volume

The procedure which creates a single volume as an example using physical disk c1t0d0 is shown.

1) Registering a disk to class

Register the physical disk with a class. When the specified class does not exist, it will be created automatically.

Example: Registering physical disk c1t0d0 to shared class "Shared", and assigning the name "Disk1".

```
# sdxdisk -M -c Shared -a type=shared,scope=node1:node2 -d c1t0d0=Disk1:single
```

2) Creating a single volume

Create a volume within the single disk.

Example: Creating a volume of 1,000 blocks within single disk "Disk 1", and assigning a volume name "Volume1".

```
# sdxvolume -M -c Shared -d Disk1 -v Volume1 -s 1000
```

(3) Creating a large-capacity volume (using concatenation)

The procedure which concatenates physical disks c1t0d0 and c2t0d0, and creates a volume as an example is shown.

1) Registering disks to class

Register the physical disks with a class. When the specified class does not exist, it will be created automatically.

Example: Registering physical disks c1t0d0 and c2t0d0 to shared class "Shared", and assigning the names "Disk1" and "Disk2" respectively.

```
# sdxdisk -M -c Shared -a type=shared,scope=node1:node2  
-d c1t0d0=Disk1,c2t0d0=Disk2
```

2) Connecting the disks to a concatenation group

Connect the disks to a concatenation group. When the specified concatenation group does not exist, it will be created automatically.

Example: Connecting "Disk1" and "Disk2" to concatenation group "Group1".

```
# sdxdisk -C -c Shared -g Group1 -d Disk1,Disk2 -a type=concat
```

3) Creating a large-capacity volume

Create a volume within the concatenation group.

Example: Creating a volume of 1,000,000,000 blocks within concatenation group "Group1", and assigning a volume name "Volume1".

```
# sdxvolume -M -c Shared -g Group1 -v Volume1 -s 1000000000 -a pslice=off
```

(4) Creating a stripe volume

The procedure which carries out striping of physical disk c1t0d0 and c2t0d0, and creates volume as an example is shown.

1) Registering disks to class

Register the physical disks with a class. When the specified class does not exist, it will be created automatically.

Example: Registering physical disks c1t0d0 and c2t0d0 to shared class "Shared", and assigning the names "Disk1" and "Disk2" respectively.

```
# sdxdisk -M -c Shared -a type=shared,scope=node1:node2 -d  
c1t0d0=Disk1,c2t0d0=Disk2
```

2) Connecting the disks to a stripe group

Connect the disks to a stripe group. When the stripe group does not exist, it will be created automatically.

Example: Connecting "Disk1" and "Disk2" to stripe group "Group1".

```
# sdxdisk -C -c Shared -g Group1 -d Disk1,Disk2 -a type=stripe,width=32
```

3) Creating a stripe volume

Create a volume within the stripe disk.

Example: Creating a volume of 1,000 blocks within stripe group "Group1", and assigning a volume name "Volume1".

```
# sdxvolume -M -c Shared -g Group1 -v Volume1 -s 1000 -a pslice=off
```

(5) Creating a mirror volume (Combining striping and mirroring)

Example: Configuring stripe groups with physical disks c1t0d0 and c2t0d0, and c3t0d0 and c4t0d0 respectively, and creating a volume by mirroring those two stripe groups.

1) Registering disks to class

Register the physical disks with a class. When the specified class does not exist, it will be created automatically.

Example: Registering physical disks c1t0d0, c2t0d0, c3t0d0 and c4t0d0 to shared class "Shared", and assigning the names "Disk1", "Disk2", "Disk3" and "Disk4" respectively.

```
# sdxdisk -M -c Shared -a type=shared,scope=node1:node2 -d  
c1t0d0=Disk1,c2t0d0=Disk2,c3t0d0=Disk3,c4t0d0=Disk4
```

2) Connecting the disks to a stripe group

Connect the disks to a stripe group. When the stripe group does not exist, it will be created automatically.

Example: Connecting "Disk1" and "Disk2" to stripe group "Group1".

```
# sdxdisk -C -c Shared -g Group1 -d Disk1,Disk2 -a type=stripe,width=32
```

Connecting Disk3 and Disk4 to stripe group Group2.

```
# sdxdisk -C -c Shared -g Group2 -d Disk3,Disk4 -a type=stripe,width=32
```

3) Connecting the stripe groups to a mirror group

Connect the stripe groups to a mirror group. When the specified mirror group does not exist, it will be created automatically.

Example: Connecting the stripe group "Group1" and "Group2" to mirror group "Group3".

```
# sdxgroup -C -c Shared -h Group3 -l Group1,Group2 -a type=mirror
```

The "-a type=mirror" option is omissible.

4) Creating a mirror volume

Create a volume within the highest level mirror group.

Example: Creating a volume of 1,000 blocks within mirror group "Group3", and assigning a volume name "Volume1"

```
# sdxvolume -M -c Shared -g Group3 -v Volume1 -s 1000 -a pslice=off
```

After returning from the command, synchronization copying will automatically begin.

5) Confirming the completion of procedure

Confirm that the synchronization copying is complete.

Example: Confirming synchronization copying of volume "Volume1" is complete.

```
# sdxinfo -S -o Volume1
OBJ CLASS GROUP DISK VOLUME STATUS
-----
slice Shared Group3 Group1 Volume1 ACTIVE
slice Shared Group3 Group2 Volume1 COPY
```

If the all the displayed slices' STATUS fields are "ACTIVE," synchronization copying is complete. If the synchronization copying is still in progress, "COPY" will be displayed in the STATUS field. Using the -e long option, you can check the progress of the synchronization copying.

For details, see "D.6 sdxinfo – Display object configuration and status information."

D.15 Snapshot Creation Using Command

In this section, operation outline of snapshot creation using command is explained. Please use it as a reference when configuring the environment. For details, see the Command Reference.

1) Joining the proxy volume with master volume

Join the proxy volume with the master volume, and copy the master volume data to the proxy volume. Before joining them, stop the proxy volume.

Example: Master volume Volume1 and proxy volume Volume2 will be joined.

```
# /usr/sbin/sdxproxy Join -c Local -m Volume1 -p Volume2
```

After returning from the command, synchronization copying will automatically be performed.



Information

-m option and -p option can indicate groups as well as volumes. When indicating a group, all volumes within the group will be copied.

2) Confirming the completion of copying

Confirm that the synchronization copying is complete.

Example: Confirming synchronization copying of proxy volume (Volume2) is complete.

```
# /usr/sbin/sdxinfo -S
OBJ CLASS GROUP DISK VOLUME STATUS
-----
slice Local Group1 Disk1 Volume1 ACTIVE
slice Local Group1 Disk2 Volume1 ACTIVE
slice Local Group2 Disk3 Volume2 STOP
slice Local Group2 Disk4 Volume2 STOP
```

If all the displayed slices' STATUS fields are "STOP," synchronization copying is complete. If the synchronization copying is still in progress, "COPY" will be displayed in the STATUS field.

3) Creating a snapshot

In order to create a snapshot, part the proxy volume from the master volume after confirming that synchronization copying is complete.

```
# /usr/sbin/sdxproxy Part -c Local -p Volume2
```



Note

In order to create a snapshot properly, you should either stop the application, or secure integrity at the file system layer or database layer that is managing the data. For instance, if you are using the file system, use the lockfs(1M) command to write-lock the file system before separating the proxy volume from the master volume.

4) Creating a backup

In order to make a data backup from the snapshot you created, follow the procedures of whatever backup tool you are using.

If you are not using the snapshot function after creating the backup, cancel the master-proxy relationship.

5) Rejoining the proxy volume with master volume

In order to synchronize the proxy volume with the master volume again, remove the created snapshot, and rejoin the proxy volume with the master volume. Before rejoining them, stop the proxy volume.

```
# /usr/sbin/sdxvolume -F -c Local -v Volume2  
# /usr/sbin/sdxproxy Rejoin -c Local -p Volume2
```

Synchronization copying will be performed as in procedure 1.

When creating more than one backup using the same proxy volume, repeat procedures 2 to 4.

6) Canceling the master-proxy relationship

Cancel the master-proxy relationship, and finish using the snapshot function.

```
# /usr/sbin/sdxproxy Break -c Local -p Volume2
```

D.16 Online Backup Making Using Command

In this section, operation outline of creating a backup with snapshot is explained.

- Making a ufs file system backup using tar
- Making a ufs file system backup using ufsdump
- Making a Global File Services file system backup using sfxdump (single partition)
- Making a Global File Services file system backup using sfxdump (multi partition)

1. Making a ufs file system backup using tar

How to create a backup of Local class volume (Volume1) with ufs file system, using the tar command is explained.

1) Joining the proxy volume

In order to create a snapshot, join the mirror proxy volume (Volume2) with the original master volume (Volume1), and synchronize master and proxy.

```
# /usr/sbin/sdxproxy Join -c Local -m Volume1 -p Volume2
```

2) Confirming the completion of copying

Confirm that the synchronization copying is complete.

Example: Confirming synchronization copying of proxy volume (Volume2) is complete.

```
# /usr/sbin/sdxinfo -S
OBJ CLASS GROUP DISK VOLUME STATUS
-----
slice Local Group1 Disk1 Volume1 ACTIVE
slice Local Group1 Disk2 Volume1 ACTIVE
slice Local Group2 Disk3 Volume2 STOP
slice Local Group2 Disk4 Volume2 STOP
```

If all the displayed slices' STATUS fields are "STOP," synchronization copying is complete.

If the synchronization copying is still in progress, "COPY" will be displayed in the STATUS field.

3) Creating a snapshot

Create a snapshot after confirming that synchronization copying is complete. Snapshot will be created by separating the proxy volume (Volume2) from the master volume (Volume1).

In order to create a snapshot safely, follow one of procedures below and prevent the data from being accessed by an application.

i) Creating a snapshot after unmounting the file system

Unmount the ufs file system from the mount point (/DATA), before creating the snapshot.

- Unmount the file system.

```
# cd /
# /usr/sbin/umount /DATA
```

- Create a snapshot.

```
# /usr/sbin/sdxproxy Part -c Local -p Volume2
```

- Mount the file system again.

```
# /usr/sbin/mount -F ufs /dev/sfdsk/Local/dsk/Volume1 /DATA
```

ii) Creating a snapshot while the file system is mounted

Write-lock the ufs file system mounted on the mount point (/DATA), before creating the snapshot.

- Write-lock the file system.

```
# /usr/sbin/lockfs -w /DATA
```

- Create a snapshot.

```
# /usr/sbin/sdxproxy Part -c Local -p Volume2
```

- Release the file system.

```
# /usr/sbin/lockfs -u /DATA
```

4) Mounting the created snapshot

When mounting the created snapshot, perform the following operation on the ufs file system.

i) Preparing to mount

- Execute fsck.

```
# /usr/sbin/fsck -F ufs /dev/sfdsk/Local/rdisk/Volume2
```

ii) Mounting the snapshot

- Execute mount.

```
# mkdir /DATA_backup  
# /usr/sbin/mount -F ufs /dev/sfdsk/Local/dsk/Volume2 /DATA_backup
```

5) Making a backup from the created snapshot

When making a data backup from the snapshot you created, follow the procedures of whatever backup tool you are using.

In this example, how to make a backup on a tape media using the tar command is explained.

```
# cd /DATA_backup  
# tar cvf /dev/rmt/0 *
```

6) Removing the snapshot

Remove the snapshot you created, and cancel proxy volume (Volume2) settings.

i) Unmounting the snapshot (When the snapshot is mounted.)

```
# cd /  
# /usr/sbin/umount /DATA_backup
```

ii) Canceling the master-proxy relationship

```
# /usr/sbin/sdxvolume -F -c Local -v Volume2  
# /usr/sbin/sdxproxy Break -c Local -p Volume2
```

2. Making a ufs file system backup using ufsdump

How to create a backup of Local class volume (Volume1) with ufs file system, using the ufsdump command is explained.

1) Joining the proxy volume

In order to create a snapshot, join the mirror proxy volume (Volume2) with the original master volume (Volume1), and synchronize master and proxy.

```
# /usr/sbin/sdxproxy Join -c Local -m Volume1 -p Volume2
```

2) Confirming the completion of copying

Confirm that the synchronization copying is complete.

Example: Confirming synchronization copying of proxy volume (Volume2) is complete.

```
# /usr/sbin/sdxinfo -S
OBJ CLASS GROUP DISK VOLUME STATUS
-----
slice Local Group1 Disk1 Volume1 ACTIVE
slice Local Group1 Disk2 Volume1 ACTIVE
slice Local Group2 Disk3 Volume2 STOP
slice Local Group2 Disk4 Volume2 STOP
```

If all the displayed slices' STATUS fields are "STOP," synchronization copying is complete.

If the synchronization copying is still in progress, "COPY" will be displayed in the STATUS field.

3) Creating a snapshot

Create a snapshot after confirming that synchronization copying is complete. Snapshot will be created by separating the proxy volume (Volume2) from the master volume (Volume1).

In order to create a snapshot safely, follow one of procedures below and prevent the data from being accessed by an application.

i) Creating a snapshot after unmounting the file system

Unmount the ufs file system from the mount point (/DATA), before creating the snapshot.

- Unmount the file system.

```
# cd /
# /usr/sbin/umount /DATA
```

- Create a snapshot.

```
# /usr/sbin/sdxproxy Part -c Local -p Volume2
```

- Mount the file system again.

```
# /usr/sbin/mount -F ufs /dev/sfdsk/Local/dsk/Volume1 /DATA
```

ii) Creating a snapshot while the file system is mounted

Write-lock the ufs file system mounted on the mount point (/DATA), before creating the snapshot.

- Write-lock the file system.

```
# /usr/sbin/lockfs -w /DATA
```

- Create a snapshot.

```
# /usr/sbin/sdxproxy Part -c Local -p Volume2
```

- Release the file system.

```
# /usr/sbin/lockfs -u /DATA
```

4) Making a backup from the created snapshot

When making a data backup from the snapshot you created, follow the procedures of whatever backup tool you are using.

In this example, how to make a backup on a tape media using the ufsdump command is explained.

i) Preparation to make a backup

```
# /usr/sbin/fsck -F ufs /dev/sfdsk/Local/rdisk/Volume2
```

ii) Making a backup

```
# ufsdump 0ucf /dev/rmt/0 /dev/sfdsk/Local/rdisk/Volume2
```

5) Removing the snapshot

Remove the snapshot you created, and cancel proxy volume (Volume2) settings.

```
# /usr/sbin/sdxvolume -F -c Local -v Volume2
# /usr/sbin/sdxproxy Break -c Local -p Volume2
```

3. Making a Global File Services file system backup using sfxdump (single partition)

How to create a backup of Local class volume (Volume1) with Global File Services file system, using the sfxdump command is explained.

1) Joining the proxy volume

In order to create a snapshot, join the mirror proxy volume (Volume2) with the original master volume (Volume1), and synchronize master and proxy.

```
# /usr/sbin/sdxproxy Join -c Local -m Volume1 -p Volume2
```

2) Confirming the completion of copying

Confirm that the synchronization copying is complete.

Example: Confirming synchronization copying of proxy volume (Volume2) is complete.

```
# /usr/sbin/sdxinfo -S
OBJ CLASS GROUP DISK VOLUME STATUS
-----
slice Local Group1 Disk1 Volume1 ACTIVE
slice Local Group1 Disk2 Volume1 ACTIVE
slice Local Group2 Disk3 Volume2 STOP
slice Local Group2 Disk4 Volume2 STOP
```

If all the displayed slices' STATUS fields are "STOP," synchronization copying is complete.

If the synchronization copying is still in progress, "COPY" will be displayed in the STATUS field.

3) Creating a snapshot

Create a snapshot after confirming that synchronization copying is complete. Snapshot will be created by separating the proxy volume (Volume2) from the master volume (Volume1).

In order to create a snapshot safely, follow one of procedures below and prevent the data from being accessed by an application.

i) Creating a snapshot after unmounting the file system

Unmount the Global File Services file system from the mount point (/DATA), before creating the snapshot.

- Unmount the file system.

```
# cd /
# /usr/sbin/umount /DATA
```

- Create a snapshot.

```
# /usr/sbin/sdxproxy Part -c Local -p Volume2
```

- Mount the file system again.

```
# /usr/sbin/mount -f sfxfs /dev/sfdsk/Local/dsk/Volume1 /DATA
```

ii) Creating a snapshot while the file system is mounted

Write-lock the Global File Services file system mounted on the mount point (/DATA), before creating the snapshot.

- Write-lock the file system.

```
# /usr/sbin/lockfs -w /DATA
```

- Create a snapshot.

```
# /usr/sbin/sdxproxy Part -c Local -p Volume2
```

- Release the file system.

```
# /usr/sbin/lockfs -u /DATA
```

4) Making a backup from the created snapshot

When making a data backup from the snapshot you created, follow the procedures of whatever backup tool you are using.

In this example, how to make a backup on a tape media using the sfxdump command is explained.

i) Preparation to make a backup

- Update the configuration information.

Update the Global File Services configuration information in the volume.

```
# /usr/sbin/sfxadm /dev/sfdsk/Local/rdsk/Volume2
```



Note

For details on sfxadm, see the Global File Services Manual.

— Execute fsck.

```
# /usr/sbin/fsck -F sfxfs /dev/sfdsk/Local/rdisk/Volume2
```

ii) Making a backup

```
# /usr/sbin/sfxdump 0ucf /dev/rmt/0 /dev/sfdsk/Local/rdisk/Volume2
```

5) Removing the snapshot

Remove the snapshot you created, and cancel proxy volume (Volume2) settings.

```
# /usr/sbin/sdxvolume -F -c Local -v Volume2
# /usr/sbin/sdxproxy Break -c Local -p Volume2
```

4. Making a Global File Services file system backup using sfxdump (multi partition)

How to create a backup of multiple Local class volumes (Volume1, Volume2 and Volume 3) with Global File Services file system, using the sfxdump command is explained.

1) Joining the proxy volume

In order to create a snapshot, join the mirror proxy volume (Volume4, Volume5 and Volume6) with the original master volume (Volume1, Volume2 and Volume3), and synchronize master and proxy.

```
# /usr/sbin/sdxproxy Join -c Local -m Volume1 -p Volume4
# /usr/sbin/sdxproxy Join -c Local -m Volume2 -p Volume5
# /usr/sbin/sdxproxy Join -c Local -m Volume3 -p Volume6
```

When using more than one volume as a file system with Global File Services, you must execute the above command for all volumes.

2) Confirming the completion of copying

Confirm that the synchronization copying is complete.

Example: Confirming synchronization copying of proxy volume (Volume4, Volume5 and Volume6) is complete.

```
# /usr/sbin/sdxinfo -S
```

OBJ	CLASS	GROUP	DISK	VOLUME	STATUS
slice	Local	Group1	Disk1	Volume1	ACTIVE
slice	Local	Group1	Disk2	Volume1	ACTIVE
slice	Local	Group1	Disk1	Volume2	ACTIVE
slice	Local	Group1	Disk2	Volume2	ACTIVE
slice	Local	Group1	Disk1	Volume3	ACTIVE
slice	Local	Group1	Disk2	Volume3	ACTIVE
slice	Local	Group2	Disk3	Volume4	STOP
slice	Local	Group2	Disk4	Volume4	STOP
slice	Local	Group2	Disk3	Volume5	STOP
slice	Local	Group2	Disk4	Volume5	STOP
slice	Local	Group2	Disk3	Volume6	STOP
slice	Local	Group2	Disk4	Volume6	STOP

If all the displayed slices' STATUS fields are "STOP," synchronization copying is complete. If the synchronization copying is still in progress, "COPY" will be displayed in the STATUS field.

3) Creating a snapshot

Create a snapshot after confirming that synchronization copying is complete. Snapshot will be created by separating the proxy volume (Volume4, Volume5 and Volume6) from the master volume (Volume1, Volume2 and Volume3).

In order to create a snapshot safely, follow one of procedures below and prevent the data from being accessed by an application.

i) Creating a snapshot after unmounting the file system

Unmount the Global File Services file system from the mount point (/DATA), before creating the snapshot.

- Unmount the file system.

```
# cd /  
# /usr/sbin/umount /DATA
```

- Create a snapshot.

When using more than one volume as a file system with Global File Services, you must separate all proxy volumes (Volume4, Volume5 and Volume6).

```
# /usr/sbin/sdxproxy Part -c Local -p Volume4,Volume5,Volume6
```

- Mount the file system again.

When using more than one volume as a file system with Global File Services, mount the representative partition (Volume1) you specified when creating the Global File Services file system.

```
# /usr/sbin/mount -F sfxfs /dev/sfdsk/Local/dsk/Volume1 /DATA
```

ii) Creating a snapshot while the file system is mounted

Write-lock the Global File Services file system mounted on the mount point (/DATA), before creating the snapshot.

- Write-lock the file system.

```
# /usr/sbin/lockfs -w /DATA
```

- Create a snapshot.

When using more than one volume as a file system with Global File Services, you must separate all proxy volumes (Volume4, Volume5 and Volume6).

```
# /usr/sbin/sdxproxy Part -c Local -p Volume4,Volume5,Volume6
```

- Release the file system.

```
# /usr/sbin/lockfs -u /DATA
```

4) Making a backup from the created snapshot

When making a data backup from the snapshot you created, follow the procedures of whatever backup tool you are using.

In this example, how to make a backup on a tape media using the sfxdump command is explained.

i) Preparation to make a backup

- Update the configuration information.

Update the Global File Services configuration information in the volume.

In this example, proxy volume (Volume4) will be regarded as the Global File Services File system's representative partition.

```
# /usr/sbin/sfxadm /dev/sfdsk/Local/rdsk/Volume4,  
/dev/sfdsk/Local/rdsk/Volume5, /dev/sfdsk/Local/rdsk/Volume6
```

Note: For details on sfxadm, see the Global File Services Manual.

- Execute fsck.

When executing fsck, specify the Global File Services system's representative partition.

ii) Make a backup.

When making a backup, specify the Global File Services system's representative partition.

```
# /usr/sbin/sfxdump 0ucf /dev/rmt/0 /dev/sfdsk/Local/rdsk/Volume4
```

5) Removing the snapshot

Remove the snapshot you created, and cancel proxy volume (Volume4, Volume5 and Volume6) settings.

```
# /usr/sbin/sdxvolume -F -c Local -v Volume4,Volume5,Volume6  
# /usr/sbin/sdxproxy Break -c Local -p Volume4  
# /usr/sbin/sdxproxy Break -c Local -p Volume5  
# /usr/sbin/sdxproxy Break -c Local -p Volume6
```

When using more than one volume as a file system with Global File Services, you must execute the above command for all volumes.

Appendix E GDS Messages

This section explains the messages for GDS driver, daemon and command.

E.1 Web-Based Admin View Messages

For message numbers 0001 to 0099, see "Web-Based Admin View Operation Guide" for details.

E.2 Driver Messages

Driver messages are output to log file or console via syslog interface.

Variable Names

Italicized words in the messages are variable names, and the actual output will vary depending on the situation. The meaning and the format of the variable names used in the message explanation are described below.

<i>v_devno</i>	Volume device number (hexadecimal number)
<i>v_maj</i>	Volume major number (decimal number)
<i>v_min</i>	Volume minor number (decimal number)
<i>p_devno</i>	Physical slice device number (hexadecimal number)
<i>p_maj</i>	Physical slice major number (decimal number)
<i>p_min</i>	Physical slice minor number (decimal number)
<i>pdevtree</i>	Physical disk device tree path name. ex) "/pci@4,2000/pci@1/sunw,ispw@4/sd@2,0"
<i>pdevinst</i>	Physical disk instance name. ex) sd77
<i>blknodk</i>	I/O requested disk block number (decimal number)
<i>blknosl</i>	I/O requested slice block number (decimal number)
<i>length</i>	I/O transfer requested block size (in byte, decimal number)
<i>resid</i>	Residing block size (in byte, decimal number)
<i>errno</i>	System call error number (decimal number)
<i>oflag</i>	Open flag (hexadecimal number)
<i>details</i>	Details

Explanation

Messages output by the driver are shown below in the order of its severity.

There are three levels of severity.

PANIC	This message is displayed when an event that will stop the system is detected.
WARNING	This message is output when an abnormal event is detected. It will not always result in an error.
NOTICE	This message is output to record the driver operation and the like.

E.2.1 Panic Messages

sfdsk: no valid root disk found, disk ID information is invalid

Explanation

Disk ID information was invalid for all root disks (disks where root file system will be mounted).

Resolution

Check for improper root disk swap or connection change of I/O cables.

When the cause cannot be specified, collect investigation material including the memory dump and contact your local customer support.

E.2.2 Warning Messages

WARNING: sfdsk: read error on slice:

volume info : devno(maj,min)=v_devno(v_maj,v_min)

device info : devno(maj,min)=p_devno(p_maj,p_min)

devname=pdevtree (pdevinst)

error info : blk in disk=blknodk, blk in slice=blknosl

length=length, resid=resid, errno=errno

Explanation

Read request sent to slice terminated abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

WARNING: sfdsk: write error on slice:

volume info : devno(maj,min)=v_devno(v_maj,v_min)

device info : devno(maj,min)=p_devno(p_maj,p_min)

devname=pdevtree (pdevinst)

error info : blk in disk=blknodk, blk in slice=blknosl

length=length, resid=resid, errno=errno

Explanation

Write request sent to slice terminated abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk driver.

WARNING: sfdsk: read error on disk:

device info : devno(maj,min)=*p_devno(p_maj,p_min)*

devname=*pdevtree (pdevinst)*

error info : blk in disk=*blknodk*, blk in slice=*blknosl*

length=*length*, resid=*resid*, errno=*errno*

Explanation

Read request sent to disk terminated abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

WARNING: sfdsk: write error on disk:

device info : devno(maj,min)=*p_devno(p_maj,p_min)*

devname=*pdevtree (pdevinst)*

error info : blk in disk=*blknodk*, blk in slice=*blknosl*

length=*length*, resid=*resid*, errno=*errno*

Explanation

Write request sent to disk terminated abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

WARNING: sfdsk: read and writeback error on slice:

volume info : devno(maj,min)=*v_devno(v_maj,v_min)*

device info : devno(maj,min)=*p_devno(p_maj,p_min)*

devname=*pdevtree (pdevinst)*

error info : blk in disk=*blknodk*, blk in slice=*blknosl*

length=*length*, resid=*resid*, errno=*errno*

Explanation

Read request and writeback request sent to slice terminated abnormally. Writeback is a process to read data from other slices in the event of a read error.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

WARNING: sfdsk: open error on slice:

volume info : devno(maj,min)=v_devno(v_maj,v_min)

device info : devno(maj,min)=p_devno(p_maj,p_min)

devname=pdevtree (pdevinst)

error info : blk in disk=blknodk, blk in slice=blknosl

oflag=oflag, errno=errno

Explanation

Open request sent to slice returned abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

WARNING: sfdsk: open error on disk:

device info : devno(maj,min)=p_devno(p_maj,p_min)

devname=pdevtree (pdevinst)

error info : oflag=oflag, errno=errno

Explanation

Open request sent to disk returned abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

WARNING: sfdsk: close error on disk:

device info : devno(maj,min)=p_devno(p_maj,p_min)

devname=pdevtree (pdevinst)

error info : oflag=oflag, errno=errno

Explanation

Close request sent to disk returned abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

WARNING: sfdsk: NVURM read error on disk:

volume info : devno(maj,min)=v_devno(v_maj,v_min)

device info : devno(maj,min)=p_devno(p_maj,p_min)

devname=pdevtree (pdevinst)

error info : blk in disk=blknodk, blk in slice=blknosl

length=length, resid=resid, errno=errno

Explanation

NVURM read request sent to disk terminated abnormally.

NVURM is volume update area map information which is stored in the disk for just resynchronization.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk driver.

WARNING: sfdsk: NVURM write error on disk:

volume info : devno(maj,min)=v_devno(v_maj,v_min)

device info : devno(maj,min)=p_devno(p_maj,p_min)

devname=pdevtree (pdevinst)

error info : blk in disk=blknodk, blk in slice=blknosl

length=length, resid=resid, errno=errno

Explanation

NVURM write request sent to disk terminated abnormally.

NVURM is volume update area map information which is stored in the disk for just resynchronization.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk driver.

WARNING: sfdsk: volume status log write error on disk:

device info : devno(maj,min)=p_devno(p_maj,p_min)

devname=pdevtree (pdevinst)

error info : blk in disk=blknodk, blk in slice=blknosl

length=length, resid=resid, errno=errno

Explanation

Write request of volume status log sent to disk terminated abnormally.

Volume status log records if volume closed normally in the event of a system crash.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk driver.

sfdsk: driver started up

Explanation

The driver has been installed into the system.

sfdsk: received shutdown request

Explanation

A shutdown request from the sdxservd daemon has been received.

sfdsk: volume status log updated successfully,*details*

Explanation

Write request sent to the volume status log terminated normally. Volume status log records if the volume closed normally when a system failure occurred.

NOTICE: sfdsk: I/O error on slice:

volume info: devno(maj,min)=*v_devno(v_maj,v_min)*

**device info: devno(maj,min)=*p_devno(p_maj,p_min)*
 devname=*device***

**error info: blk in disk=*blknodk*, blk in slice=*blknosl*
 length=*length*, resid=*resid*, errno=*errno*
 flags=*b_flags***

Explanation

I/O request sent to slice terminated abnormally.

NOTICE: sfdsk: read error and writeback success on slice:

volume info : devno(maj,min)=*v_devno(v_maj,v_min)*

**device info : devno(maj,min)=*p_devno(p_maj,p_min)*
 devname=*pdevtree (pdevinst)***

**error info : blk in disk=*blknodk*, blk in slice=*blknosl*
 length=*length***

Explanation

Read request sent to slice terminated abnormally, but has been recovered by writeback process. Writeback is a process to read data from other slices in the event of a read error.

NOTICE: sfdsk: copy timeout. no response from sdxservd daemon:

volume info: devno(maj,min)=v_devno(v_maj,v_min)

Explanation

Synchronization copying was not performed since there was no response from the sdxservd daemon.

E.2.4 Internal Error Messages

severity: sfdsk: internal error, details

Explanation

Internal error occurred. *details* indicates the cause of the error. *severity* indicates the severity of the message.

Resolution

Collect investigation material and contact your local customer support.

E.3 Daemon Messages

Daemon messages are output to GDS log file.

`/var/opt/FJSVsdx/msglog/sdxserved.log`

Output format to log file is as follows.

Mon Day HH:MM:SS SDX:daemon: severity: message

Mon gives the month the message was output, *Day* the date, *HH* the hour, *MM* the minute, *SS* the second, *daemon* the daemon program name, *severity* the severity of the message, and *message* the message text body.

Depending on the settings of the syslog, daemon message will also be output to syslog log file and console.

Variable Names

Italicized words in the messages are variable names, and the actual output will vary depending on the situation. The meaning and the format of the variable names used in the message explanation is described below.

<i>class</i>	Class name
<i>disk</i>	Disk name
<i>group</i>	Group name
<i>lgroup</i>	Lower level group name
<i>hgroup</i>	Higher level group name
<i>volume</i>	Volume name
<i>disk.volume</i>	Slice name
<i>object.volume</i>	Slice name
<i>proxy</i>	Proxy object name
<i>object</i>	Object name
<i>status</i>	Object status
<i>device</i>	Physical disk name (cCtTdD, mphdI, mplbI,emcpowerN) where C is the controller number, T is the target ID, D is the disk number, I is the instance number and N is the emcpower device number.
<i>pslice</i>	Physical slice name(cCtTdDs{0-7}, mphd/s{0-7},mplb/s{0-7},emcpowerN{a-h}) where C is the controller number, T is the target ID, D is the disk number, I is the instance number, N is the emcpower device number, {0-7} is the slice number with 1 digit from 0 to 7, and {a-h} is the slice number with one alphabet from a to h. May be indicated by device tree file name. ex) "sd@2,0:a,raw"
<i>v_devno</i>	Volume device number (hexadecimal number)
<i>v_maj</i>	Volume major number (decimal number)
<i>v_min</i>	Volume minor number (decimal number)
<i>p_devno</i>	Physical slice device number (hexadecimal number)
<i>p_maj</i>	Physical slice major number (decimal number)
<i>p_min</i>	Physical slice minor number (decimal number)
<i>blknodk</i>	I/O requested disk block number (decimal number)
<i>blknosl</i>	I/O requested slice block number (decimal number)
<i>length</i>	I/O transfer requested block size (in byte, decimal number)

<i>resid</i>	Residing block size (in byte, decimal number)
<i>errno</i>	System call error number (decimal number)
<i>sdxerrno</i>	GDS defined internal error number (decimal number)
<i>node</i>	Node identifier, or node name
<i>oflag</i>	Open flag (hexadecimal number)
<i>ioctlcmd</i>	The ioctl command name
<i>timeofday</i>	Character string indicating current time and date
<i>details</i>	Details
<i>sdxfunc</i>	GDS function name
<i>exitstat</i>	Value indicating command end status (decimal number)
<i>cmdline</i>	Character indicating command line

Explanation

Messages output by daemon are shown below in the order of its severity. There are four levels of severity.

HALT	This message is output when an event that will halt all services provided by GDS is detected.
ERROR	This message is output when an event that will halt some services is detected.
WARNING	This message is output when an abnormal event is detected. It will not always result in a service halt.
INFO	This message is output to record the daemon operation and the like. Usually, it can be ignored.

E.3.1 Halt Messages

HALT: failed to create a new thread, *errno=errno*

Explanation

Function `pthread_create()` terminated abnormally. Process cannot be continued. It will exit the daemon process.

Resolution

When error number information *errno* is insufficient in identifying the cause and to recover, collect investigation material and contact your local customer support.

HALT: cannot open driver administrative file, *errno=errno*

Explanation

GDS driver(`sfds`) administrative file cannot be opened. Process cannot be continued. It will exit the daemon process.

This message is output when files under `/dev/sfdsk` directory cannot be accessed.

Resolution

Collect investigation material, and contact your local customer support.

HALT: startup failed, no enough address space

Explanation

Startup failure due to unsuccessful memory allocation. Process cannot be continued. It will exit the daemon process.

Resolution

Confirm you have sufficient memory or swap area.

HALT: failed to respawn *daemon* daemon, *osfunc=osfunc*, *errno=errno*

Explanation

The *daemon* terminated abnormally and failed to restart. The failure was caused by abnormal termination of Solaris(TM) Operating Environment *osfunc* function.

The error number is *errno*. This message is output via *syslog*.

Resolution

When error number information is insufficient in identifying the cause, collect investigation material, and contact your local customer support.

HALT: cannot start node-down recovery for remote node *node*, no enough space, *osfunc=osfunc*, *errno=errno*

Explanation

Could not recover the crashed remote node *node* due to unsuccessful memory allocation. Process cannot be continued. It will exit the daemon process.

Resolution

Solaris(TM) Operating Environment function *osfunc* terminated abnormally, and error number is *errno*. Confirm you have sufficient memory or swap area, and recover.

E.3.2 Error Messages

ERROR: read error on *status slice object.volume*, *class=class*:

volume info:devno(maj,min)=*v_devno(v_maj,v_min)*

device info:devno(maj,min)=*p_devno(p_maj,p_min)*

devname=*device*

error info: blk in disk=*blknodk*, blk in slice=*blknosl*,

length=*length*, resid=*resid*, *errno=errno*

Explanation

Read request sent to slice *object.volume* in *status* status terminated abnormally. Read request sent to volume configured by this slice, or to slice accessible in isolation was returned to error. You must attempt recovery as soon as possible since application may stop normal operation.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and promptly recover the disk.

ERROR: write error on *status slice object.volume*, class=*class*:

volume info:devno(maj,min)=*v_devno(v_maj,v_min)*

device info:devno(maj,min)=*p_devno(p_maj,p_min)*

devname=*device*

error info:blk in disk=*blknodk*, blk in slice=*blknosl*,

length=*length*, resid=*resid*, errno=*errno*

Explanation

Write request sent to slice *object.volume* in *status* status terminated abnormally. Write request sent to volume configured by this slice, or to slice accessible in isolation was returned to error.

You must attempt recovery as soon as possible since application may stop normal operation.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and promptly recover the disk.

ERROR: open error on *status slice object.volume*, class=*class*:

volume info:devno(maj,min)=*v_devno(v_maj,v_min)*

device info:devno(maj,min)=*p_devno(p_maj,p_min)*

devname=*device*

error info:oflag=*oflag*, errno=*errno*

Explanation

Open request sent to slice *object.volume* in *status* status terminated abnormally. Open request sent to volume configured by this slice, or to slice accessible in isolation was returned to error.

You must attempt recovery as soon as possible since application may stop to operate normally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and promptly recover the disk.

ERROR: read and writeback error on *status slice object.volume*, class=*class*:

volume info:devno(maj,min)=*v_devno(v_maj,v_min)*

device info:devno(maj,min)=*p_devno(p_maj,p_min)*

devname=*device*

error info:blk in disk=*blknodk*, blk in slice=*blknosl*,

length=*length*, resid=*resid*, errno=*errno*

Explanation

Read request and writeback request sent to slice *object.volume* in *status* status terminated abnormally. Read request sent to volume configured by this slice, or to slice accessible in isolation was returned to error.

You must attempt recovery as soon as possible since application may not operate normally.

Writeback is a process to read data from other slices in the event of a read error.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and promptly recover the disk.

ERROR: *volume*: closed down volume, class=*class*

Explanation

volume is closed down.

Resolution

Promptly attempt recovery by identifying the cause of failure by referring to GDS log message that is output immediately before the error.

For information on recovery, see "F.1.3 Volume Status Abnormality."

ERROR: *class*: cannot startup *class*, no valid configuration database, sdxerno=*errno*

Explanation

Could not startup *class*, since valid *class* configuration database could not be found.

This message is output when all disks (or the majority of disks) registered with *class* is unavailable.

Resolution

See "F.1.4 Class Status Abnormality".

ERROR: *class*: cannot startup *class*, too few valid configuration database replicas, sdxerno=*errno*

Explanation

Could not startup *class*, due to insufficient number of valid database. This message is output when many disks registered with *class* is unavailable.

Resolution

See "F.1.4 Class Status Abnormality".

ERROR: *class*: closing down *class*, no valid configuration database

Explanation

class was closed since no valid *class* configuration database was found.

This message is output when all disks (or the majority of disks) registered with *class* is unavailable.

Resolution

See "F.1.4 Class Status Abnormality".

ERROR: *class*: closing down class, too few valid configuration database replicas

Explanation

class was closed due to insufficient number of valid database.
This message is output when many disks registered with *class* is unavailable.

Resolution

See "F.1.4 Class Status Abnormality".

ERROR: failed to send request message on node *node*, *details*

Explanation

Sending request message from *node* was unsuccessful.

Resolution

Collect investigation material and contact your local customer support.

ERROR: timeout on receiving reply message from node *node*, *details*

Explanation

Timeout occurred while receiving a reply message from a remote node *node*.

Resolution

Collect investigation material and contact your local customer support.

ERROR: rejected request message on remote node *node*, *details*

Explanation

Processing a request message on a remote node *node* was unsuccessful.

Resolution

Investigate the *node* message log and take necessary actions. If recovery is impossible, collect investigation material and contact your local customer support.

ERROR: *class*: failed to start shared volumes, *status* volume *volume* exists, node=*node*

Explanation

Starting a shared volume within *class* was unsuccessful on node due to *volume* in *status* status.

Resolution

volume is in abnormal status. First, you must recover normal status. For information on recovery, see "F.1.3 Volume Status Abnormality."

ERROR: *class*: failed to start and standby shared volumes, *status* volume *volume* exists, node=*node*

Explanation

Starting and performing a standby process on a shared volume within *class* was unsuccessful on *node* due to *volume* in *status* status.

Resolution

volume is in abnormal status. First, you must recover normal status. For information on recovery,

see "F.1.3 Volume Status Abnormality".

ERROR: *class*: failed to stop and standby shared volumes, *status* volume *volume* exists, *node*=*node*

Explanation

Stopping a shared volume within class was unsuccessful on *node* due to *volume* in *status* status.

Resolution

volume is in abnormal status. First, you must recover normal status. For information on recovery, see "F.1.3 VolumeStatus Abnormality".

ERROR: *class*: failed to stop shared volumes, *status* volume *volume* exists, *node*=*node*

Explanation

Starting a shared volume was unsuccessful due to *class* in closed status.

Resolution

Recover closed *class*. There may be multiple disk failures. Identify the cause based on object status, GDS log message, and syslog message.

For information on recovery, see "F.1.4 Class Status Abnormality."

ERROR: *class*: failed to start shared volumes, *class* closed down, *node*=*node*

Explanation

Starting a shared volume and performing a standby process on a shared volume was unsuccessful due to *class* in closed status.

Resolution

Recover closed *class*. There may be multiple disk failures. Identify the cause based on object status, GDS log message, and syslog message.

For information on recovery, see "F.1.4 Class Status Abnormality."

ERROR: *class*: failed to start and standby shared volumes, *class* closed down, *node*=*node*

Explanation

Starting a shared volume and performing a standby process was unsuccessful on a shared volume due to *class* in closed status.

Resolution

Recover closed *class*. There may be multiple disk failures. Identify the cause based on object status, GDS log message, and syslog message.

For information on recovery, see "F.1.4 Class Status Abnormality".

ERROR: *class*: failed to stop and standby shared volumes, *class* closed down, *node*=*node*

Explanation

Stopping a shared volume was unsuccessful due to *class* in closed status.

Resolution

Recover closed *class*. There may be multiple disk failures. Identify the cause based on object status, GDS log message, and syslog message.

For information on recovery, see "F.1.4 Class Status Abnormality".

ERROR: *class*: failed to stop shared volumes, class closed down, node=*node***Explanation**

Stopping a shared volume was unsuccessful due to *class* in closed status.

Resolution

Recover closed *class*. There may be multiple disk failures. Identify the cause based on object status, GDS log message, and syslog message.

For information on recovery, see "F.1.4 Class Status Abnormality".

ERROR: *class*: closing down class, cluster-wide lock failure, *sdxfcrmo*=*sdxfcrmo***Explanation**

Abnormal exclusive control between cluster system nodes occurred.

Since the process cannot be continued, *class* will be closed.

Resolution

Collect investigation material and contact your local customer support.

ERROR: *class*: cannot startup class, cluster-wide lock failure, *sdxfcrmo*=*ermo***Explanation**

Abnormal exclusive control between cluster system nodes occurred.

Since the process cannot be continued, *class* could not be started.

Resolution

Collect investigation material and contact your local customer support.

ERROR: *class*: closing down class, cluster communication failure, *sdxfcrmo*=*sdxfcrmo***Explanation**

Transmission failure between cluster system nodes occurred. Since the process cannot be continued, *class* will be closed.

Resolution

Collect investigation material and contact your local customer support.

ERROR: *class*: cannot operate in cluster environment, created when cluster control facility not ready**Explanation**

The class *class* cannot be used in the cluster environment since it was created when the cluster control was disabled. One of the following operation mistakes can be the possible reason.

- a) After *class* was created cluster initialization was performed on a node where the cluster has not been initialized.

- b) In a cluster environment where cluster initialization is complete, *class* was created in the single-user mode.

Resolution

Boot the node in the single-user mode, create a backup copy of data on a volume according to need and delete *class*. Then, reboot the node in the multi-user mode, and re-create *class* and a volume as appropriate for the volume data to be restored.

ERROR: *proxy*: failed to copy with OPC, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

While performing copying between proxy volume *proxy* and master volume with the OPC function, an I/O error occurred, and the copying process failed. *source* specifies the original slice name, *target* the mirror slice name, and *class* to which the original slice and mirror slice belong.

Resolution

Identify the cause by referring to the log messages for GDS, disk driver, GR700 series and ETERNUS3000 that was output right before the occurrence of the error, and restore.

ERROR: *proxy*: failed to copy with EC, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

While performing copying between proxy volume *proxy* and master volume with the EC function, an I/O error occurred, and the copying process failed. *source* specifies the original slice name, *target* the mirror slice name, and *class* to which the original slice and mirror slice belong.

Resolution

Identify the cause by referring to the log messages for GDS, disk driver, and GR disk unit that was output right before the occurrence of the error, and restore.

ERROR: *proxy*: failed to copy with TimeFinder, source=*disk*, target=*disk*, class=*class*

Explanation

Copying failed due to an I/O error caused while conducting copying between the proxy group *proxy* and the master group with TimeFinder. "source" is the name of the copy source disk, "target" is the name of the copy destination disk, and "class" is the name of the class to which the copy source and destination disks belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or Symmetrix device, and promptly recover the disk.

ERROR: *proxy*: failed to copy with SRDF, source=*disk*, target=*disk*, class=*class*

Explanation

Copying failed due to an I/O error caused while conducting copying between the proxy group *proxy* and the master group with SRDF. "source" is the name of the copy source disk, "target" is the name of the copy destination disk, and "class" is the name of the class to which the copy source and destination disks belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or

Symmetrix device, and promptly recover the disk.

ERROR: *proxy*: failed to start OPC, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

When starting OPC between proxy volume *proxy* and master volume, an error occurred. *source* specifies the original slice name, *target* the mirror slice name, and *class* to which the original slice and mirror slice belong.

Resolution

Identify the cause by referring to the log messages for GDS, disk driver, GR700 series and ETERNUS3000 that was output right before the occurrence of the error, and restore.

ERROR: *proxy*: failed to stop OPC, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

When stopping OPC between proxy volume *proxy* and master volume, an error occurred. *source* specifies the original slice name, *target* the mirror slice name, and *class* to which the original slice and mirror slice belong.

Resolution

Identify the cause by referring to the log messages for GDS, disk driver, GR700 series and ETERNUS3000 that was output right before the occurrence of the error, and restore.

ERROR: *proxy*: failed to start EC session, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

When starting EC session between proxy volume *proxy* and master volume, an error occurred. *source* specifies the original slice name, *target* the mirror slice name, and *class* to which the original slice and mirror slice belong.

Resolution

Identify the cause by referring to the log messages for GDS, disk driver, and GR disk unit that was output right before the occurrence of the error, and restore.

ERROR: *proxy*: failed to stop EC session, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

When stopping EC session between proxy volume *proxy* and master volume, an error occurred. *source* specifies the original slice name, *target* the mirror slice name, and *class* to which the original slice and mirror slice belong.

Resolution

Identify the cause by referring to the log messages for GDS, disk driver, and GR disk unit that was output right before the occurrence of the error, and restore.

ERROR: *proxy*: failed to suspend EC session, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

When temporarily suspending EC session between proxy volume *proxy* and master volume, an error occurred. *source* specifies the original slice name, *target* the mirror slice name, and *class* to which the original slice and mirror slice belong.

Resolution

Identify the cause by referring to the log messages for GDS, disk driver, and GR disk unit that was output right before the occurrence of the error, and restore.

ERROR: *proxy*: failed to resume EC session, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

When resuming EC session between proxy volume *proxy* and master volume, an error occurred. *source* specifies the original slice name, *target* the mirror slice name, and *class* to which the original slice and mirror slice belong.

Resolution

Identify the cause by referring to the log messages for GDS, disk driver, and GR disk unit that was output right before the occurrence of the error, and restore.

ERROR: *proxy*: failed to establish BCV pair, STD=*disk*, BCV=*disk*, class=*class*

Explanation

Error occurred while establishing a BCV pair between the proxy group *proxy* and the master group. "STD" is the disk name of the standard device, "BCV" is the disk name of the BCV device, and "class" is the name of the class to which the standard and BCV devices belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or Symmetrix device, and promptly recover the disk.

ERROR: *proxy*: failed to cancel BCV pair, STD=*disk*, BCV=*disk*, class=*class*

Explanation

Error occurred while canceling a BCV pair between the proxy group *proxy* and the master group. "STD" is the disk name of the standard device, "BCV" is the disk name of the BCV device, and "class" is the name of the class to which the standard and BCV devices belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or Symmetrix device, and promptly recover the disk.

ERROR: *proxy*: failed to split BCV pair, STD=*disk*, BCV=*disk*, class=*class*

Explanation

Error occurred while splitting a BCV pair between the proxy group *proxy* and the master group. "STD" is the disk name of the standard device, "BCV" is the disk name of the BCV device, and "class" is the name of the class to which the standard and BCV devices belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or Symmetrix device, and promptly recover the disk.

ERROR: *proxy*: failed to re-establish BCV pair, STD=*disk*, BCV=*disk*, class=*class*

Explanation

Error occurred while re-establishing a BCV pair between the proxy group *proxy* and the master group. "STD" is the disk name of the standard device, "BCV" is the disk name of the BCV device, and "class" is the name of the class to which the standard and BCV devices belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or Symmetrix device, and promptly recover the disk.

ERROR: *proxy*: failed to establish SRDF pair, STD=*disk*, BCV=*disk*, class=*class*

Explanation

Error occurred while establishing an SRDF pair between the proxy group *proxy* and the master group. "source" is the name of the source disk, "target" is the name of the target disk, and "class" is the name of the class to which the source and target disks belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or Symmetrix device, and promptly recover the disk.

ERROR: *proxy*: failed to cancel SRDF pair, STD=*disk*, BCV=*disk*, class=*class*

Explanation

Error occurred while canceling an SRDF pair between the proxy group *proxy* and the master group. "source" is the name of the source disk, "target" is the name of the target disk, and "class" is the name of the class to which the source and target disks belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or Symmetrix device, and promptly recover the disk.

ERROR: *proxy*: failed to split SRDF pair, STD=*disk*, BCV=*disk*, class=*class*

Explanation

Error occurred while splitting an SRDF pair between the proxy group *proxy* and the master group. "source" is the name of the source disk, "target" is the name of the target disk, and "class" is the name of the class to which the source and target disks belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or Symmetrix device, and promptly recover the disk.

ERROR: *proxy*: failed to re-establish SRDF pair, STD=*disk*, BCV=*disk*, class=*class*

Explanation

Error occurred while re-establishing an SRDF pair between the proxy group *proxy* and the master group. "source" is the name of the source disk, "target" is the name of the target disk, and "class" is the name of the class to which the source and target disks belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or Symmetrix device, and promptly recover the disk.

E.3.3 Warning Messages

WARNING: read error on *status slice object.volume*, class=*class*:

volume info:devno(maj,min)=*v_devno(v_maj,v_min)*

device info:devno(maj,min)=*p_devno(p_maj,p_min)*

devname=*device*

error info:blk in disk=*blknodk*, blk in slice=*blknosl*,

length=*length*, resid=*resid*, errno=*errno*

Explanation

Read request sent to slice *object.volume* in *status* status terminated abnormally. Slice with abnormality will be detached.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

WARNING: write error on *status slice object.volume*, class=*class*:

volume info:devno(maj,min)=*v_devno(v_maj,v_min)*

device info:devno(maj,min)=*p_devno(p_maj,p_min)*

devname=*device*

error info:blk in disk=*blknodk*, blk in slice=*blknosl*,

length=*length*, resid=*resid*, errno=*errno*

Explanation

Write request sent to slice *object.volume* in *status* status terminated abnormally. Slice with abnormality will be detached.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

WARNING: open error on *status* slice *object.volume*, class=*class*:

volume info:devno(maj,min)=*v_devno(v_maj,v_min)*

device info:devno(maj,min)=*p_devno(p_maj,p_min)*

devname=*device*

error info:oflag=*oflag*, errno=*errno*

Explanation

Open request sent to slice *object.volume* in *status* status terminated abnormally. Slice with abnormality will be detached.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

WARNING: read and writeback error on *status* slice *object.volume*, class=*class*:

volume info:devno(maj,min)=*v_devno(v_maj,v_min)*

device info:devno(maj,min)=*p_devno(p_maj,p_min)*

devname=*device*

error info:blk in disk=*blknodk*, blk in slice=*blknosl*,

length=*length*, resid=*resid*, errno=*errno*

Explanation

Read request and writeback request sent to slice *object.volume* in *status* status terminated abnormally. Slice with abnormality will be detached.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk driver.

WARNING: NVURM write error on disk *disk*, class=*class*:

volume info:devno(maj,min)=*v_devno(v_maj,v_min)*

device info:devno(maj,min)=*p_devno(p_maj,p_min)*

devname=*device*

error info:blk in disk=*blknodk*, blk in slice=*blknosl*,

length=*length*, resid=*resid*, errno=*errno*

Explanation

NVURM write request sent to disk *disk* terminated abnormally.

Although just resynchronization process on *volume* will be temporarily disabled, it will automatically attempt recovery.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

WARNING: *object.volume*: detached status slice by an I/O error, class=*class*

Explanation

Since an I/O error occurred on slice *object.volume* in *status* status, the slice was detached from the volume.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk driver.

WARNING: open error on private slice *pslice*, oflag=*oflag*, ermo=*ermo*

Explanation

Open request sent to disk private slice *pslice* terminated abnormally. It will automatically search for a normal alternate disk and attempt recovery.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

WARNING: read error on private slice *p_devno(p_maj,p_min)*, offset=*blkno*, length=*length*, resid=*resid*, ermo=*ermo*

Explanation

Read request sent to disk private slice *p_devno(p_maj,p_min)* terminated abnormally. It will automatically search for a normal alternate disk and attempt recovery.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk driver.

WARNING: write error on private slice *p_devno(p_maj,p_min)*, offset=*blkno*, length=*length*, resid=*resid*, ermo=*ermo*

Explanation

Write request sent to disk private slice *p_devno(p_maj,p_min)* terminated abnormally. It will automatically search for a normal alternate disk and attempt recovery.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

WARNING: close error on private slice *p_devno(p_maj,p_min)*, ermo=*ermo*

Explanation

Close request sent to disk private slice *p_devno(p_maj,p_min)* terminated abnormally. It will automatically search for a normal alternate disk and attempt recovery.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk driver.

WARNING: *sdxfunc: pslice: open error, ermo=ermo*

Explanation

Open request sent to physical slice *pslice* terminated abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk driver.

WARNING: *sdxfunc: pslice: read error, ermo=ermo*

Explanation

Read request sent to physical slice *pslice* terminated abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk driver.

WARNING: *sdxfunc: pslice: write error, ermo=ermo*

Explanation

Write request sent to physical slice *pslice* terminated abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

WARNING: *sdxfunc: pslice: ioctl error, request=ioctlcmd, ermo=ermo*

Explanation

The ioctl request sent to physical slice *pslice* terminated abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

WARNING: *sdxfunc: pslice: close error, ermo=ermo*

Explanation

Close request sent to physical slice *pslice* terminated abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

WARNING: *volume*: failed to enable JRM, no available NVURM space, class=*class*

Explanation

Resuming just resynchronization process on *volume* was unsuccessful due to insufficient number of normal disks.

Resolution

Check the disk status within *class*. Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

WARNING: *volume*: failed to retrieve NVURM from disk *disk*, class=*class*

Explanation

NVURM read request of *volume* from *disk* was unsuccessful. Just resynchronization will switch to copying of the entire disk.

NVURM is a volume update area map information stored in the disk for just resynchronization.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk driver.

WARNING: *disk: device*: disabled disk, class=*class*

Explanation

disk is disabled since its data is invalid. *device* is the physical disk name corresponding to *disk*.

Resolution

Disk configuration may have been wrongly changed or disk data could be damaged. Check for improper connection change of I/O cables and disk swap.

Also check for disk driver log message regarding the disk in question, and attempt recovery.

WARNING: *volume*: volume synchronization failed, class=*class*

Explanation

Synchronization copying of *volume* was unsuccessful.

Resolution

Attempt recovery by identifying the cause of failure by referring to GDS log message and disk driver message that was output immediately before the error.

WARNING: *volume*: volume just resynchronization failed, class=*class*

Explanation

Just resynchronization copying of *volume* was unsuccessful.

Resolution

Attempt recovery by identifying the cause of failure by referring to GDS log message and disk driver message that was output immediately before the error.

WARNING: class: unknown class file found

Explanation

Class file *class* which does not exist in class database file was found while booting the system.

**WARNING: invalid configuration database ID information, sdxermo=*sdxermo*,
class=*class*:**

/devices/sbus@1f,0/QLGC,isp@0,10000/sd@1,0:a,raw

Explanation

class configuration database ID information was invalid, and was not used as configuration database.

Resolution

Disk configuration may have been wrongly changed or disk data could be damaged. Check for improper connection change of I/O cables and disk swap. Also check for disk driver log message regarding the disk in question, and attempt recovery.

WARNING: class: too few valid configuration database replicas

Explanation

There is not sufficient valid configuration database for *class*.
This message is displayed when majority of disks registered with *class* is unavailable.
Leaving it as is may cause serious problems.

Resolution

For details, see "F.1.4 Class Status Abnormality."

WARNING: cannot open message logging file, ermo=*ermo*

/var/opt/FJSVsdx/msglog/daemon.log

Explanation

Opening GDS log file was unsuccessful. This message will be output via syslog. Although the message will not be output on GDS logfile, it does not affect other processes.

Resolution

Collect investigation material, and contact your local customer support regarding recovery.

WARNING: cannot write message logging file, ermo=*ermo*

/var/opt/FJSVsdx/msglog/sdxservd.log

Explanation

Writing to GDS log file was unsuccessful. This message will be output via syslog. Although the message will not be output on GDS logfile, it does not affect other processes.

Resolution

Collect investigation material and contact your local customer support.

WARNING: failed to reply message to node *node*, details

Explanation

Replying to remote node *node* was unsuccessful.

Resolution

Collect investigation material and contact your local customer support.

WARNING: class: failed to change class resource status on remote node *node*, status=*new-status*, sdxerno=*sdxerno*

Explanation

Changing class resource status to *new-status* on a remote node *node* was unsuccessful.

Resolution

Collect investigation material and contact your local customer support.

WARNING: sdxfunc: no enough address space, osfunc=*osfunc*, ermo=*ermo*

Explanation

Function *osfunc* provided by Solaris(TM) Operating Environment resulted in an error. Error number is *ermo*.

Resolution

Confirm you have sufficient memory or swap area.

WARNING: respawned *daemon* daemon successfully

Explanation

Although *daemon* terminated abnormally, it completed normally by restarting. This message will be output via syslog.

Resolution

To identify the cause of *daemon* terminating abnormally, collect investigation material and contact your local customer support.

WARNING: device: failed to restore VTOC on disk, sdxerno=*sdxerno*

Explanation

Recovering physical disk *device* format information was unsuccessful. Disk failure may have occurred.

Resolution

Use format command to recover format information.

WARNING: cannot set use-nvramrc? parameter to true, NVRAMRC not empty

Explanation

Since a command is already registered with NVRAMRC, use-nvramrc? parameter could not be set to "true."

Resolution

For details, see "A.2.9 System Disk Mirroring."

WARNING: *volume*: cannot copy to one or more slices in volume

Explanation

Copying process could not be performed on some slices within *volume*.

Resolution

Execute synchronization copying as needed.

E.3.4 Information Messages

INFO: read error and writeback success on *status* slice *object.volume*, class=*class*:

volume info:devno(maj,min)=*v_devno(v_maj,v_min)*

device info:devno(maj,min)=*p_devno(p_maj,p_min)*

devname=*device*,

error info:blk in disk=*blknodk*, blk in slice=*blknosl*,

length=*length*

Explanation

Read request sent to slice *object.volume* in status *status* was unsuccessful, but was restored by writeback process.

Writeback is a process to read data from other slices in the event of a read error.

INFO: NVRAM configuration parameter has been updated:

parameter=*'param'*

old =*'old_value'*

new =*'new_value'*

Explanation

Parameter *param* value stored in NVRAM (non-volatile memory) on the body unit was updated from *old_value* to *new_value*.

INFO:*volume*: temporarily disabled JRM, class=*class*

Explanation

Although just resynchronization process on *volume* will be temporarily disabled, it will automatically attempt recovery.

**INFO: *disk*: failed to connect spare disk for disk *disk*, group=*group*, class=*class*,
sdxerno=*sdxerno***

Explanation

Connecting spare disk *disk* to group instead of *disk* was unsuccessful.

**INFO: *disk*: failed to connect spare disk for group *lgroup*, the highest level
group=*hgroup*, class=*class*, sdxerno=*sdxerno***

Explanation

The attempt to connect spare disk *disk* to the highest level group *hgroup* in place of group *lgroup* failed.

INFO: *disk*: connected spare disk for disk *disk*, group=*group*, class=*class*

Explanation

Connected spare disk *disk* to group instead of *disk*.

**INFO: *disk*: connected spare disk for group *lgroup*, the highest level group=*hgroup*,
class=*class***

Explanation

Spare disk *disk* was connected to the highest level group *hgroup* in place of group *lgroup*.

INFO: *group*: free blocks are reduced in group, class=*class*

Explanation

Free blocks of *group* was reduced.

INFO: *volume*: reallocated NVURM space and enabled JRM successfully, class=*class*

Explanation

Just resynchronization process of *volume* was resumed.

INFO: *volume*: retrieved NVURM from disk *disk* successfully, class=*class*

Explanation

NVURM read request of *volume* sent from *disk* was successful. Just resynchronization process will be resumed.

NVURM is a volume update area map information stored in the disk for just resynchronization.

INFO: *volume*: no need to retrieve NVURM, *sdxerno*=*sdxerno*, *class*=*class*

Explanation

NVURM was not retrieved. Entire copying of volume will start.

INFO: *disk: pslice*: failed to open physical special files exclusively, *erno*=*erno*

Explanation

Opening *disk* physical slice *pslice* exclusively was unsuccessful.

INFO: *disk: device*: disk ID information is invalid, *sdxerno*=*sdxerno*

Explanation

Disk ID information of *disk* is invalid. *disk* will be automatically disabled.

INFO: *disk*: enabled disk, *class*=*class*

Explanation

Disabled *disk* was enabled.

INFO: *volume*: volume synchronization started, *class*=*class*

Explanation

Started synchronization copying process of *volume*.

INFO: *volume*: volume just resynchronization started, *class*=*class*

Explanation

Started just resynchronization process of *volume*.

INFO: *volume*: volume synchronization canceled, *class*=*class*

Explanation

Canceled synchronization copying of *volume*.

INFO: *volume*: volume just resynchronization canceled, *class*=*class*

Explanation

Canceled just resynchronization process of *volume*.

INFO: *volume*: volume synchronization completed successfully, class=*class*

Explanation

Synchronization copying process of *volume* completed successfully.

INFO: *volume*: volume just resynchronization completed successfully, class=*class*

Explanation

Just resynchronization process of *volume* completed successfully.

INFO: *object*: failed to update configuration database, class=*class*

Explanation

Updating configuration database was unsuccessful, since all (or majority of) disks registered with *class* was unavailable. Usually, class will be closed following this message.

INFO: sdxservd daemon started up

Explanation

The sdxservd daemon started up. GDS will be started.

INFO: started local volumes, *timeofday*

Explanation

All volumes under local class has been started.

INFO: started root volumes

Explanation

All volumes under root class has been started.

INFO: started local volumes, *timeofday*

Explanation

All local volumes were started.

INFO: stopped all services by shutdown request, *timeofday*

Explanation

All GDS services were stopped in response to a shutdown request.

INFO: cannot open class database file, ermo=*ermo*

Explanation

Could not open class database file while booting the system. Will automatically attempt recovery.

INFO: class database file corrupted

Explanation

Corrupted class database file was detected while booting the system.
Will automatically attempt recovery.

INFO: *class*: cannot open class file, ermo=*ermo*

Explanation

Could not open class file *class* while booting the system. Will automatically attempt recovery.

INFO: *class*: class file corrupted

Explanation

Corrupted class file *class* was detected while booting the system. Will automatically attempt recovery.

INFO: class database file updated successfully

Explanation

Class database file was updated.

INFO: *class*: class file updated successfully

Explanation

Class file *class* was updated.

INFO: cannot write class database file, ermo=*ermo*

Explanation

Could not write to class database file. Will automatically attempt recovery.

INFO: *class*: cannot write class file, ermo=*ermo*

Explanation

Could not write to class file *class*. Will automatically attempt recovery.

INFO: cannot check configuration database ID information, sdxerno=*sdxerno*, class=*class*:

/devices/sbus@1f,0/QLGC,isp@0,10000/sd@1,0:a,raw

Explanation

Opening or reading private slice was unsuccessful, and could not check configuration database ID information of *class*.

INFO: cannot check configuration database, sdxerno=*sdxerno*, class=*class*:

/devices/sbus@1f,0/QLGC,isp@0,10000/sd@1,0:a,raw

Explanation

Opening or reading private slice was unsuccessful, and could not check configuration database of *class*.

INFO: configuration database corrupted, sdxerno=*sdxerno*, class=*class*:

/devices/sbus@1f,0/QLGC,isp@0,10000/sd@1,0:a,raw

...

Explanation

Did not use the configuration database, since checksum of *class* configuration database was invalid.

INFO: configuration database defeated, sdxerno=*sdxerno*, class=*class*:

/devices/sbus@1f,0/QLGC,isp@0,10000/sd@1,0:a,raw

...

Explanation

Configuration database of *class* was regarded invalid.

INFO: *class*: valid configuration database replicas exist on:

/devices/sbus@1f,0/QLGC,isp@0,10000/sd@1,0:a,raw

...

Explanation

Valid *class* configuration database was determined.

INFO: *class*: starting up class

Explanation

class will be started.

INFO: cannot update configuration database replica, *sdxermo*=*sdxermo*, *class*=*class*:

/devices/sbus@1f,0/QLGC,isp@0,10000/sd@1,0:a,raw

Explanation

Could not update *class* configuration database replica.

INFO: *class*: relocated configuration database replicas on:

/devices/sbus@1f,0/QLGC,isp@0,10000/sd@1,0:a,raw

...

Explanation

Relocated *class* configuration database replica.

INFO: *disk*: disconnected spare disk from group *group*, *class*=*class*

Explanation

Disconnected spare disk *disk* from *group*.

INFO: *group*: free blocks are increased in group, *class*=*class*

Explanation

Number of free blocks in *group* increased.

INFO: failed to create a new thread, *ermo*=*ermo*

Explanation

Function `pthread_create()` terminated abnormally.

INFO: cannot open configuration parameter file *filename*, *ermo*=*ermo*

Explanation

Opening configuration parameter *filename* was unsuccessful.

INFO: cannot read configuration parameter file, errno=*errno*

Explanation

Reading configuration parameter files was unsuccessful.

INFO: received unexpected data from sfdsk driver and ignored

Explanation

Unexpected data was received from sfdsk driver, and has been ignored.

INFO: received unexpected event from sfdsk driver and ignored, *details*

Explanation

Unexpected event was received from sfdsk driver, and has been ignored.
details displays the details of the event.

INFO: *class*: class closed down, node=*node*

Explanation

class on *node* was closed.

INFO: command executed:

cmdline

Explanation

Command *cmdline* was executed.

INFO: command exited, exit-status=*exitstat*:

cmdline

Explanation

Processing *cmdline* is completed.

**INFO: *class*: changed class resource status on remote node *node*, *old-status* =>
*new-status***

Explanation

Class resource status on remote node *node* was changed from *old-status* to *new-status*.

INFO: *class*: changed class resource status on current node *node*, *old-status* => *new-status*

Explanation

Class resource status on current node *node* was changed from *old-status* to *new-status*.

INFO: *class*: started shared volumes, node=*node*

Explanation

Starting process of all the shared volumes associated with *class* on *node* is complete.

INFO: *class*: started and stood by shared volumes, node=*node*

Explanation

Starting and performing a standby process of shared volume associated with *class* on *node* is complete.

INFO: *class*: stopped and stood by shared volumes, node=*node*

Explanation

Stopping and standby process of shared volume associated with *class* on *node* is complete.

INFO: *class*: stopped shared volumes, node=*node*

Explanation

Stopping process of shared volume associated with *class* on *node* is complete.

INFO: cannot connect spare disk, cluster-wide lock failure, class=*class*, sdxerno=*sdxerno*

Explanation

Due to the occurrence of an abnormal exclusive control between cluster system nodes, the spare disk could not be connected.

INFO: cannot connect spare disk, too few valid configuration database replicas, class=*class*, disk=*disk*

Explanation

Could not connect spare disk due to insufficient number of valid configuration databases.

INFO: cannot connect spare disk, hot spare disabled, class=*class*, disk=*disk*

Explanation

Spare disk *disk* could not be connected since the hot spare is disabled.

INFO: *class*: started class-down recovery for remote node *node*

Explanation

Closed class *class* on remote node *node* will be recovered.

INFO: *class*: class-down recovery failed, already class-down on current node *node*

Explanation

Attempted recovery of closed class *class* on remote node *node*. Recovery was unsuccessful since the class was also closed on current node *node*.

INFO: *class*: class-down recovery failed, sdxerno=*sdxerno*

Explanation

Recovering closed class *class* was unsuccessful.

INFO: *class*: class-down recovery completed successfully

Explanation

Recovering closed class *class* was successful.

INFO: *class*: started node-down recovery for remote node *node*

Explanation

Started node-down recovery on remote node *node*.

INFO: *class*: started shutdown recovery for remote node *node*

Explanation

Started shutdown recovery on remote node *node*.

INFO: *class*: node-down recovery failed, already class-down on current node *node*

Explanation

Recovering node-down was unsuccessful since *class* was in closed status on current node.

INFO: *class*: shutdown recovery failed, already class-down on current node *node*

Explanation

Recovering shutdown was unsuccessful since *class* was in closed status on current node.

INFO: *class*: node-down recovery failed, *sdxermo*=*sdxermo*

Explanation

Recovering *class* from node-down was unsuccessful.

INFO: *class*: shutdown recovery failed, *sdxermo*=*sdxermo*

Explanation

Recovering *class* from shutdown was unsuccessful.

INFO: *class*: node-down recovery completed successfully

Explanation

Recovering *class* from node-down completed successfully.

INFO: *class*: shutdown recovery completed successfully

Explanation

Recovering *class* from shutdown completed successfully.

INFO: *object.volume*: failed to update slice error information, class closed down, class=*class*

Explanation

Updating error information of slice *object.volume* was unsuccessful due to *class* in closed status.

INFO: *volume*: failed to disable JRM, class closed down, class=*class*

Explanation

Disabling just resynchronization process on *volume* was unsuccessful due to *class* in closed status.

INFO: *object.volume*: failed to detach slice, class closed down, class=*class*

Explanation

Detaching slice *object.volume* was unsuccessful due to *class* in closed status.

INFO: *volume*: failed to restart volume, class closed down, class=*class*

Explanation

Restarting *volume* was unsuccessful due to *class* in closed status.

INFO: open error on *status slice object.volume*, class=*class*

Explanation

Open request sent to slice *object.volume* in *status* status terminated abnormally.

INFO: *class*: trying to identify class master, *details*

Explanation

Trying to identify the class master for shared class *class*.

INFO: *class*: identified class master, node=*node*

Explanation

The master for shared class *class* has been identified as node *node*.

INFO: *class*: searching class master

Explanation

Searching the master for shared class *class*.

INFO: *class*: class master found, node=*node*

Explanation

The master for shared class *class* was found to be node *node*.

INFO: *class*: class master not found

Explanation

The master for shared class *class* cannot be found.

INFO: *class*: got class master privilege

Explanation

Trying to obtain master privilege for shared class *class*.

INFO: *class*: broadcasted class master information to remote nodes

Explanation

The master information for shared class *class* has been broadcast to the remote nodes.

INFO: *class*: received confirmations of class master information from remote node *node*

Explanation

Received master confirmation for shared class *class* from remote node *node*.

INFO: waiting for outstanding event operations, *details*

Explanation

Waiting for event operations in process.

INFO: completed outstanding event operations

Explanation

Event operations in process has been completed.

INFO: *class*: trying to release class master privilege, *details*

Explanation

Trying to release master privilege for shared class *class*.

INFO: *class*: released class master privilege

Explanation

Master privilege for shared class *class* has been released.

**NFO: *proxy*: started to copy with OPC, source=*disk.volume*, target=*disk.volume*,
class=*class***

Explanation

Copying between proxy volume *proxy* and master volume with the OPC function started. *source* specifies the original slice name, *target* the mirror slice name, and *class* to which the original slice and mirror slice belong.

**INFO: *proxy*: completed copying with OPC, source=*disk.volume*, target=*disk.volume*,
class=*class***

Explanation

Copying between proxy volume *proxy* and master volume with the OPC function is completed. *source* specifies the original slice name, *target* the mirror slice name, and *class* to which the original slice and mirror slice belong.

**INFO: *proxy*: canceled copying with OPC, source=*disk.volume*, target=*disk.volume*,
class=*class***

Explanation

Copying between proxy volume *proxy* and master volume with the OPC function was cancelled. *source* specifies the original slice name, *target* the mirror slice name, and *class* to which the original slice and mirror slice belong.

INFO: *proxy*: EC session started, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

EC session between proxy volume *proxy* and master volume started. *source* specifies the original slice name, *target* the mirror slice name, and *class* to which the original slice and mirror slice belong.

**INFO: *proxy*: completed copying with EC, source=*disk.volume*, target=*disk.volume*,
class=*class***

Explanation

Copying between proxy volume *proxy* and master volume with the EC function is completed. *source* specifies the original slice name, *target* the mirror slice name, and *class* to which the original slice and mirror slice belong.

**INFO: *proxy*: canceled copying with EC, source=*disk.volume*, target=*disk.volume*,
class=*class***

Explanation

Copying between proxy volume *proxy* and master volume with the EC function was cancelled. *source* specifies the original slice name, *target* the mirror slice name, and *class* to which the original slice and mirror slice belong.

**INFO: *proxy*: EC session stopped, source=*disk.volume*, target=*disk.volume*,
class=*class***

Explanation

EC session between proxy volume proxy and master volume with the EC function was stopped. source specifies the original slice name, target the mirror slice name, and class to which the original slice and mirror slice belong.

INFO: proxy: EC session suspended, source=disk.volume, target=disk.volume, class=class

Explanation

EC session between proxy volume proxy and master volume with the EC function has been temporarily suspended. source specifies the original slice name, target the mirror slice name, and class to which the original slice and mirror slice belong.

INFO: proxy: EC session resumed, source=disk.volume, target=disk.volume, class=class

Explanation

EC session between proxy volume proxy and master volume with the EC function was resumed. source specifies the original slice name, target the mirror slice name, and class to which the original slice and mirror slice belong.

INFO: proxy: established BCV pair, STD=disk, BCV=disk, class=class

Explanation

BCV pair between the proxy group proxy and the master group was established. "STD" is the disk name of the standard device, "BCV" is the disk name of the BCV device, and "class" is the name of the class to which the standard and BCV devices belong.

INFO: proxy: completed copying with TimeFinder, source=disk, target=disk, class=class

Explanation

Copying between the proxy group proxy and the master group with TimeFinder is complete. "source" is the name of the copy source disk, "target" is the name of the copy destination disk, and "class" is the name of the class to which the copy source and destination disks belong.

INFO: proxy: canceled copying with TimeFinder, source=disk, target=disk, class=class

Explanation

Copying between the proxy group proxy and the master group with TimeFinder was canceled. "source" is the name of the copy source disk, "target" is the name of the copy destination disk, and "class" is the name of the class to which the copy source and destination disks belong.

INFO: *proxy*: canceled BCV pair, STD=*disk*, BCV=*disk*, class=*class*

Explanation

BCV pair between the proxy group *proxy* and the master group was canceled. "STD" is the disk name of the standard device, "BCV" is the disk name of the BCV device, and "class" is the name of the class to which the standard and BCV devices belong.

INFO: *proxy*: split BCV pair, STD=*disk*, BCV=*disk*, class=*class*

Explanation

BCV pair between the proxy group *proxy* and the master group was split. "STD" is the disk name of the standard device, "BCV" is the disk name of the BCV device, and "class" is the name of the class to which the standard and BCV devices belong.

INFO: *proxy*: re-established BCV pair, STD=*disk*, BCV=*disk*, class=*class*

Explanation

BCV pair between the proxy group *proxy* and the master group was re-established. "STD" is the disk name of the standard device, "BCV" is the disk name of the BCV device, and "class" is the name of the class to which the standard and BCV devices belong.

INFO: *proxy*: established SRDF pair, source=*disk*, target=*disk*, class=*class*

Explanation

SRDF pair between the proxy group *proxy* and the master group was established. "source" is the name of the source disk, "target" is the name of the target disk, and "class" is the name of the class to which the source and target disks belong.

INFO: *proxy*: completed copying with SRDF, source=*disk*, target=*disk*, class=*class*

Explanation

Copying between the proxy group *proxy* and the master group with SRDF is complete. "source" is the name of the copy source disk, "target" is the name of the copy destination disk, and "class" is the name of the class to which the copy source and destination disks belong.

INFO: *proxy*: canceled copying with SRDF, source=*disk*, target=*disk*, class=*class*

Explanation

Copying between the proxy group *proxy* and the master group with SRDF was canceled. "source" is the name of the copy source disk, "target" is the name of the copy destination disk, and "class" is the name of the class to which the copy source and destination disks belong.

INFO: *proxy*: canceled SRDF pair, source=*disk*, target=*disk*, class=*class*

Explanation

SRDF pair between the proxy group *proxy* and the master group was canceled. "source" is the name of the source disk, "target" is the name of the target disk, and "class" is the name of the class to which the source and target disks belong.

INFO: *proxy*: split SRDF pair, source=*disk*, target=*disk*, class=*class*

Explanation

SRDF pair between the proxy group *proxy* and the master group was split. "source" is the name of the source disk, "target" is the name of the target disk, and "class" is the name of the class to which the source and target disks belong.

INFO: *proxy*: re-established SRDF pair, source=*disk*, target=*disk*, class=*class*

Explanation

SRDF pair between the proxy group *proxy* and the master group was re-established. "source" is the name of the source disk, "target" is the name of the target disk, and "class" is the name of the class to which the source and target disks belong.

E.3.5 Internal Error Messages

***daemon*: severity: module: internal error, details**

Explanation

Internal error occurred. *details* gives the cause of the error, *daemon* gives the name of the daemon program, *severity* gives the severity of the message, and *module* the module name (usually the internal function name) where the event was detected.

Resolution

Collect investigation material and contact your local customer support.

E.4 Command Messages

Command messages will be sent to standard output or standard error output. Output format is as follows.

SDX:command: severity: message

The *command* gives the command name, *severity* the severity of the message, and *message* the message text body.

Command message will also be output to the following log file.

/var/opt/FJSVsdx/msglog/sdxservd.log

Output format is as follows.

Mon Day HH:MM:SS SDX:daemon: severity: message

Mon gives the month the message was output, *Day* the date, *HH* the hour, *MM* the minute, *SS* the second, *daemon* the daemon program name, *severity* the severity of the message, and *message* the message text body.

Depending on the syslog settings, message will be output to the console, and to the syslog log file.

Variable Names

Italicized words in the messages are variable names, and the actual output will vary depending on the situation. The meaning and the format of the variable names used in the message explanation is described below.

<i>class</i>	Class name
<i>disk</i>	Disk name
<i>group</i>	Group name
<i>igroup</i>	Lower level group name
<i>hgroup</i>	Higher level group name
<i>volume</i>	Volume name
<i>disk.volume</i>	Slice name
<i>object.volume</i>	Slice name
<i>object</i>	Object name (or physical disk name)
<i>master</i>	Master object name
<i>proxy</i>	Proxy object name
<i>status</i>	Object status
<i>device</i>	Physical disk name(cCtTdD, mphdI ,mplbI, emcpowerN) where <i>C</i> is the controller number, <i>T</i> is the target ID, <i>D</i> is the disk number, <i>I</i> is the instance number, and <i>N</i> is the emcpower device number.
<i>pslice</i>	Physical slice name(cCtTdDs{0-7}, mphd/s{0-7},mplb/s{0-7},emcpowerN{a-h}) where <i>C</i> is the controller number, <i>T</i> is the target ID, <i>D</i> is the disk number, <i>I</i> is the instance number, <i>N</i> is the emcpower device number, {0-7} is the slice number with 1 digit from 0 to 7, and {a-h} is the slice number with one alphabet from a to h.
<i>node</i>	Node identifier or node name.
<i>attribute</i>	Attribute name
<i>value</i>	Attribute value
<i>param</i>	Parameter name
<i>val</i>	Parameter value
<i>size</i>	Number of blocks (sectors), (decimal number)

<i>option</i>	Command options.
<i>usage</i>	Syntax when using command.
<i>letter</i>	Characters
<i>details</i>	Details
<i>string</i>	Other character strings

Explanation

Messages output by command are shown below in the order of its severity.

There are four levels of severity.

ERROR	This message is output to indicate the cause when a command terminates abnormally.
WARNING	This message is output as a warning when an abnormal event is detected. It will not always result in an abnormal termination of the command.
INFO	This message is output to record the command operation.
TO FIX	This message is output to show you how to fix the command. Usually it outputs a command syntax.

E.4.1 Error Messages

ERROR: connection timeout

Explanation

Connection failed due to no response from sdxservd daemon.

Resolution

Confirm that sdxservd daemon process is started normally.

ERROR: not privileged

Explanation

Executed user is not a superuser.

Resolution

You need to have superuser authority for execution.

ERROR: *option*: illegal option

Explanation

option is illegal.

Resolution

See the Fix Messages described later, or the "AppendixD Command Reference " in the manual.

ERROR: syntax error

Explanation

Executed command has a syntax error.

Resolution

See the Fix Messages described later, or the " AppendixD Command Reference " in the manual.

ERROR: *string*: name too long**Explanation**

Object name or node identifier specified with *string* is too long.

Resolution

Specify the proper object name or node identifier.

ERROR: *object*: name contains invalid character "*letter*"**Explanation**

Object name *object* contains an invalid character *letter*.

Resolution

You can only use alphanumeric characters, -(minus), and _(underscore) when specifying an object name.

ERROR: *object*: name starting with "_" or "-" is invalid**Explanation**

Object name starting with "-" (minus) or "_" (underscore) is invalid.

Resolution

Specify an object name that starts with an alphanumeric character.

ERROR: *device*: illegal physical disk name**Explanation**

Physical disk name *device* is illegal.

Resolution

See the " AppendixD Command Reference " in the manual

ERROR: *object*: object names must be unique within a class**Explanation**

Object names within a class must be unique.

Resolution

Specify a unique object name.

ERROR: *attribute*: invalid attribute name**Explanation**

The same *attribute* has already been specified.

Resolution

You can only specify an attribute once.

ERROR: *attribute=value*: invalid attribute value

Explanation

Attribute value *value* is invalid.

Resolution

See the " AppendixD Command Reference " in the manual.

ERROR: *attribute*:attribute name duplicated

Explanation

The same attribute *attribute* has already been specified.

Resolution

You can only specify an attribute once.

ERROR: *node*: node name duplicated

Explanation

The same node identifier has already been specified.

Resolution

You can only specify a node identifier once.

ERROR: too many nodes in scope

Explanation

There are too many nodes specified in scope.

Resolution

Up to 16 nodes can be specified in a scope.

ERROR: class: cannot operate shared objects, cluster control facility not ready

Explanation

Cannot operate shared object since cluster control is not operating.

Resolution

Start cluster control and try again.

ERROR: *node*: remote node cannot be specified for local class

Explanation

Remote node was specified for a local class.

Resolution

Do not specify a node identifier other than the one for current node for local class.

ERROR: node: remote node cannot be specified for root class

Explanation

Remote node was specified for root class.

Resolution

Do not specify a node identifier other than the one for current node for root class.

ERROR: current node must be specified as scope

Explanation

Current node is not specified for shared class scope.

Resolution

Always specify a node set which includes current node as scope.

ERROR: multi-nodes must be specified for shared class

Explanation

Only one node is specified for shared class scope.

Resolution

Multiple node identifiers must be specified in a scope.

ERROR: *node*: unknown node

Explanation

Node *node* does not exist.

Resolution

First check the cluster system environment. Then, try again after changing the node identifier.

ERROR: *class*: not shared by the current node

Explanation

Scope of *class* does not include the current node.

In the cluster system, *class* created in another node is specified.

Resolution

Execute command on a node sharing *class*. Or, change the specified class name.

ERROR: *node*: current node not active in *class*

Explanation

Current node *node* is a standby node in shared class *class*.

Resolution

Execute the command from another active node.

ERROR: *class*: too many disks in *class*

Explanation

There are too many disks registered with *class*.

Resolution

Create a new class.

ERROR: *device*: not connected to *node*

Explanation

The physical disk *device* has not been connected to the node *node*, or registration of the cluster resources is not completed properly.

Resolution

Check the system configuration and the resource configuration. If *device* has not been connected to *node*, specify the proper physical device name and re-execute the command. If no disk resource exists, execute automatic resource registration of the cluster and re-execute the command.

ERROR: *device*: not a shared device

Explanation

Physical disk *device* is not a shared disk unit, or cluster shared disk definition may not be correct.

Resolution

Check the system configuration and execute the command.

ERROR: *device*: already exists in *class*

Explanation

Physical disk *device* is already registered with *class*.

Resolution

Same physical disk cannot be registered with more than one class. Specify a correct physical disk.

ERROR: *device*: already exists in another class

Explanation

Physical disk *device* is already registered with another class. This class is not shared by current node.

Resolution

Same physical disk cannot be registered with more than one class. Specify the correct physical disk.

ERROR: *device*: no such device

Explanation

There is no physical disk *device*.

Resolution

Specify the correct physical disk name.

ERROR: *device*: cannot open, erro=*ermo***Explanation**

Cannot open physical disk *device*.

Resolution

Confirm that physical disk *device* is operating normally.

ERROR: *device*: not a hard disk**Explanation**

Physical disk *device* is not a hard disk.

Resolution

GDS cannot manage devices other than hard disks.

ERROR: *device*: disk driver *driver* not supported**Explanation**

Physical disk unit with the driver name *driver* is not supported.

Resolution

There is no resolution.

ERROR: *device*: illegal format**Explanation**

Physical disk *device* format is illegal.

Resolution

Check the format status.

ERROR: *object*: device busy**Explanation**

Object *object* is in use.

Resolution

Change to unused status and execute the command again.

ERROR: *object*: linked to a cluster service**Explanation**

Object *object* is used in a cluster application.

Resolution

Check the cluster environment settings.

ERROR: *object*: name already assigned to another object

Explanation

Object name *object* already exists within class. You cannot create multiple objects with the same name.

Resolution

Specify another name and execute the command again.

ERROR: cannot connect to sdxcl

Explanation

Connecting to sdxcl, a GDS cluster linked daemon was unsuccessful.

Resolution

When unable to identify the cause, collect investigation material and contact your local customer support.

ERROR: physical device driver returned an error, *ermo=ermo*

Explanation

Physical disk driver returned an error.

Resolution

Identify the cause by referring to error number and message log.

ERROR: special file operation failed, *ermo=ermo*

Explanation

Operating special file was unsuccessful.

Resolution

Identify the cause by referring to error number and GDS log message.

ERROR: sfdsk driver returned an error, *ermo=ermo*

Explanation

GDS driver returned an error.

Resolution

Identify the cause by referring to error number, GDS log message, and syslog message.

ERROR: sfdsk driver returned a temporary error, try again for a while

Explanation

GDS driver returned a temporary error.

Resolution

Execute the command again after a while.

ERROR: *class*: class closed down

Explanation

class is closed. You cannot operate objects within a closed class.

Resolution

A number of disk failures may have occurred. Identify the cause by referring to object status, GDS log message, and syslog message. For information on recovery, see "F.1.4 Class Status Abnormality."

ERROR: *class*: class closed down on another node

Explanation

class is closed on another node. You cannot operate objects within a closed class.

Resolution

Recover closed *class*. A number of disk failures may have occurred. Identify the cause by referring to object status, GDS log message, and syslog message. For information on recovery, see "F.1.4 Class Status Abnormality."

ERROR: keep disk cannot be specified for local or shared class

Explanation

Keep disk cannot be registered with local or shared class.

Resolution

Register keep disk with root class.

ERROR: too many keep disks specified

Explanation

The number of specified keep disks exceeds that of undefined disks

Resolution

When specifying keep disks, the same number or more should be specified as undefined disks.

ERROR: *class*: already root class exists

Explanation

You tried to create a new class when there is already a root *class*. You can only create one root class within a node.

Resolution

Specify the existing root *class* or a local class.

ERROR: *pslice* is not backup slice, correct slice tag and size by format(1M) command

Explanation

Physical slice *pslice* is not a backup slice.

Resolution

Use the format(1M) command to correct the tag and the size of physical slice *pslice*.

ERROR: *device* contains overlapping slices

Explanation

Physical disk *device* contains a physical slice with overlapping cylinder.

Resolution

Use the format (1M) command and correct the physical slice configuration.

ERROR: *device*: one or more free slices are required, change a slice size to 0 by format(1M) command

Explanation

There are no free physical slices, in other words a size 0 slice, in physical disk *device*.

Resolution

Use the format(1M) command to correct the physical slice configuration.

ERROR: *pslice*: invalid partition tag, tag=*tag_num*

Explanation

Partition tag value *tag_num* for physical slice *pslice* is invalid.

Resolution

Use the format (1M) command and correct the physical slice configuration.

ERROR: *device*: no unassigned disk space nor swap space, reserve enough space by format(1M) command

Explanation

Physical disk *device* does not have space area or swap area.

Resolution

Use the format (1M) command and create sufficient space area or swap area.

ERROR: *device*: no enough unassigned disk space nor swap space, reserve enough space by format(1M) command

Explanation

Physical disk *device* does not have sufficient space area or swap area.

Resolution

Use the format (1M) command and create sufficient space area or swap area.

ERROR: *device*: invalid physical device, managed by *driver*

Explanation

Specified physical disk is managed by *driver*, and is therefore invalid.

Resolution

Confirm the I/O and cluster system configuration, and specify the correct physical disk name.

ERROR: *device*: disk driver *driver* not supported for root class

Explanation

Physical disk device with driver name *driver* cannot be registered with root class.

Resolution

Register the disk with local class or shared class as needed.

ERROR: *device*: IDE disk cannot be specified as spare disk

Explanation

The *device* is an IDE disk. It can not be used as a spare disk.

Resolution

For spare disks, use disks other than IDE disks.

ERROR: *class*: no such class

Explanation

Cannot find *class*.

Resolution

Check GDS configuration.

ERROR: *group*: not a group

Explanation

group is not a group name.

Resolution

There is another object with the name *group* within class. Check configuration.

ERROR: *group* is a lower level stripe group

Explanation

Group *group* is a stripe group connected to another group. Disks and groups cannot be connected to or disconnected from *group*.

Resolution

Disconnect *group* from the higher level group as necessary.

ERROR: *group*: connected to a lower level stripe group

Explanation

Group *group* is connected to a lower level stripe group. Disks and groups cannot be connected to or disconnected from *group*.

Resolution

Disconnect the higher level stripe group of *group* from its higher level group as necessary.

ERROR: too many groups in *class*

Explanation

Class *class* already has the maximum number of groups possible.
You can create up to 100 groups within the root class, and up to 1024 groups within a local class or a shared class.

Resolution

Create a new class.

ERROR: too many disks and/or groups are connected to *group*

Explanation

A maximum number of disks or lower level groups are already connected to group *group*.

Resolution

There is no resolution. With mirror groups, the maximum number of objects you can connect is eight. With stripe groups and concatenation groups, it is sixty-four.

ERROR: *object*: smaller than stripe width of group *group*

Explanation

Since the available size of disk or lower level group indicated by *object* is smaller than the stripe width of group *group*, *object* cannot be connected to *group*.

Resolution

Either execute the command after indicating a disk or lower level group with sufficient size, or remove group *group* and adjust the stripe width to a smaller size.

ERROR: *disk*: no such disk

Explanation

There is no *disk*.

Resolution

Check GDS configuration.

ERROR: *object*: already connected to *group*

Explanation

Disk or group indicated by *object*, is already connected to group *group*.

Resolution

Indicate the correct disk name or group name.

ERROR: *disk* is a spare disk

Explanation

disk is a spare disk. You cannot connect a spare disk to a group.

Resolution

Change the disk attributes to undefined, and execute the command again.

ERROR: *object* not in *status* status**Explanation**

Object status is not *status* status.

Resolution

Confirm that object status is *status*, and execute the command again.

ERROR: *object* too small**Explanation**

The size of *object* is too small.

Resolution

Check the necessary object size and specify a larger object.

ERROR: another disk must be connected to *group***Explanation**

You must connect another disk to *group*.

Resolution

Connect another disk and execute the command again.

ERROR: invalid physical slice number, *pslice_num***Explanation**

Invalid physical slice number *pslice_num* is specified.

Resolution

For physical slice number *pslice_num*, specify an integer from 0 to 7, excluding 2.

ERROR: *pslice* is backup slice**Explanation**

Physical slice *pslice* is a backup slice.

Resolution

For physical slice number, specify an integer from 0 to 7, excluding 2.

ERROR: *pslice* is private slice**Explanation**

Physical slice *pslice* is a private slice.

Resolution

Do not specify a private slice for physical slice number.

ERROR: *pslice*: corresponding volume attributes must be specified

Explanation

Volume attributes corresponding to physical slice *pslice* is not specified.

Resolution

See the " AppendixD Command Reference " in the manual.

ERROR: *object*: invalid size

Explanation

Object size of *object* is invalid.

Resolution

See the " AppendixD Command Reference " in the manual.

ERROR: two or more keep disks cannot be connected to a group

Explanation

More than one keep disk cannot be connected to group.

Resolution

See the " AppendixD Command Reference " in the manual.

ERROR: two or more single disks cannot be connected to a group

Explanation

More than one single disk cannot be connected to a group.

Resolution

For details, see the " AppendixD Command Reference " in the manual.

ERROR: both keep and single disks cannot be connected to a group

Explanation

Both keep and single disks cannot be connected to a group at the same time.

Resolution

For details, see the " AppendixD Command Reference " in the manual.

ERROR: keep disk cannot be connected to existing group

Explanation

Keep disk cannot be connected to an existing group.

Resolution

See the " AppendixD Command Reference " in the manual.

ERROR: *disk*: single disk cannot be connected to existing group

Explanation

Single disk *disk* cannot be connected to the existing group.

Resolution

For details, see the " AppendixD Command Reference " in the manual.

ERROR: two or more IDE disks cannot be connected to a group

Explanation

More than one IDE disk cannot be connected to a group.

Resolution

When mirroring an IDE disk, use it with disks other than IDE disks.

ERROR: *disk*: not a bootable device

Explanation

Disk *disk* cannot be booted.

Resolution

Confirm the disk configuration. When you cannot specify the cause, collect investigation material and contact your local customer support.

ERROR: too few valid configuration database replicas

Explanation

There is not sufficient valid configuration database for class.

This message is displayed when majority of disks registered with class is unavailable.

Leaving it as is may cause serious problems.

Resolution

For details, see "F.1.4 Class Status Abnormality."

ERROR: *group*: no ENABLE disk in group

Explanation

Group *group* does not have a disk in ENABLE status connected.

Resolution

See "F.1.2 Disk Status Abnormality" and recover disks connected to the group.

ERROR: *msec*: invalid delay value

Explanation

Delay time *msec* is invalid.

Resolution

See the " AppendixD Command Reference " in the manual.

ERROR: *lgroup*: mirror group cannot be connected to another group

Explanation

Since group *lgroup* is a mirror group, it cannot be connected to another group.

Resolution

There is no resolution.

ERROR: *hgroup*: any group cannot be connected to a concatenation group

Explanation

Since group *hgroup* is a concatenation group, any group cannot be connected to *hgroup*.

Resolution

There is no resolution.

ERROR: *lgroup*: same type as higher level group *hgroup*

Explanation

Since the type attributes for group *lgroup* and higher level group *hgroup* are the same, *lgroup* cannot be connected to *hgroup*.

Resolution

There is no resolution.

ERROR: *hgroup*: same name as a lower level group

Explanation

For the higher level group name, name *hgoup*, which is the same as the lower level group was indicated.

Resolution

Different names must be indicated for higher level group and lower level group.

ERROR: *size*: must be more than zero

Explanation

size must be a positive integer.

Resolution

Specify the correct size.

ERROR: *size*: invalid size

Explanation

The indicated size *size* is invalid.

Resolution

See the " AppendixD Command Reference " in the manual.

ERROR: *group*: no such group

Explanation

There is no *group*.

Resolution

Specify the correct group name.

ERROR: *group*: not the highest level group

Explanation

Group *group* is not the highest level group..

Resolution

Disconnect *group* from the higher level group as necessary.

ERROR: too many volumes exist in *object*

Explanation

The maximum number of volumes, already exists in the object *object*, that is a class, a group, or a single disk.

Resolution

No resolution.

ERROR: too many volumes with physical slices in *object*

Explanation

In the object *object*, the number of volumes which consist of physical slices has reached maximum, that is six. *object* is a name of group or single disk.

Resolution

No resolution.

ERROR: *group*: volume with physical slice(s) cannot be created

Explanation

Although a physical slice cannot be created in group *group*, an attempt was made to create a volume with one or more physical slices.

Resolution

When creating a volume in *group* using the `sdxvolume -M` command, specifying the `-a pslice=off` option will allow you to create a volume without a physical slice. Also if *group* is a mirror group, connecting more than one disk to *group* will allow you to create a volume with physical slice.

ERROR: *object*: no enough space

Explanation

There is not enough space in the group or single disk indicated by *object*.

Resolution

Change the size you specify as needed.

ERROR: *disk*: not a single disk

Explanation

The disk *disk* is not a single disk.

Resolution

Specify an appropriate disk.

ERROR: *disk*: *status* disk connected to *group*

Explanation

disk in *status* status is connected to *group*.

Resolution

Restore the status of *disk* and execute the command again.

ERROR: *status* disk exists in *group*

Explanation

The *disk* in *status* status is connected to *group*, or it connected to the lower level group of *group*.

Resolution

Confirm the disk status, and please execute a command again after canceling the *status* status if necessary.

ERROR: *disk*: no such disk in *group*

Explanation

disk is not connected to *group*.

Resolution

Specify the correct disk name or group name.

ERROR: *disk*: *device*: device busy

Explanation

Disk *disk* (physical disk name is *device*) is in use.

Resolution

Change to unused status, and execute the command again.

ERROR: too many nodes specified

Explanation

Too many nodes were specified.

Resolution

Specify the nodes included in the scope of the class.

ERROR: *option*: cannot be specified for root nor local class

Explanation

Command option *option* cannot be specified for a root or local class.

Resolution

See "Command Reference."

ERROR: *node*: not in scope, class=*class*

Explanation

Node *node* is not included in the scope of the class *class*.

Resolution

Confirm your GDS configuration, specify a node included in the scope of the class, and try the command again.

ERROR: *volume*: cannot start, class closed down, node=*node*

Explanation

Starting volume <volume> was unsuccessful due to the class to which *volume* belongs in closed status.

Resolution

Recover the closed class. A number of disk failures may have occurred. Identify the cause by referring to object status, GDS log message, and syslog message. For information on recovery, see "F.1.4 Class Status Abnormality."

ERROR: *volume*: no such volume

Explanation

There is no *volume*.

Resolution

Specify the correct volume name.

ERROR: *object* in *status* status

Explanation

object is in *status* status.

Resolution

Confirm the object status, and cancel the status if necessary.

ERROR: lock is set on volume *volume*, node=*node*

Explanation

The "Lock volume" mode for volume *volume* on node *node* is turned on.

Resolution

Turn off the "Lock volume" mode, or use the -e unlock option as necessary.

ERROR: *volume*: cannot stop, class closed down, node=*node*

Explanation

Stopping volume *volume* was unsuccessful due to the class to which *volume* belongs in closed status.

Resolution

Recover the closed class. A number of disk failures may have occurred. Identify the cause by referring to object status, GDS log message, and syslog message. For information on recovery, see "F.1.4 Class Status Abnormality."

ERROR: *object.volume*: status slice exists in object

Explanation

Slice *object.volume* in *status* status exists in *object*.

Resolution

Confirm the object status, and cancel the status if necessary.

ERROR: *object* in *status* status, node=*node*

Explanation

object on *node* is in *status* status.

Resolution

Confirm the object status, and cancel the status if necessary.

ERROR: spare disk connected for *disk*

Explanation

Spare disk is connected instead of *disk*.

Resolution

First recover the *disk* status.

ERROR: *object.volume* is only valid slice

Explanation

object.volume is the only valid slice within volume *volume*. You cannot continue your operation because data within *volume* may be lost if you continue your operation.

Resolution

If *volume* is a mirror volume, you will be able to continue your operation by recovering the mirroring (for example, by connecting a new disk to group). If *volume* is not a mirror volume, remove *volume* as necessary.

ERROR: *object*: the last disk or group in lower level group *group*

Explanation

Object *object* is the only disk or group connected to the lower level group *group*. You cannot disconnect *object* from *group*.

Resolution

Disconnect *group* from the higher level group as necessary.

ERROR: *disk*: not connected to the end of concatenation group *group***Explanation**

Disk *disk* is not the disk that was concatenated last to the concatenation group *group*.
You cannot disconnect *disk* from *group*.

Resolution

There is no resolution.

ERROR: *disk*: disk space is assigned to volume *volume***Explanation**

Area on disk *disk* is allocated to volume *volume*. You cannot disconnect *disk* from the concatenation group.

Resolution

First remove *volume* if necessary.

ERROR: *lgroup*: no such group in *hgroup***Explanation**

Group *lgroup* is not connected to group *hgroup*.

Resolution

Indicate the correct group name.

ERROR: one or more volumes exist in *group***Explanation**

One or more volumes exist in group *group*.

Resolution

First, remove the volume if necessary.

ERROR: *disk* connected to *group***Explanation**

disk is connected to *group*.

Resolution

First, disconnect the disk from group if necessary.

ERROR: *disk*: The last ENABLE disk in class cannot be removed

Explanation

When there is a disk in SWAP or DISABLE status in class, you cannot remove the last disk in ENABLE status.

Resolution

First recover the disk in SWAP or DISABLE status. Or, you can register a new disk with class.

ERROR: one or more volumes exist in class

Explanation

Volume exists in *disk*.

Resolution

First, remove the volume as needed.

ERROR: one or more groups exist in class

Explanation

Group exists in *class*.

Resolution

First, remove the group as needed.

ERROR: *disk: status* disk exists in class

Explanation

disk in *status* status exists in *class*.

Resolution

First, restore the *disk* in *status* status.

ERROR: *volume: status* volume exists in class, node=*node*

Explanation

volume in *status* status on *node* exists in *class*.

Resolution

First, change the volume status as needed.

ERROR: *disk*: no such disk

Explanation

There is no disk *disk*.

Resolution

Specify the correct disk name.

ERROR: *volume*: not associated with *object*

Explanation

You cannot specify slice by combining disk or group indicated by *object*, and volume indicated by *volume*.

Resolution

See the Command Reference and indicate the correct disk name, group name or volume name.

ERROR: *volume* is single volume

Explanation

The *volume* is a single volume.

Resolution

Specify the correct parameter.

ERROR: *volume*: not a mirror volume

Explanation

Volume *volume* is not a mirror volume.

Resolution

There is no resolution.

ERROR: *object*: not connected to the highest level group

Explanation

Disk or group indicated by *object* is not connected to the highest-level group.

Resolution

See the Command Reference and indicate the correct disk name or group name.

ERROR: *disk*: The last ENABLE disk in class cannot be swapped

Explanation

You cannot swap the last ENABLE status disk in class.

Resolution

Depending on the configuration, use a different method to avoid such situations. For example, you can register a new disk with class.

ERROR: *disk*: keep disk cannot be swapped out

Explanation

Keep disk *disk* cannot be swapped.

Resolution

Change the type attribute of *disk*, and execute the command again.

ERROR: *disk: volume in status status*

Explanation

There is a volume related to disk in *status* status.

Resolution

First, recover the volume in *status* status.

ERROR: *disk: the highest level group is not a mirror group*

Explanation

The highest level group of disk *disk* is not a mirror group. You cannot swap *disk*.

Resolution

There is no resolution.

ERROR: *disk: device: cannot open, errno=errno*

Explanation

Cannot open physical disk *device*.

Resolution

Confirm that physical disk *device* is operating normally.

ERROR: *disk: device: not a hard disk*

Explanation

Physical disk *device* is not a hard disk.

Resolution

Specify a hard disk.

ERROR: *disk: device: illegal format*

Explanation

Physical disk *device* format is illegal.

Resolution

Check the format of the physical disk.

ERROR: *disk busy -/dev/rdisk/pslice*

Explanation

Physical disk *pslice* is in use.

Resolution

Change to unused status and execute the command again.

ERROR: *disk: device: linked to a cluster service*

Explanation

Physical disk *device* is used in a cluster application.

Resolution

Check the settings for cluster environment.

ERROR: *disk: device: not enough size*

Explanation

Physical disk size is too small.

Resolution

Specify a physical disk with sufficient size.

ERROR: *disk: device: not connected to node*

Explanation

Physical disk *device* is not connected to *node*. Or, cluster shared disk definition may not be correct.

Resolution

Check the system configuration and execute the command.

ERROR: *disk: device: invalid disk, managed by driver*

Explanation

Specified disk (corresponding physical disk name is *device*) is managed by *driver*, and is therefore invalid.

Resolution

Confirm the I/O and cluster system configuration, and specify the correct physical disk name.

ERROR: *object.volume: status slice exists in group*

Explanation

Slice *object.volume* in *status* status already exists in group *group*.

Resolution

Check the object status and cancel *status* status if necessary.

ERROR: *object: read error, ermo=ermo*

Explanation

Read error occurred on the disk indicated by *object*, or on a disk connected to group indicated by *object*, or on a disk connected to a lower level of the group indicated by *object*.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and promptly recover the disk. Otherwise, indicate another disk or group that is normal.

ERROR: *object.volume* : no such slice

Explanation

You cannot specify slice by combining disk or group indicated by *object*, and volume indicated by *volume*.

Resolution

See the Command Reference and indicate the correct disk name, group name or volume name.

ERROR: *disk*: I/O error not occur

Explanation

An I/O error has not been occurred on the disk *disk*. There is no need to restore *disk* using the `sdxfix -D` command.

Resolution

There is no resolution.

ERROR: *volume*: cannot restart to copy, cancel current interrupted copy operation by `sdxcopy` command with `-C` option

Explanation

Copying process within volume could not be resumed.

Resolution

Execute `sdxcopy -C` as needed to cancel the interrupted copying process and resume copying.

ERROR: *attribute=value*: cannot modify type attribute of root class

Explanation

Type attribute of root class cannot be changed.

Resolution

There is no resolution.

ERROR: *class*: class names must be unique within a domain

Explanation

Class with the same name as *class* already exists within the cluster domain.

Resolution

Specify a different class name.

ERROR: one or more disks not connected to *node*

Explanation

There is a disk not connected to *node*.

Resolution

Check the hardware configuration and the disk configuration within class.

ERROR: *disk*: IDE disk cannot be specified as spare disk

Explanation

The *disk* is an IDE disk. It can not be used as a spare disk.

Resolution

For spare disks, use disks other than IDE disks.

ERROR: *object*: no such object

Explanation

There is no *object*.

Resolution

Specify the correct object name.

ERROR: *class*: shared objects information not yet available, try again for a while

Explanation

Share object information is not available yet.

Resolution

Wait until cluster control is activated, and execute the command again.

ERROR: *node*: node in stopped status

Explanation

You cannot proceed with operation since *node* is in stopped status.

Resolution

Start the *node* and try again.

ERROR: *node*: node in abnormal status

Explanation

node is in an abnormal status and cannot proceed with operation.

Resolution

Confirm the normal startup of *node* and try again.

ERROR: cluster communication failure

Explanation

Cannot proceed with operation since communication failed with cluster.

Resolution

Check that the cluster system and GDS are operating normally.
After recovering, try again.

ERROR: cluster communication failure, sdxerno=*sdxerno*

Explanation

Cannot proceed with operation since communication failed with cluster.

Resolution

Check that the cluster system and GDS are operating normally.
After recovering, try again.

ERROR: cluster communication failure, remote-node=*node*, sdxerno=*sdxerno*

Explanation

Cluster communication with remote node *node* failed. Operation cannot be performed.

Resolution

Check that the cluster system and GDS are operating normally.
After recovering, try again.

ERROR: *class*: not a root class

Explanation

Class *class* is not a root class.

Resolution

See the " AppendixD Command Reference " in the manual and specify the correct class name.
cluster communication failure

ERROR: *disk*: not a keep disk

Explanation

Disk *disk* is not a keep disk.

Resolution

See the " AppendixD Command Reference " in the manual and specify the correct disk name.

ERROR: *disk*: not connected to any group

Explanation

Disk *disk* is not connected to a group.

Resolution

See the " AppendixD Command Reference " in the manual and specify the correct disk name.

ERROR: *volume: status* volume exists in *group*

Explanation

There is *volume* in *status* status in *group*.

Resolution

Recover volume status, and execute the command again.

ERROR: *disk*: not a system disk

Explanation

Disk *disk* is not a system disk.

Resolution

Specify the correct disk name, and execute the command again.

ERROR: *device*: all of system disk(s) must be in *class*

Explanation

System disk *device* is not registered with *class*.

Resolution

See "5.2.1 System Disk Settings" and complete preparation correctly. Then execute the command again.

ERROR: *disk*: all of system disk(s) must be specified

Explanation

System disk *disk* is not specified.

Resolution

Specify the correct disk, and execute the command again.

ERROR: *disk*: two or more disks connected to *group*

Explanation

group to which *disk* is connected has two or more disks connected.

Resolution

See "5.2.1 System Disk Settings" and complete preparation correctly. Then execute the command again.

ERROR: root file system not mounted on volume

Explanation

Root file system is not mounted on volume.

Resolution

Confirm the configuration, and see the " AppendixD Command Reference " in the manual.

ERROR: illegal slice name

Explanation

Slice name includes a "." (period).

Resolution

Specify a correct slice name.

ERROR: *disk.volume* cannot be operated on the current node, take over by *sdxslice* command with -T option

Explanation

Slice *disk.volume* cannot be operated on the current node.

Resolution

Take over the slice by executing *sdxslice -T* command.

ERROR: *volume*: physical slice attribute value is off

Explanation

The physical slice attribute value of volume *volume* is "off." A slice in a volume without physical slices cannot be detached.

Resolution

Retry the command after turning the physical slice attribute of *volume* to be "on" using *sdxattr -V* command.

ERROR: *object*: device busy on node *node*

Explanation

object is in use on *node*.

Resolution

Change to unused status and execute the command again.

ERROR: *class*: not a shared class

Explanation

class is not a shared class.

Resolution

Specify a shared class.

ERROR: *param*: invalid parameter name

Explanation

Parameter name *param* is invalid.

Resolution

See the " AppendixD Command Reference " in the manual.

ERROR: *param =val*: invalid parameter value

Explanation

Parameter value *val* is invalid

Resolution

See the " AppendixD Command Reference " in the manual.

ERROR: *param*: parameter name duplicated

Explanation

The same parameter name *param* has already been specified.

Resolution

You can only use a parameter name once.

ERROR: *copy_concurrency=val*: value more than or equal to the number of actually running copy operations must be specified

Explanation

Value smaller than the number of copying currently in process has been specified for *copy_concurrency* parameter.

Resolution

Set a value more than or equal to the number of copying currently in process for *copy_concurrency* parameter.

ERROR: *mode=string*: access mode duplicated

Explanation

Multiple access modes have been specified.

Resolution

Specify only one access mode.

ERROR: *mode=string*: invalid access mode

Explanation

Access mode value *string* is invalid.

Resolution

See command reference manual.

ERROR: *volume*: already started with different access mode, *node=node*

Explanation

sdxvolume -M command was executed with *-e mode=val* option on node *node*, but volume *volume* had already been started with a different access mode.

Resolution

Use *sdxvolume -F* command to stop volume *volume*, and execute again.

ERROR: *volume*: related to proxy volume *proxy*

Explanation

Volume *volume* is a master volume related to proxy volume *proxy*.

Resolution

Use *sdxproxy Break* command to cancel the relationship of master volume *volume* and proxy volume *proxy*, and execute again.

ERROR: *volume*: related to master volume *master*

Explanation

Volume *volume* is a proxy volume related to master volume *master*.

Resolution

Use `sdxproxy Break` command to cancel the relationship of master volume *volume* and proxy volume *volume*, and execute again.

ERROR: *volume*: related to proxy volume *proxy* with EC

Explanation

EC session exists between volume *volume* and proxy volume *proxy*.

Resolution

Use `sdxproxy Break` command to cancel the relationship of master volume *master* and proxy volume *volume*, and execute again.

ERROR: *volume*: related to master volume *master* with EC

Explanation

EC session exists between volume *volume* and master volume *master*.

Resolution

Use `sdxproxy Break` command to cancel the relationship of master volume *master* and proxy volume *volume*, and execute again.

ERROR: *volume*: related to proxy volume *proxy* with TimeFinder

Explanation

There is a BCV pair between the volume *volume* and the proxy volume *proxy*.

Resolution

Break the relation between the master volume *volume* and the proxy volume *proxy* with the `sdxproxy Break` command as necessary, and try the command again.

ERROR: *volume*: related to master volume *master* with TimeFinder

Explanation

There is a BCV pair between the volume *volume* and the master volume *master*.

Resolution

Break the relation between the master volume *volume* and the proxy volume *proxy* with the `sdxproxy Break` command as necessary, and try the command again.

ERROR: *volume*: related to proxy volume *proxy* with SRDF

Explanation

There is an SRDF pair between the volume *volume* and the proxy volume *proxy*.

Resolution

Break the relation between the master volume *volume* and the proxy volume *proxy* with the `sdxproxy Break` command as necessary, and try the command again.

ERROR: *volume*: related to master volume *master* with SRDF

Explanation

There is an SRDF pair between the volume *volume* and the master volume *master*.

Resolution

Break the relation between the master volume *volume* and the proxy volume *proxy* with the `sdxproxy Break` command as necessary, and try the command again.

ERROR: *volume*: parted proxy volume

Explanation

Volume *volume* is a proxy volume parted from master volume.

Resolution

Use `sdxproxy Rejoin` command to rejoin the volume with master volume, or `sdxproxy Break` command to cancel the relationship with master volume, and execute again.

ERROR: *group*: related to proxy group *proxy*

Explanation

Group *group* is a master group related to proxy group *proxy*.

Resolution

Use `sdxproxy Break` command to cancel the relationship between master group *group* and proxy group *proxy*, and execute again.

ERROR: *group*: related to master group *master*

Explanation

Group *group* is a proxy group related to master group *master*.

Resolution

Use `sdxproxy Break` command to cancel the relationship between master group *group* and proxy group *proxy*, and execute again.

ERROR: *volume*: related to master or proxy volume

Explanation

Volume volume is related to either master volume or proxy volume.

Resolution

Use `sdxproxy Break` command to cancel the relationship between master and proxy, and execute again.

ERROR: *volume*: joined to master volume *master*

Explanation

Volume volume is a proxy volume joined to master volume *master*.

Resolution

Use `sdxproxy Part` command to part volume volume from master volume *master*, or `sdxproxy Break` command to cancel the relationship with master volume, and execute again.

ERROR: *volume*: copying from master volume *master*

Explanation

Data is being copied from master volume *master* to volume volume.

Resolution

After the copying process is complete, execute again.

ERROR: *volume*: copying from proxy volume *proxy*

Explanation

Data is being copied from proxy volume *proxy* to volume volume.

Resolution

After the copying process is complete, execute again.

ERROR: *class* is root class

Explanation

Class *class* is root class. Root class objects cannot be related as master or proxy.

Resolution

There is no resolution.

ERROR: *object*: not volume nor group

Explanation

Object *object* is neither a volume or a group.

Resolution

Check the GDS configuration and specify the correct volume name or group name. Execute again.

ERROR: different types of objects, master=*master*, proxy=*proxy*

Explanation

Different types of objects were specified as master and proxy.

Resolution

Specify a pair of volumes or groups as master and proxy.

ERROR: *object*: same name as master

Explanation

The object name specified for proxy is the same as the master's name, *object*.

Resolution

Different object names must be specified for master and proxy.

ERROR: *group*: not a mirror group

Explanation

Group *group* is not a mirror group.

Resolution

There is no resolution.

ERROR: no volume exists in *group*

Explanation

There are no volumes in group *group*.

Resolution

Use `sdxvolume -M` command to create a volume within group *group*, and execute again.

ERROR: too many proxy volumes are related to *master*

Explanation

There are too many proxy volumes related to master object *master*.

Resolution

Check the GDS configuration. For details on the number of proxy volumes, see "A.1.10 Number of Proxy Volumes."

ERROR: *master*: corresponding proxy volume name must be specified

Explanation

There is no corresponding proxy volume name specified for master volume *master*.

Resolution

See the "D.13 `sdxproxy` – Proxy object operations."

ERROR: *proxy*: proxy volume name duplicated

Explanation

The same proxy volume name *proxy* is specified for more than one master volume.

Resolution

Specify a different name, and execute again.

ERROR: *volume*: no such volume in *group*

Explanation

Volume *volume* does not exist.

Resolution

Specify the correct volume name or group name, and execute again.

ERROR: *object*: object name duplicated

Explanation

Proxy object name *object* has been specified more than once, or the group to which the proxy volume *object* belongs, and proxy volume *object* itself have been specified simultaneously.

Resolution

Proxy object can be specified only once.

ERROR: *proxy*: already parted

Explanation

Proxy volume *proxy* is already parted.

Resolution

Nothing needs to be done.

ERROR: one point copy not available

Explanation

Instant snapshot cannot be created, since OPC (One Point Copy) function is not available.

Resolution

The disk unit does not provide OPC function, or the OPC function is not available for some reason. Check the hardware configuration of the disk unit.

ERROR: *proxy*: already joined

Explanation

Proxy volume *proxy* is already joined.

Resolution

Nothing needs to be done.

ERROR: *proxy*: not joined to master

Explanation

Proxy volume *proxy* is not joined to master volume.

Resolution

Use `sdxproxy Rejoin` command to join the proxy volume *proxy* to master volume, and execute again.

ERROR: *proxy*: no such proxy object

Explanation

Proxy object *proxy* cannot be found.

Resolution

Specify the correct proxy object name.

ERROR: *master*: no such master object

Explanation

Master object *master* cannot be found.

Resolution

Specify the correct master object name.

ERROR: *volume*: exists in proxy group

Explanation

Volume *volume* is a proxy volume in proxy group.

Resolution

Specify a proxy group and execute again.

ERROR: *group*: not a proxy group

Explanation

group is not a proxy group name.

Resolution

Check the GDS configuration, and specify a proxy group name.

ERROR: *volume*: copying with EC

Explanation

Volume *volume* is in the process of EC copy.

Resolution

After the copying process is complete, execute again.

ERROR: *volume*: copying with OPC

Explanation

Volume *volume* is in the process of OPC copy.

Resolution

After the copying process is complete, execute again.

ERROR: *volume*: copying with TimeFinder

Explanation

Volume *volume* is being copied with TimeFinder.

Resolution

Execute the command again after the copying process is complete.

ERROR: *volume*: copying with SRDF

Explanation

Volume *volume* is being copied with SRDF.

Resolution

Execute the command again after the copying process is complete.

ERROR: *volume*: related to same master volume as proxy *proxy*, master=*master*

Explanation

Volume *volume* is related to the same master volume *master* as proxy volume *proxy*.

Resolution

Specify the correct volume name. For details, see "Appendix D Command Reference" manual.

ERROR: master and proxy exist in same group *group*

Explanation

The specified master volume and proxy volume exist in the same group *group*.

Resolution

Specify volumes that exist in different groups.

ERROR: *proxy*: joined to master with EC, rejoin them by soft copy and try again

Explanation

Slices cannot be swapped since EC session in process between proxy *proxy* and master.

Resolution

First, part master and proxy, or cancel the master-proxy relationship. Rejoin the two by specifying the `-e softcopy` option, and execute `sdxproxy Swap` command again.

ERROR: *proxy*: joined to master with TimeFinder, rejoin them by soft copy and try again

Explanation

Slices cannot be swapped because there is a BCV pair between the proxy *proxy* and the master.

Resolution

Part the master and the proxy once or break their relation as necessary, rejoin them indicating the `-e softcopy` option, and try the `sdxproxy Swap` command again.

ERROR: *proxy*: joined to master with SRDF, rejoin them by soft copy and try again

Explanation

Slices cannot be swapped because there is an SRDF pair between the proxy *proxy* and the master.

Resolution

Part the master and the proxy once or break their relation as necessary, rejoin them indicating the `-e softcopy` option, and try the `sdxproxy Swap` command again.

ERROR: *volume*: proxy volume cannot be specified when using TimeFinder

Explanation

You cannot specify *volume* when you perform parting, re-joining or restoring because there is an BCV pair between the proxy volume *volume* and the master.

Resolution

If you wish to perform parting, re-joining, or restoring, you must specify the group to which *volume* belongs when executing the command.

ERROR: *volume*: proxy volume cannot be specified when using SRDF

Explanation

You cannot specify *volume* when you perform parting, re-joining or restoring because there is an SRDF pair between the proxy volume *volume* and the master.

Resolution

If you wish to perform parting, re-joining, or restoring, you must specify the group to which *volume* belongs when executing the command.

E.4.2 Warning Messages

WARNING: spare disk *disk* too small

Explanation

The size of *disk* is too small and may not function as a spare disk.

Resolution

Specify a larger disk, and execute the command again.

WARNING: *group*: free blocks are reduced

Explanation

Free blocks on *group* was reduced.

Resolution

You may not be able to create a volume with sufficient capacity.
Execute commands as needed and attempt recovery.

WARNING: another disk must be connected to *group*

Explanation

You must connect another disk to *group*.

Resolution

Connect another disk.

WARNING: *object*: copying not completed successfully

Explanation

Synchronization copying did not complete successfully.

Resolution

Disk failure may have occurred. Identify the caused by referring to GDS log message and syslog message.

WARNING: *object*: gave up wait for the completion of copying by a cancel request

Explanation

Synchronization copying was canceled.

Resolution

Check the *object* status. If it is in the process of copying, wait for its completion. If not, retry the copying again. For details, see the " AppendixD Command Reference " in the manual.

WARNING: cannot set use-nvramrc? parameter to true, NVRAMRC not empty

Explanation

Since a command is already registered with NVRAMRC, use-nvramrc? parameter cannot be set to "true."

Resolution

For details, see "A.2.9 System Disk Mirroring."

WARNING: *group*: no spare disk available

Explanation

There is no valid spare disk in *group*.

Resolution

Define a spare disk as needed.

WARNING: *object.volume*: cannot attached due to in *status* status

Explanation

Slice *object.volume* could not be attached since it is in *status* status.

Resolution

Check the slice status and cancel *status* status as needed.

WARNING: *disk.volume*: special file(s) not found

Explanation

Special file for slice *disk.volume* could not be found.

Resolution

No particular resolution required.

WARNING: *node*: node in stopped *status*

Explanation

node is in STOP status.

Resolution

No particular resolution required. However, promptly activating the node is recommended.

WARNING: *node*: node in abnormal status

Explanation

node is in abnormal status.

Resolution

No particular resolution required. However, promptly recovering the node and activating it normally is recommended.

WARNING: *object*: already in *status* status

Explanation

object is already in *status* status.

Resolution

No particular resolution required.

WARNING: *object*: no need to update attribute value

Explanation

There was no need to change *object* attributes.

Resolution

No particular resolution required.

WARNING: *pslice*: entry in /etc/vfstab not updated, unsupported file system type *fstype*

Explanation

Entry in the /etc/vfstab file relevant to the physical slice *pslice* was not updated since *fstype* is a file system type not supported by the sdxroot command.

Resolution

Edit the /etc/vfstab file or change the file system configuration according to need. If *fstype* is sfxfs or sfcs, refer to "A.2.11 GFS Partition on a Keep Disk."

WARNING: correct /etc/vfstab and file system configuration before rebooting, or system may not be booted

Explanation

Invalid entry is included in the /etc/vfstab file. If such file system entry as /(root), /usr, and /var essential to system start-up is invalid, the system cannot be started when rebooted.

Resolution

Refer to a GDS's WARNING message output just before this message and pinpoint the invalid entry in the /etc/vfstab file. Be sure to repair the /etc/vfstab file and the file system configuration as appropriate prior to rebooting the system. If entry in a GFS file system is invalid, refer to "A.2.11 GFS Partition on a Keep Disk."

WARNING: *volume*: entry in /etc/vfstab not updated, unsupported file system type *fstype*

Explanation

Since file system type *fstype* is not supported, entry in /etc/vfstab file for *volume* could not be updated.

Resolution

Directly edit /etc/vfstab as necessary.

WARNING: ignored parameter *param* for EC

Explanation

Since EC function will be used for copying, parameter *param* was ignored.

WARNING: ignored parameter *param* for OPC

Explanation

Since EC function will be used for copying, parameter *param* was ignored.

WARNING: ignored parameter *param* for TimeFinder

Explanation

Parameter *param* was ignored for copying with TimeFinder.

WARNING: ignored parameter *param* for SRDF

Explanation

Parameter *param* was ignored for copying with SRDF.

E.4.3 Information Messages

INFO: waiting for a response from sdxservd daemon...

Explanation

It is waiting for response from sdxservd daemon.

INFO: *class*: created class

Explanation

class was created.

INFO: *disk*: created disk

Explanation

disk was registered.

INFO: *device*: disabled access to physical special files

/dev/rdisk/devices*

/dev/dsk/devices*

Explanation

Physical special files can no longer be accessed.

INFO: *class*: removed class

Explanation

class was removed.

INFO: *group*: created group

Explanation

group was created.

INFO: *disk*: connected disk to group *group*

Explanation

disk was connected to *group*.

INFO: *igroup*: connected group to another group *hgroup*

Explanation

Group *igroup* was connected to another group *hgroup*.

INFO: *group*: removed group

Explanation

group was removed.

INFO: *object*: waiting for the completion of copying...

Explanation

Waiting for synchronization copying to complete.

INFO: *object*: copying completed successfully

Explanation

Synchronization copying completed successfully.

INFO: *volume*: created volume

Explanation

volume was created.

INFO: *volume*: started volume on node *node*

/dev/sfdsk/class/rdisk/volume

/dev/sfdsk/class/dsk/volume

Explanation

volume was started on *node*. You can now access via special file.

INFO: *volume*: stopped volume on node *node*

Explanation

volume was stopped on *node*.

INFO: *volume*: removed volume

Explanation

volume was removed.

INFO: *disk*: disconnected disk from group *group*

Explanation

disk was disconnected from *group*.

INFO: *lgroup*: disconnected group from group *hgroup*

Explanation

Group *lgroup* was disconnected from *hgroup*.

INFO: *group*: free blocks are increased

Explanation

Free blocks on *group* increased.

INFO: *disk*: removed disk

Explanation

disk was removed from class.

INFO: *device*: enabled access to physical special files

/dev/rdisk/devices*

/dev/dsk/devices*

Explanation

You can now access to physical special files.

INFO: *disk.volume*: temporarily detached slice from volume *volume*

Explanation

Slice *disk.volume* was temporarily disconnected from *volume*.

INFO: *disk.volume*: enabled slice

/dev/sfdsk/class/rdisk/disk.volume

/dev/sfdsk/class/dsk/disk.volume

Explanation

Slice *disk.volume* can now be accessed partly from volume.

Use the special file displayed.

INFO: *disk.volume*: disabled slice

Explanation

Slice *disk.volume* cannot be accessed any more.

INFO: *disk.volume*: reattached slice to volume *volume*

Explanation

Slice *disk.volume* is attached to *volume* again.

INFO: *disk: device*: ready for swapping

Explanation

disk is ready for swapping. When using a file unit device supporting active swap, you may remove the physical disk *device* at this point for swapping. If active swap is not supported, shut down the system and turn off the power before removing the physical disk.

INFO: *disk*: restored disk

Explanation

disk restored its original status.

INFO: *volume*: recovered volume with disk *disk*

Explanation

volume was recovered based on *disk*.

INFO: *volume*: recovered volume with lower level group *lgroup*

Explanation

Volume *volume* was recovered based on lower level group *lgroup*.

INFO: *volume* : recovered volume in group *group*

Explanation

Volume *volume* in group *group* has been recovered.

INFO: *disk*: recovered disk

Explanation

disk *disk* was recovered.

INFO: no unsynchronized slice in *volume*

Explanation

There is no slice that needs synchronization copying to be started, or resumed in *volume*.

INFO: *volume*: copying canceled

Explanation

Synchronization copying of *volume* has been cancelled.

INFO: no copying slice exists in *volume*

Explanation

There is no slice in the process of copying in *volume*.

INFO: *volume*: copying interrupted

Explanation

Synchronization copying on *volume* was canceled.

INFO: *class*: updated attribute value of class

Explanation

Attribute value of *class* was updated.

INFO: *disk*: updated attribute value of disk

Explanation

Attribute value of *disk* was updated.

INFO: *group*: updated attribute value of group

Explanation

Attribute value of *group* was updated.

INFO: *volume*: updated attribute value of volume

Explanation

Attribute value of *volume* was updated.

INFO: *slice*: updated attribute value of slice

Explanation

Attribute value of *slice* was updated.

INFO: updated /etc/vfstab file

Explanation

/etc/vfstab file has been updated.

INFO: updated /kernel/drv/sfdsk.conf file

Explanation

/kernel/drv/sfdsk.conf file has been updated.

INFO: updated /etc/system file

Explanation

/etc/system file has been updated.

INFO: completed definitions of system disk(s) mirroring

Explanation

Definition for system disk mirroring is complete.

INFO: system must be rebooted: shutdown -y -g0 -i6

Explanation

Reboot the system as follows.
shutdown -y -g0 -i6

INFO: completed definitions of system disk(s) mirroring cancellation

Explanation

Canceling definition for system disk mirroring is complete.

INFO: *disk.volume*: started slice on node *node*

/dev/sfdsk/class/rdsk/disk.volume

/dev/sfdsk/class/dsk/disk.volume

Explanation

Slice *disk.volume* was activated on *node*.

You may access slice by using the displayed special file.

INFO: *disk.volume*: stopped slice on node *node*

Explanation

Slice *disk.volume* was stopped on *node*.

You cannot access slice until it is activated again.

INFO: *disk.volume*: disabled slice on remote node *node*

Explanation

Slice *disk.volume* has become unavailable on remote node *node*.

INFO: *disk.volume*: current node *node* took over and started slice

/dev/sfdsk/class/rdsk/disk.volume

/dev/sfdsk/class/dsk/disk.volume

Explanation

Current node took over slice *disk.volume* and activated it.

You may access to slice by using the displayed special file.

INFO: *volume*: copying parameter modification completed

Explanation

Modification of parameter related to the copying process of *volume* has completed.

INFO: configuration parameters modification completed

Explanation

Modification of configuration parameter has completed

INFO: *proxy*: related and joined to master *master*

Explanation

Proxy *proxy* has been related and joined to master *master*.

INFO: *proxy*: created proxy volume

Explanation

Proxy volume *proxy* has been created.

INFO: *proxy*: broken relation with master

Explanation

Relationship between master and proxy *proxy* has been cancelled.

INFO: *proxy*: forced to break relation with master

Explanation

Relationship between master and proxy *proxy* was forcibly cancelled.

INFO: *proxy*: must be recovered by `sdxfix(1M)` command

Explanation

In order to restore data of proxy *proxy*, volume must be restored with `sdxfix(1M)` command.

INFO: *volume*: started volume on node *node*,...

`/dev/sfdsk/class/rdisk/volume`

`/dev/sfdsk/class/dsk/volume`

Explanation

Volume was started on node *node*,.... It can be accessed with the special file.

INFO: *proxy*: parted from master

Explanation

Proxy *proxy* was parted from master.

INFO: *proxy*: parted from master instantly

Explanation

Proxy *proxy* was immediately parted from master.

INFO: proxy: rejoined to master master

Explanation

Proxy *proxy* has been rejoined to master *master*.

INFO: master: restored data from proxy proxy

Explanation

Data of master *master* was restored using proxy *proxy*.

INFO: master: restored data from proxy proxy instantly

Explanation

Data of master *master* was immediately restored using proxy *proxy*.

INFO: swapped slice components between master and proxy

Explanation

Slice comprising proxy and slice comprising master have been swapped.

E.4.4 Fix Messages

TO FIX: Usage: usage

Resolution

Follow the *usage* syntax and execute the command.

E.4.5 Internal Error Messages

command: severity: module: internal error, details

Explanation

Internal error occurred. *details* gives you the cause of error, *command* the command name, *severity* the severity of message, and *module* the module name (usually the internal function name) where the event was detected.

Resolution

Collect investigation material and contact your local customer support.

E.5 Operation Management View Messages

This section describes the GDS Management View messages.

E.5.1 Error Messages

Executing the command was unsuccessful. Proceed?

Explanation

Executing the command was unsuccessful.

Resolution

To proceed with operation press <Yes>, to cancel press <No>.

The following exceptions were received: (details)

Explanation

Exceptions given in (details) occurred.

Resolution

Contact your local customer support.

Command line output: (details)

Explanation

Error given in (details) occurred when executing the command.

Resolution

Identify the cause referring to the command line output. When you cannot identify the cause, contact your local customer support.

Object name you typed is already defined. Change the object name.

Explanation

The object name you typed is already defined.

Resolution

Change the object name.

Abnormality of log monitoring mechanism occurred on node *node*. Monitoring task will be stopped. Check that the node is in operation, and update to the latest information.

Explanation

Abnormality of log monitoring mechanism such as a stopping of the node occurred.

Resolution

Check that the node is in operation, and update to the latest information.

Reading GDS configuration information was unsuccessful. Contact the system administrator.

Explanation

Reading GDS configuration information was unsuccessful.

Resolution

Contact your local customer support.

Invalid class name was specified. Type another name.

Explanation

When typing the class name, an invalid character was used.

Resolution

Press <OK> and type a group name again. For information on characters usable for group names, see "A.1.1 Object Name."

Invalid group name was specified. Type another name.

Explanation

When typing the group name, an invalid character was used.

Resolution

Press <OK> and type a group name again. For information on characters usable for group names, see "A.1.1 Object Name."

Invalid volume name was specified. Type another name.

Explanation

When typing the volume name, an invalid character was used.

Resolution

Press <OK> and type a volume name again. For information on characters usable for volume names, see "A.1.1 Object Name."

Invalid disk name was specified. Type another name.

Explanation

When typing the disk name, an invalid character was specified.

Resolution

Press <OK> and type a disk name again. For information on characters usable for disk names, see "A.1.1 Object Name."

File system type or mount point is not set select a file system type or type in mount point

Explanation

<Create> was pressed before specifying the file system type or mount point.

Resolution

Press <Create> after selecting a file system type and typing a mount point.

The mount point is already in use. Change the mount point.

Explanation

The mount point you typed is already in use with a different file system.

Resolution

Press <Create> after changing the mount point.

System disk cannot be deleted.

Explanation

In [System Disk Settings] screen, deletion of system disk was attempted.

Resolution

System disk cannot be deleted from group configuration disk.

Physical disk cannot be added Maximum of eight disks may be added to group configurations disks.

Explanation

In [System Disk Settings] screen, adding ninth disk to group configuration disks was attempted.

Resolution

A maximum of 8-way multiplex mirroring is possible with system disk. You must limit the group configuration disks to four.

There are no empty slices in the system disk. Secure an empty slice before setting the system disk.

Explanation

There is no empty slice on system disk that may be used for mirror settings.

Resolution

Secure empty slice on system disk by executing format (1M) command and such.
After securing an empty slice, set the system disk again.

A non-supported device exists in the configuration information on node XXX.

Confirm system environment.

Explanation

Start up of GDS operation management failure. Cause of failure has the GDS configuration information which is not supported by GDS operation management.

Resolution

Contact your local customer support.

There is an error in order of removal of the disk which constitutes a group or the group.

Remove according to the right order of removal.

Explanation

The disk or group which cannot be removed was selected.

Resolution

When you remove the disk which constitutes a group or the group, remove in order the disk or group added later.

An error is in specification of stripe width.

Specify a value of two raised to the power within the limits of an integer.

Explanation

In the input of stripe width, the value of the character which cannot be used, or the range which cannot be specified was specified..

Resolution

Click <OK>. Input stripe width again.

E.5.2 Warning Messages

Class name you typed is already defined. Register the disk with this class?

Explanation

The class name you typed is already defined.

Resolution

Press <Yes> to add the disk to the specified class. Press <No> to register the disk with a different class.

Group name you typed is already defined. Register the disk with this group?

Explanation

The group name you typed is already defined.

Resolution

Press <Yes> to add the disk to the specified group. Press <No> to connect the disk with a different group.

If you register a physical disk with class, all data contained in the disk will be lost.

Proceed?

Explanation

All disk data will be lost upon registering the physical disk with class.

Resolution

Press <Yes> to register the physical disk, otherwise <No>.

All volume data will be lost upon removing the volume. Proceed?

Explanation

All volume data will be lost upon removing the volume.

Resolution

Press <Yes> to remove the volume, otherwise <No>.

Volume will be recovered. Proceed?

Explanation

Volume that cannot be accessed will be recovered.

Resolution

Press <Yes> to recover the volume, otherwise <No>.

Selected disk is normal. Swap disks?

Explanation

A normal disk is selected for swapping.

Resolution

Press <Yes> to swap the disks, otherwise <No>.

In order to perform automatic configuration settings, at least two I/O adapters should be installed. The system is not configured to perform automatic configuration settings. Go to Class Configuration screen.

Explanation

With automatic configuration settings, disks connected to different I/O adapter will be mirrored. Automatic configuration settings cannot be performed when all the selected disks are connected to one I/O adapter.

Resolution

From the Settings menu, click the desired configuration settings (Class Configuration, Group Configuration, Volume Configuration) and proceed with operation.

All data contained in the selected disk will be lost. Proceed?

Explanation

Data of the disk you selected for automatic configuration will be lost.

Resolution

Press <Yes> to proceed, otherwise <No>.

Physical disk will be excluded from Global Disk Services. (Physical Disk: XXXXXXXXXXXX Affected Volume: XXXXXXXXXX) To swap the disks, related slice will stop its operation. Excluding the disk will affect data reliability.

Explanation

When swapping physical disks, related slice will be stopped until restore is completed. Volume reliability will also be affected.

Resolution

Press <Yes> to swap physical disks , otherwise <No>.

Physical disk will be incorporated into Global Disk Services. (Physical Disk: XXXXXXXXXXXX Affected Volume: XXXXXXXXXX) All data within physical disk that will be lost. Proceed?

Explanation

Data within physical disk will be lost by restoring the physical disk.

Resolution

Press <Yes> to restore the physical disk , otherwise <No>.

Slice will be detached from the volume. (Target Slice: XXXXXXXXXXXX Affected Volume: XXXXXXXXXX) Detaching the slice will affect data reliability. Proceed?

Explanation

Detaching the slice will affect reliability.

Resolution

Press <Yes> to detach the slice, otherwise <No>.

The detached slice will be attached to the volume. (Target Slice: XXXXXXXXXXXX Affected Volume: XXXXXXXXXX) All data within the slice will be lost. Proceed?

Explanation

Data contained in the slice will be overwritten by attaching slice.

Resolution

Press <Yes> to attach the slice, otherwise <No>.

Volume will be stopped. Proceed?

Explanation

Volume will be stopped.

Resolution

Press <Yes> to stop volume, otherwise <No>.

Removing all disks configuring a class will result in removing the class as well.

Proceed?

Explanation

Removing all disks configuring a class will result in removing the class as well.

Resolution

Press <Yes> to remove class, otherwise <No>.

Removing all disks configuring a group will result in removing the group as well.

Proceed?

Explanation

Removing all disks configuring a group will result in removing the group as well.

Resolution

Press <Yes> to remove group, otherwise <No>.

File system will be created within the selected volume. Proceed?

Explanation

A file system will be created within the selected volume. Proceeding with this operation will result in losing all data within the volume.

Resolution

Press <Yes> to create a file system, otherwise <No>.

Mount and unmount information of the selected volume will be removed and will not be available as a file system. Proceed?

Explanation

Mount and unmount information of the selected volume will be removed and will not be available as a file system.

Resolution

Press <Yes> to remove the file system, otherwise <No>.

Configuration change of GDS configuration information was detected. Update the configuration information?

Explanation

Commands other than those from GDS Management View changed the GDS configuration information.

Resolution

Press <Yes> to update the GDS Management View display, otherwise <No>.

Physical disk will be excluded from Global Disk Services. Proceed?

Explanation

In order to swap physical disks, unused disk or spare disk will be excluded. This will not affect the volume.

Resolution

Press <Yes> to swap the physical disks , otherwise <No>.

File system is configured within the selected volume. Remove the file system from File System Configuration.

Explanation

The volume cannot be removed due to a file system configured within the volume.

Resolution

Remove the file system from the [File System Configuration] screen and then remove the volume.

Physical disk has not been swapped. Confirm if swapping took place. Proceed with incorporating disk?

Explanation

Swapping of physical disk may not have been performed.

Resolution

Press <Yes> to incorporate physical disk, <No> to cancel.

Status of physical disks connected to system will be updated. Updating may take some time. Proceed?

Explanation

Status of physical disks connected to system will be updated. During update, you cannot operate GDS.

Resolution

Press <Yes> to update physical disk information, <No> to cancel.

There is a volume with a file system within the class. You cannot change attributes since it will result in changing the path. First, delete the file system, and then change the attributes.

Explanation

Name of the class containing a volume with a file system was changed.

Resolution

Changing the class name will result in changing the access path. First delete the file system.

Volume has a file system. You cannot change attributes since it will result in changing the path. First, delete the file system, and then change the attributes.

Explanation

Name of the volume with file system was changed.

Resolution

Changing the volume name will result in changing the access path. First delete the file system.

There are no disks available as system disk after unmirroring. Recover the abnormal volume before unmirroring the system disk.

Explanation

Unmirroring of system disk was attempted when one of the disks configuring system disk mirroring was abnormal.

Resolution

First swap the abnormal physical disk. After recovering the physical disk, unmirror the system disk.

Unmirroring the system disk will affect system reliability. Proceed with unmirroring?

Explanation

<Unmirror> was pressed in the Unmirror System Disk screen.

Resolution

Press <Yes> to unmirror system disk, <No> to cancel.

Since synchronization copying is in process, settings for system disk mirroring are not complete. After synchronization copying is complete, select [System Disk Settings] again.

Explanation

While setting the system disk, a volume in process of synchronization copying was detected.

Resolution

Press <OK> and close the message. From the GDS Management screen, monitor the root class volume in process of synchronization copying. Confirm that the volume status changes from "warnings" to "active," and try setting the system disk again.

Settings for system disk mirroring are not complete. You must rewrite the mount information. Proceed with system disk mirroring settings?

Explanation

While setting the system disk, a volume in process of synchronization copying was detected and system disk setting was interrupted.

Resolution

Press <Yes> and proceed with system disk mirroring settings.

Slice access will be inhibited. Proceed?

Explanation

Access to detached slice will be inhibited.

Resolution

Press <Yes> to inhibit access to slice, <No> to allow access.

Slice will become accessible. Proceed?

Explanation

Slice that has become inaccessible as a result of stopping slice will become accessible.

Resolution

Press <Yes> to allow access to slice, <No> to inhibit access.

Copying to slice will be stopped. You will not be able to resume copying. Proceed?

Explanation

Resynchronization copying of slice will be cancelled.
Copying cannot be resumed from the point where it stops.

Resolution

Press <Yes> to cancel resynchronization copying, <No> not to cancel.

Will start (resume) copying to slice. Proceed?

Explanation

Copying of slice will be performed to match the data of slices configuring volume.
Also, copying of the interrupted slice will resume.

Resolution

Press <Yes> to start (resume) resynchronization copying, <No> not to start(resume).

Remove group within "root" class?

Explanation

A group within "root" class is about to be removed.

Resolution

Press <Yes> to remove the group. Press <No> if you do not.

Abnormal event detected on *node*. Display *node* information?

Explanation

An object abnormality was detected on another node within the cluster system.

Resolution

Press <Yes> to display the information about node with an abnormality. Press <No> if you do not wish to display.

Abnormal event detected on *node*. After completing the current operation, switch to the node with the abnormality using the node selection function.

Explanation

An object abnormality was detected on another node within the cluster system.

Resolution

After completing the current operation, switch to the node with the abnormality by selecting [Select Node] on the [General] menu.

E.5.3 Information Messages

The slice is attached and the volume will be restored.

Explanation

The slice is attached. In order to restore the volume, synchronization copying will start.

Resolution

Press <OK> and close the message. Confirm the completion of synchronization copying by referring to the main screen volume information field.

Volume is restored.

Explanation

The volume is restored.

Resolution

Press <OK> and close the message.

Slice has been detached. The following access path is available for use.

Access path: /dev/??dsk/???????????

Explanation

Slice has been detached to make it accessible.

Resolution

Confirm the displayed access path and press <OK>.

Disk is excluded from Global Disk Services. Contact our engineer (CE) to swap physical disk.

Explanation

To swap the physical disk, it was excluded from GDS.

Resolution

Press <OK>, and contact your local customer support for swapping the physical disk.

Physical disk is incorporated into Global Disk Services. Volume is now online.

Explanation

The swapped physical disk is placed under GDS.

Resolution

Press <OK> and close the message.

Optimum disk pair cannot be found. Create a volume from Class Configuration screen.

Explanation

The number of disks or disk size was inappropriate and could not configure optimum mirroring environment with automatic configuration.

Resolution

From the [Settings] menu, click the desired configuration settings (Class Configuration, Group Configuration, Volume Configuration) and proceed with operation.

File system was created.

Explanation

A file system was created and mount information was added.

Resolution

Press <OK> and close the message.

File system was removed.

Explanation

The file system and mount information was removed.

Resolution

Press <OK> and close the message.

GDS configuration information is not defined. Create a volume from System Disk Configuration or Automatic Configuration or Class Configuration screen.

Explanation

GDS configuration information is not defined.

Resolution

Create a volume from either Class Configuration or Automatic Configuration screen.

Status monitoring will be updated every ?? minutes.

Explanation

Object monitoring intervals was changed.

Resolution

Press <OK> and close the message.

Exit Global Disk Services?

Explanation

Exits GDS.

Resolution

Press <Yes> to exit, <No> to continue operation.

The following files were updated to complete the system disk mirror settings.

/etc/system

/etc/vfstab

/kernel/drv/sfdsk.conf

Restart the system.

Explanation

In [System Disk Settings] screen, mirroring definition has been completed.

Resolution

Press <OK> and close the message. Then promptly reboot the system.

Unmirroring of system disk is complete. To validate the unmirroring, reboot the system.

Explanation

In [Unmirror System Disk] screen, unmirroring has been completed.

Resolution

Press <OK> and close the message. Then promptly reboot the system.

Slice access inhibited. Activate the slice in order to gain access.

Explanation

Access to detached slice has been inhibited.

Resolution

Press <OK> and close the message.

Slice is now accessible.

Explanation

Detached slice has become accessible.

Resolution

Press <OK> and close the message.

The selected slice is neither TEMP or TEMP-STOP status. You cannot perform any operations on the selected slice.

Explanation

[Stop/Activate Slice] operation has been performed on slices not in TEMP or TEMP-STOP status

Resolution

Press <OK> and close the message.

Cancelled copying to slice.

Explanation

Resynchronization copying of slice has been cancelled.

Resolution

Press <OK> and close the message.

Started (Resumed) copying to slice.

Explanation

Resynchronization copying has started (resumed).

Resolution

Press <OK> and close the message.

The selected slice is neither ACTIVE or STOP status. You cannot perform any operations on the selected volume.

Explanation

Copying operation has been performed when slices configuring volume is not in COPY status.

Resolution

Press <OK> and close the message.

The selected slice is neither ACTIVE or STOP status. You cannot perform any operations on the selected slice.

Explanation

[Detach Slice] operation has been performed on slices that are not in ACTIVE or STOP status.

Resolution

Press <OK> and close the message.

There is a slice in COPY status in the selected volume. You cannot perform operation on the selected volume.

Explanation

[Start Copying] operation has been performed on a volume with a slice in process of synchronization copying.

Resolution

Press <OK> and close the message.

There are no slices in INVALID status in the selected volume. You cannot perform any operations on the selected volume.

Explanation

[Start Copying] operation has been performed on a volume without a slice in INVALID status.

Resolution

Press <OK> and close the message.

The selected volume is not in COPY status. You cannot perform any operations on the selected volume.

Explanation

[Cancel Copying] operation has been performed on a volume without a slice in process of synchronization copying.

Resolution

Press <OK> and close the message.

Appendix F Troubleshooting

Since GDS automatically attempts recovery when various abnormalities are detected, users do not have to be aware of most of the abnormalities. However, depending on the type of abnormality, or when multiple abnormalities occur at the same time, automatic recovery may not be possible.

In this section, resolutions for system administrator, in the event of abnormalities with objects and physical disks managed by GDS, are explained.



Note

As a general rule, all operations explained here must be executed with superuser authority.

F.1 Resolving Problems

F.1.1 Slice Status Abnormality

(1) Mirror slice configuring the mirror volume is in INVALID status.

Explanation

You can check the status of the slice configuring the volume as follows.

# sdxinfo -S -o Volume1					
OBJ	CLASS	GROUP	DISK	VOLUME	STATUS
slice	Class1	Group1	Object1	Volume1	ACTIVE
slice	Class1	Group1	Object2	Volume1	INVALID

In this example, among the slices that exist in volume Volume1, the slice within Object2 is in INVALID status, as shown in the STATUS field. Object2 is a disk or lower level group connected to the highest level mirror group Group1.

The following five occasions can possibly cause the INVALID status of the mirror slice Volume1.Object2.

- An I/O error occurred on the mirror slice Volume1.Object2.
(Cause a)
A disk component relevant to Object2 failed to operate properly, and an I/O error occurred on the mirror slice Volume1.Object2.
(Cause b)
A component other than disks relevant to Object2 (such as an I/O adapter, I/O cable, I/O controller, power, and fan) failed to operate properly, and an I/O error occurred on the mirror slice Volume1.Object2.
- An I/O error occurred on the mirror slice Volume1.Object1 during synchronization copying to the mirror slice Volume1.Object2.
(Cause a')
A disk component relevant to Object1 failed to operate properly during synchronization copying to the mirror slice Volume1.Object2, and an I/O error occurred on the copy source slice Volume1.Object1.
(Cause b')
A component other than disks relevant to Object1 (such as an I/O adapter, I/O cable, I/O controller, power, and fan) failed to operate properly during synchronization copying to the mirror slice Volume1.Object2, and an I/O error occurred on the mirror slice Volume1.Object1 that is the copy source.

- Others

(Cause c)

Synchronization copying to the mirror slice Volume1.Object2 was cancelled as a result of a cause such as [Cancel Copying] selection from the GDS Management View, sdxcopy command execution, and a power outage.

Resolution

- 1) Identify the physical disk name of a faulty disk using the sdxinfo command.

(Example A1)

```
# sdxinfo -G -o Volume1
OBJ  NAME  CLASS  DISKS          BLKS    FREEBLKS  SPARE
-----
group Group1  Class1  Object1:Object2 17596416 17498112  0

# sdxinfo -D -o Volume1 -e long
OBJ  NAME  TYPE  CLASS  GROUP  DEVNAM  DEVBLKS  FREEBLKS  DEVCONNECT  STATUS  E
-----
disk  Object1  mirror  Class1  Group1  c1t1d0  17596416  *          node1        ENABLE  *
disk  Object2  mirror  Class1  Group1  c2t3d0  17596416  *          node1        ENABLE  *
```

In this example, Object2 is a disk connected with the highest level group Group1. As indicated in the E field, an I/O error occurred on the disk Object2, and the possible cause is (Cause a) or (Cause b). The physical disk name of the disk Object2 is c2t3d0 as shown in the DEVNAM field.

In example A1, if a value is 0 in the E field of the disk Object2 including a slice in the INVALID status and if a value is 1 in the E field of the disk Object1 that is mirrored with the disk Object2, it indicates an I/O error occurred on the disk Object1 and the possible cause is (Cause a') or (Cause b'). In such a case, see "(2) The copy destination slice was made INVALID due to an I/O error generated on the copy source slice during synchronization copying" and perform restoration.

If the E field of any disk does not contain the value 1 in example A1, the possible cause is (Cause c).

(Example B1)

```
# sdxinfo -G -o Volume1
OBJ  NAME  CLASS  DISKS          BLKS    FREEBLKS  SPARE
-----
group Group1  Class1  Object1:Object2 35127296 35028992  0
group Object1  Class1  Disk1:Disk2      35127296 *          *
group Object2  Class1  Disk3:Disk4      35127296 *          *

# sdxinfo -D -o Volume1 -e long
OBJ  NAME  TYPE  CLASS  GROUP  DEVNAM  DEVBLKS  FREEBLKS  DEVCONNECT  STATUS  E
-----
disk  Disk1  stripe  Class1  Object1  c1t1d0  17596416  *          node1        ENABLE  0
disk  Disk2  stripe  Class1  Object1  c1t2d0  17596416  *          node1        ENABLE  0
disk  Disk3  stripe  Class1  Object2  c2t3d0  17682084  *          node1        ENABLE  1
disk  Disk4  stripe  Class1  Object2  c2t4d0  17682084  *          node1        ENABLE  0
```

In this example, Object2 is a lower level group connected with the highest level group Group1. As indicated in the E field, an I/O error occurred on the disk Disk3 connected with Object2 and it is relevant to (Cause a) or (Cause b). The physical disk name corresponding to Disk3 is c2t3d0 as shown in the DEVNAM field.

In example B1, if a value is 0 in the E field of the disks (Disk3 and Disk4) connected with the lower level group Object2 including a slice in the INVALID status and if a value is 1 in the E field of the disk (Disk1 or Disk2) connected with the lower level group Object1 that is mirrored with the lower level group Object2, it indicates an I/O error occurred on the disk connected with Object1 and the possible cause is (Cause a') or (Cause b'). In such a case, see "(2) The copy destination slice was made INVALID due to an I/O error generated on the copy source slice during synchronization copying" and perform restoration.

In example B1, if there are no disks with "1" in the E field, INVALID status would be a result of (Cause c.)

2) Refer to disk driver log message and check the physical disk abnormalities.

Failure of the disk may not be limited to that of the disk itself, but also includes failure or a defect of other components such as I/O adapter, I/O cable, I/O controller, power, and fan.

Contact your local customer support and specify which component failed, or might be defective.

If there are no failures or defective components, INVALID status will be a result of (Cause c.) If it is a result of (Cause a,) follow procedures 3a to 5a, and 6. If it is a result of (Cause b,) follow procedures 3b and 6. If it is a result of (Cause c,) follow procedures 3c and 6.

3a) If it is a result of (Cause a,) you must perform the following operations, before and after the disk swap. For procedures on swapping disks from Operation Management View, see " 5.3.3 In Case of Disk Abnormality."

Before swapping the disks, execute the following command.

(Example A3)

```
# sdxswap -O -c Class1 -d Object2
```

In this example, disk Object2 connected to the highest level group Group1 will be swapped.

(Example B3)

```
# sdxswap -O -c Class1 -d Disk3
```

In this example, disk Disk3 will be swapped. Disk3 is a disk connected to lower level group Object2, which is a lower level group of the highest level group Group1.

4a) Swap disks.

5a) After swapping disks, execute the following command.

(Example A5)

```
# sdxswap -I -c Class1 -d Object2
```

or

(Example B5)

```
# sdxswap -I -c Class1 -d Disk3
```

3b) If it is the result of (Cause b,) you must shut down the system. Recover the failed component and boot your system. Mirroring status will be recovered by automatic synchronization copying.

3c) If it is the result of (Cause c,) perform synchronization copying of mirror volume.

```
# sdxcopy -B -c Class1 -v Volume1
```

6) You can confirm the recovery of the slice configuring the volume, as shown below.

# sdxinfo -S -o Volume1						
OBJ	CLASS	GROUP	DISK	VOLUME	STATUS	
slice	Class1	Group1	Object1	Volume1	ACTIVE	
slice	Class1	Group1	Object2	Volume1	ACTIVE	

In this example, slice within Object1 and Object2 are both in ACTIVE status. This indicates that the recovery process has completed successfully.

(2) The copy destination slice was made INVALID due to an I/O error generated on the copy source slice during synchronization copying.

Explanation

When an I/O error occurs on the copy source slice during synchronization copying, the copy destination slice becomes INVALID while the source slice is still ACTIVE.

The following two occasions are possible causes.

(Cause a')

A disk component of the copy source failed to operate properly, and an I/O error occurred on the copy source slice.

(Cause b')

A component other than disks of the copy source (such as an I/O adapter, I/O cable, I/O controller, power, and fan) failed to operate properly during synchronization copying, and an I/O error occurred on the copy source slice.

For details on determining whether the status is relevant to one of these events and identifying the physical disk name of a faulty disk, see [Explanation] and procedure 1) of [Resolution] described in "(1) Mirror slice configuring the mirror volume is in INVALID status."

Resolution

First examine the physical disk abnormalities referring to disk driver log messages and so on. Then contact your local customer support and locate the disabled or faulty part.

When the possible cause is (Cause b'), shut down the system once, repair the disabled part, and boot the system. Consequently, synchronization copying is performed and the mirroring status is restored.

When the possible cause is (Cause a'), take the following procedures for restoration. The procedures are illustrated for each of the following three cases.

- A. In the case of /(root), /usr, or /var
- B. In the case of the swap area
- C. In the case of others (other than /(root), /usr, /var, and swap)

The following illustrates restoration procedures when the class name is Class1, the volume name is Volume1, the name of a faulty disk of the copy source is Disk1, and the name of a disk of the copy destination is Disk2 as examples.

```

# sdxinfo -S -o Volume1
OBJ CLASS GROUP DISK VOLUME STATUS
-----
slice Class1 Group1 Disk1 Volume1 ACTIVE
slice Class1 Group1 Disk2 Volume1 INVALID

# sdxinfo -G -o Volume1
OBJ NAME CLASS DISKS BLKS FREEBLKS SPARE
-----
group Group1 Class1 Disk1:Disk2 17596416 17498112 0

# sdxinfo -D -o Volume1 -e long
OBJ NAME TYPE CLASS GROUP DEVNAM DEVBLKS FREEBLKS DEVCONNECT STATUS E
-----
disk Disk1 mirror Class1 Group1 c1t1d0 17596416 * node1 ENABLE 1
disk Disk2 mirror Class1 Group1 c2t3d0 17596416 * node1 ENABLE 0

```

A. In the case of /(root), /usr, or /var
 Contact your local customer support.

B. In the case of the swap area

B.1) Remove the volume from the swap area.

```
# swap -d /dev/sfdsk/Class1/dsk/Volume1
```

B.2) Inactivate the volume.

```
# sdxvolume -F -c Class1 -v Volume1
```

B.3) Restore the status of the copy destination slice in the INVALID status.

```
# sdxfix -V -c Class1 -d Disk2 -v Volume1
```

B.4) Verify that the restored copy destination slice is in the STOP status and the copy source slice is in the INVALID status now.

```

# sdxinfo -S -o Volume1
OBJ CLASS GROUP DISK VOLUME STATUS
-----
slice Class1 Group1 Disk1 Volume1 INVALID
slice Class1 Group1 Disk2 Volume1 STOP

```

B.5) Activate the volume.

```
# sdxvolume -N -c Class1 -v Volume1
```

B.6) Add the volume to the swap area again.

```
# swap -a /dev/sfdsk/Class1/dsk/Volume1
```

B.7) Remove the faulty copy source disk from GDS management to make it a replaceable status.

```
# sdxswap -O -c Class1 -d Disk1
```

B.8) Swap the faulty copy source disk.

B.9) Put the swapped disk back in control of GDS management to make it available.

```
# sdxswap -I -c Class1 -d Disk1
```

C. In the case of others (other than /(root), /usr, /var, and swap)

C.1) Exit applications using the volume.

C.2) Unmount the file system on the volume when it has been mounted.

```
# umount /dev/sfdsk/Class1/dsk/Volume1
```

If the umount(1M) command fails due to an I/O error, specify the -f option and re-execute the command.

C.3) Inactivate the volume.

```
# sdxvolume -F -c Class1 -v Volume1
```

C.4) Restore the status of the copy destination slice in the INVALID status.

```
# sdxfix -V -c Class1 -d Disk2 -v Volume1
```

C.5) Verify that the restored copy destination slice is in the STOP status and the copy source slice is the INVALID status now.

```
# sdxinfo -S -o Volume1
OBJ CLASS GROUP DISK VOLUME STATUS
-----
slice Class1 Group1 Disk1 Volume1 INVALID
slice Class1 Group1 Disk2 Volume1 STOP
```

C.6) Activate the volume.

```
# sdxvolume -N -c Class1 -v Volume1
```

C.7) The consistency of volume data may have lost. Restore the backup data or perform repair using the fsck(1M) command if necessary.



Note

If the I/O error occurred on the copy source slice during resynchronization copying after the system down, restoration may possibly be performed with the fsck(1M) command.

C.8) Remove the faulty copy source disk from GDS management to make it a replaceable status.

```
# sdxswap -O -c Class1 -d Disk1
```

C.9) Swap the faulty copy source.

C.10) Put the swapped disk back in control of GDS management to make it available.

```
# sdxswap -I -c Class1 -d Disk1
```

(3) Slice configuring the volume is in TEMP status.

Explanation

Slice is not attached after it has been detached with the sdxslice command. Or else, you have not performed [Attach Slice] after performing [Detach Slice] from Operation Management View.

Resolution

Attach the slice again with the sdxslice command, or perform [Attach Slice] from Operation Management View as necessary.

(4) Slice configuring volume is in TEMP-STOP status.

Explanation

Slice is not activated after it has been stopped with the `sdxslice` command, or the detached node is not current node.

Or else, you have not performed [Stop Slice] after performing [Activate Slice] from Operation Management View.

Resolution

Activate slice or take over slice with the `sdxslice` command as needed. Or, perform [Activate Slice] from Operation Management View.

(5) Slice configuring the volume is in COPY status.

Explanation

In order to attach slice, synchronization copying is in process.

Or, synchronization copying is in process between master and proxy.

Resolution

Wait until synchronization copying is complete. Note that a slice in process of synchronization copying will not restrict you from accessing an active volume.

(6) Slice configuring the volume is in NOUSE status.

Explanation

When the status of disk related to slice is either in DISABLE or SWAP status, the slice becomes NOUSE to inhibit slice operation.

Resolution

Recover disk in DISABLE or SWAP status. For details, see "F.1.2 Disk Status Abnormality".

F.1.2 Disk Status Abnormality

(1) Disk is in DISABLE status.

Explanation

Disk ID information is checked when booting the system. If the ID information is regarded invalid at the time, the disk will become DISABLE.

There are three reasons that may cause this DISABLE status.

(Cause a)

You changed the I/O cable connection and booted the system.

(Cause b)

You swapped the disk and booted the system.

(Cause c)

Disk failure occurred.

Resolutions are described for each of the above causes.

Resolution

a) When it is due to (Cause a), shut down the system, connect the I/O cables correctly and boot again.

b) When it is due to (Cause b), shut down the system, return the original disk and boot again. If swapping the disks is necessary, follow the procedures and swap the disks. For procedures on disk swapping, see "D.8 sdxswap - Swap disk", or "5.3.3 In Case of Disk Abnormality".

c) When it is due to (Cause c), follow the procedures and swap the disks.

(2) Disk is in SWAP status.

Explanation

If the sdxswap -O command is executed, or [Swap Physical Disk] is performed from Operation Management View, disk will be in SWAP status.

Resolution

Complete the disk swap and restore the disk by either executing the sdxswap -I command, or by performing [Restore Physical Disk] from Operation Management View.

F.1.3 Volume Status Abnormality

(1) Mirror volume is in INVALID status.

Explanation

You can confirm the status of the volume as shown below.

# sdxinfo -V -o Volume1									
OBJ	NAME	CLASS	GROUP	SKIP	JRM	1STBLK	LASTBLK	BLOCKS	STATUS
volume	*	Class1	Group1	*	*	0	65535	65536	PRIVATE
volume	Volume1	Class1	Group1	off	on	65536	17596415	17530880	INVALID

In this example, volume Volume1 that exists in the highest level group Group1 is in INVALID status, as shown in the STATUS field.

If none of the mirror slices consisting the mirror volume contains valid data (ACTIVE or STOP), the mirror volume becomes INVALID. You cannot start a volume in INVALID status.

There are two reasons that may cause this INVALID status.

(Cause a) Disk is in DISABLE status.

(Cause b) Master-proxy relationship was cancelled forcibly while master data was being copied to proxy.

Resolution

1) Confirm that there is a disk in DISABLE status within the group with which the volume is associated as follows.

(Example A1)

# sdxinfo -G -o Volume1									
OBJ	NAME	CLASS	DISKS	BLKS	FREEBLKS	SPARE			
group	Group1	Class1	Disk1:Disk2	17596416	0	0			

# sdxinfo -D -o Volume1									
OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	DEVCONNECT	STATUS	
disk	Disk1	mirror	Class1	Group1	c1t1d0	17596416	node1	ENABLE	
disk	Disk2	mirror	Class1	Group1	c2t3d0	17596416	node1	DISABLE	

In this example, disks Disk1 and Disk2 are connected to the highest level mirror group Group1, and Disk2 is in DISABLE status as shown in the STATUS field.

(Example B1)

# sdxinfo -G -o Volume1							
OBJ	NAME	CLASS	DISKS	BLKS	FREEBLKS	SPARE	
group	Group1	Class1	Group2:Group3	35127296	17530880	0	
group	Group2	Class1	Disk1:Disk2	35127296	*	0	
group	Group3	Class1	Disk3:Disk4	35127296	*	0	

# sdxinfo -D -o Volume1								
OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	DEVCONNECT	STATUS
disk	Disk1	stripe	Class1	Group2	c1t1d0	17596416	node1	ENABLE
disk	Disk2	stripe	Class1	Group2	c1t2d0	17596416	node1	DISABLE
disk	Disk3	stripe	Class1	Group3	c2t3d0	17682084	node1	ENABLE
disk	Disk4	stripe	Class1	Group3	c2t4d0	17682084	node1	ENABLE

In this example, lower level stripe groups, Group2 and Group3 are connected to the highest level mirror group Group1. Disk Disk2 which is connected to Group2 is in DISABLE status as shown in the STATUS field.

2) When it is due to (Cause a), restore the disk by following the procedures given in section "F.1.2 Disk Status Abnormality."

3) From the disks and lower level groups connected to the highest level mirror group, determine the disk or lower level group to which the slice you will use to recover data belongs. Then, execute the sdxfix command to recover data.

(Example A3)

```
# sdxfix -V -c Class1 -d Disk1 -v Volume1
```

In this example, Volume1 is recovered after a slice in disk Disk1.

(Example B3)

```
# sdxfix -V -c Class1 -g Group3 -v Volume1
```

In this example, Volume1 is recovered after a slice in lower level stripe group Group3.

4) Start the volume.

```
# sdxvolume -N -c Class1 -v Volume1 -e nosync
```

5) Access Volume1 and check its contents. Restore backup data or run fsck to regain data integrity as necessary.

6) Perform synchronization copying on volume.

```
# sdxcopy -B -c Class1 -v Volume1
```


(2) Single volume is in INVALID status.

Explanation

You can confirm the status of the volume as shown below.

# sdxinfo -V -o Volume1									
OBJ	NAME	CLASS	GROUP	SKIP	JRM	1STBLK	LASTBLK	BLOCKS	STATUS
volume	*	Class1	Disk1	*	*	0	32767	32768	PRIVATE
volume	Volume1	Class1	Disk1	off	on	32768	65535	32768	INVALID
volume	*	Class1	Disk1	*	*	65536	8421375	8355840	FREE

In this example, single volume Volume1 that exists in single disk Disk1 is in INVALID status, as shown in the STATUS field.

You cannot start a volume in INVALID status.

There are two reasons that may cause this INVALID status.

(Cause a) Single disk is in DISABLE status. In this case, the single slice becomes NOUSE status.

(Cause b) Master-proxy relationship was cancelled forcibly while master data was being copied to proxy.

Resolution

1) Confirm that the single disk is in DISABLE status as shown below.

# sdxinfo -D -o Volume1								
OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	DEVCONNECT	STATUS
disk	Disk1	single	Class1	*	c1t11d0	8355840	node1	DISABLE

In this example, single disk Disk1 is in DISABLE status, as shown in the STATUS field.

2) When it is due to (Cause a), restore the disk by following the procedures given in section "F.1.2 Disk Status Abnormality."

3) Execute the sdxfix command to recover the single volume's data.

```
# sdxfix -V -c Class1 -d Disk1 -v Volume1
```

4) Activate the volume.

```
# sdxvolume -N -c Class1 -v Volume1
```

5) Access Volume1 and check its content. Restore backup data or run the fsck command to regain data integrity as necessary.

(3) Stripe volume or volume in concatenation group is in INVALID status.

Explanation

You can confirm the status of the volume as shown below.

# sdxinfo -V -o Volume1									
OBJ	NAME	CLASS	GROUP	SKIP	JRM	1STBLK	LASTBLK	BLOCKS	STATUS
volume	*	Class1	Group1	*	*	0	65535	65536	PRIVATE
volume	Volume1	Class1	Group1	off	on	65536	17596415	17530880	INVALID

In this example, volume Volume1 that exists in the highest level group Group1 is in INVALID status, as shown in the STATUS field.

If any of the disks related to volume is in DISABLE status, the slice consisting that volume becomes NOUSE status, and the volume becomes INVALID. You cannot start a volume in INVALID status.

Resolution

1) You can confirm the status of the disk related to volume as shown below.

# sdxinfo -G -o Volume1 -e long									
OBJ	NAME	CLASS	DISKS	BLKS	FREEBLKS	SPARE	MASTER	TYPE	WIDTH
group	Group1	Class1	Group2:Group3	70189056	65961984	*	*	stripe	32
group	Group2	Class1	Disk1:Disk2	35127296	*	*	*	concat	*
group	Group3	Class1	Disk3:Disk4	35127296	*	*	*	concat	*

# sdxinfo -D -o Volume1								
OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	DEVCONNECT	STATUS
disk	Disk1	concat	Class1	Group2	c1t1d0	17596416	node1	ENABLE
disk	Disk2	concat	Class1	Group2	c1t2d0	17596416	node1	DISABLE
disk	Disk3	concat	Class1	Group3	c2t3d0	17682084	node1	ENABLE
disk	Disk4	concat	Class1	Group3	c2t4d0	17682084	node1	ENABLE

In this example, lower level concatenation groups Group2 and Group3 are connected to the highest level stripe group Group1, and Disk2 connected to Group2 is in DISABLE status as shown in STATUS field.

2) Follow the procedures in “F.1.2 Disk Status Abnormality” and restore the disk status.

3) Execute the sdxfix command to recover the volume’s data. With -g option, indicate the highest level group name (in this example, Group1).

```
# sdxfix -V -c Class1 -g Group1 -v Volume1
```

4) Activate the volume.

```
# sdxvolume -N -c Class1 -v Volume1
```

5) Access Volume1 and check its content. Restore backup data or run the fsck command to regain data integrity as necessary.

(4) Master volume is in INVALID status.

Explanation

If the copying process fails while copying data from proxy volume to master volume because of I/O error or such, the status of master volume to which the data is being copied becomes INVALID.

Resolution

- 1) Check if there is a DISABLE status disk in the group to which the volume belongs with the following command.

# sdxinfo -D -o Volume1								
OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	DEVCONNECT	STATUS
disk	Disk1	mirror	Class1	Group1	c1t1d0	8421376	*	ENABLE
disk	Disk2	mirror	Class1	Group1	c1t2d0	8421376	*	DISABLE

In this example, Disk2 is in DISABLE status.

If there is a disk in DISABLE status, see section "Disk is in DISABLE status" in "F.1.2 Disk Status Abnormality," and check which of the causes listed in that section applies. If it is due to (Cause a) or (Cause b), follow the procedures and restore the disk.

- 2) Follow the procedures given in section "Slice configuring the volume is in INVALID status" in "F.1.1 Slice Status Abnormality," and check if there is a disk hardware abnormality. If there is, identify the faulty part. When the abnormality is caused by a failed or defective non-disk component, repair the faulty part.

- 3) Procedures to restore the data for different scenarios are given below.

Non-disk component failure

- When recovering data using the proxy volume:
-> Follow procedures (a) to restore.
- When recovering data using backup data on media such as tapes:
-> Follow procedures (b) to restore.

When caused by a disk component failure:

- When the disk does not belong to master group:
 - some disks connected to the group have failed
-> Follow procedures (c) to restore.
 - all disks connected to the group have failed
-> Follow procedures (d) to restore.
- When the disk belongs to master group:
-> Follow procedures (e) to restore.

a) Procedures to recover master volume data using proxy volume.

a1) In order to check if the proxy volume that will be used to recover data is separated from the master volume, execute the sdxinfo -V -e long command, and check the PROXY field.

a2) If the proxy volume is not separated, execute the following command.

```
# sdxproxy Part -c Class1 -p Volume2
```

a3) Exit all applications accessing the proxy volume. When using the proxy volume as a file system, execute unmount. When using the proxy volume as a file system, execute unmount.

a4) If the proxy volume is started, execute the following command.

```
# sdxvolume -F -c Class1 -v Volume2
```

a5) Recover master volume data using the proxy volume's data.

```
# sdxproxy RejoinRestore -c Class1 -p Volume2
```

b) Procedures to recover data using backup data.

b1) When the volume is in INVALID status, you must first change it to STOP status. Decide on the disk (slice) you wish to use to recover data, and execute the `sdxfix` command.

```
# sdxfix -V -c Class1 -d Disk1 -v Volume1
```

In this example, Volume 1 is restored after a slice in Disk 1.

b2) When the volume to be restored is stopped, start it with the following command.

```
# sdxvolume -N -c Class1 -v Volume1 -e nosync
```

b3) Access the volume to be restored and check its contents. Restore backup data or run `fsck` to regain data integrity as necessary.

b4) When mirroring is configured with the volume, perform synchronization copying.

```
# sdxcopy -B -c Class1 -v Volume1
```

c) Procedures to swap some disks connected to the group.

c1) If you restore the INVALID master volume later using data of a proxy volume related to the master volume, or use data of proxy volumes related to the master volume after restoring it, part the proxy volumes using the `sdxproxy Part` command.

```
# sdxproxy Part -c Class1 -p Volume2
```

c2) When there is a volume in INVALID status in the group, change it to STOP status with the `sdxfix -V` command. `-d` option indicates the disk without abnormality.

```
# sdxfix -V -c Class1 -d Disk1 -v Volume1
```

c3) Follow the procedures and swap the disks. For details, see the "D.8 sdxswap - Swap disk", or section "5.3.3 In Case of Disk Abnormality."

c4) Recover the master volume data. If data will be recovered using the proxy volume, follow procedures described in (a). If data will be recovered using backup data on media such as tapes, follow procedures described in (b).

d) Procedures to swap all disks connected to the group.

d1) Exit all applications accessing the master volume and the proxy volume that will be used to recover data. When using the proxy volume or the master volume as a file system, execute `umount`.

d2) Stop the master volume and proxy volume in (d1).

```
# sdxvolume -F -c Class1 -v Volume1  
# sdxvolume -F -c Class1 -v Volume2
```

d3) Execute the `sdxproxy RejoinRestore` command and restore the master volume data using proxy volume in (d1). If the command terminates normally and the master volume is not in INVALID status, restoration process is complete, and you do not need to perform procedures (d4) and after.

```
# sdxproxy RejoinRestore -c Class1 -p Volume2
```

d4) Execute the `sdxproxy Swap` command and swap slices of master volume with proxy volume in (d1).

```
# sdxproxy Swap -c Class1 -p Volume2
```

d5) By performing procedure (d4), the status of master volume will not be in INVALID status, and the status of the proxy volume becomes INVALID. Follow the procedures given in section "Proxy volume is in INVALID status" in "F.1.3 Volume Status Abnormality," and restore the proxy volume in INVALID status.

d6) Execute the `sdxproxy Swap` command and swap the slices of master volume and proxy volume you swapped in procedure (d4).

```
# sdxproxy Swap -c Class1 -p Volume2
```

e) Procedures to swap disks connected to the master group.

e1) Exit all applications accessing the master group, and volumes in the proxy group that will be used to recover data. When using the volume as a file system, execute umount.

e2) Stop all volumes in the master group and proxy group in (e1).

```
# sdxvolume -F -c Class1 -v Volume1
# sdxvolume -F -c Class1 -v Volume2
```

e3) Execute the sdxproxy RejoinRestore command and restore the master group data using proxy group in (e1). If the command terminates normally and all master volumes are not in INVALID status, restoration process is complete, and you do not need to perform procedures (e4) and after.

```
# sdxproxy RejoinRestore -c Class1 -p Volume2
```

e4) Execute the sdxproxy Swap command and swap slices of master group and proxy group in (e1).

```
# sdxproxy Swap -c Class1 -p Group2
```

e5) By performing procedure (e4), master volume will not be in INVALID status and the status of the proxy volume becomes INVALID. Follow the procedures given in section "Proxy volume is in INVALID status" in "F.1.3 Volume Status Abnormality," and restore the proxy volume in INVALID status.

e6) Execute the sdxproxy Swap command and swap the slices of master group and proxy group you swapped in procedure (e4).

```
# sdxproxy Swap -c Class1 -p Group2
```

(5) Proxy volume is in INVALID status.

Explanation

If the copying process fails while copying data from master volume to proxy volume because of I/O error or such, the status of the proxy volume to which the data is being copied becomes INVALID.

Resolution

1) Check if there is a DISABLE status disk in the group to which the volume belongs with the following command.

```
# sdxinfo -D -o Volume1
OBJ  NAME  TYPE  CLASS  GROUP  DEVNAM  DEVBLKS  DEVCONNECT  STATUS
----  ----  ----  ----  ----  ----  ----  ----  ----
disk  Disk1  mirror  Class1  Group1  c1t1d0  8421376  *  ENABLE
disk  Disk2  mirror  Class1  Group1  c1t2d0  8421376  *  DISABLE
```

In this example, Disk2 is in DISABLE status.

If there is a disk in DISABLE status, see section "Disk is in DISABLE status" in "F.1.2 Disk Status Abnormality," and check which of the cause (There are three causes a, b and c listed.) listed in that section applies. If it is due to (Cause a) or (Cause b), follow the procedures and restore the disk.

2) Follow the procedures given in section "Slice configuring the volume is in INVALID status" in "F.1.1 Slice Status Abnormality," and check if there is a disk hardware abnormality. If there is, identify the faulty part. When the abnormality was caused by a failed or defective non-disk component, repair the faulty part.

3) Procedures to restore the data for different scenarios are given below.

- When caused by a non-disk component failure:
 - > Follow procedures (a) to restore.
- When caused by a disk component failure:
 - disk does not belong to proxy group
 - some disks connected to the group have failed
 - > Follow procedures (b) to restore.
 - When all disks connected to the group have failed:
 - > Follow procedures (c) to restore.
 - When the disk belongs to proxy group:
 - > Follow procedures (d) to restore.

a) Procedures to recover proxy volume data using master volume.

a1) In order to check if the proxy volume is separated from the master volume, execute the `sdxinfo -V -e long` command, and check the PROXY field.

a2) If the proxy volume is not separated, execute the following command.

```
# sdxproxy Part -c Class1 -p Volume2
```

a3) Rejoin proxy volume with master volume.

```
# sdxproxy Rejoin -c Class1 -p Volume2
```

b) Procedures to swap some disks connected to the group.

b1) Cancel the relationship with master volume using the `sdxproxy Break` command.

```
# sdxproxy Break -c Class1 -p Volume2
```

b2) Separate the volumes that are in INVALID status in the group with the `sdxfix -V` command, and change them to STOP status. `-d` option indicates the disk without abnormality.

```
# sdxfix -V -c Class1 -d Disk1 -v Volume2
```

b3) Follow the procedures and swap the disks. For details, see the "D.8 `sdxswap` - Swap disk", or section "5.3.3 In Case of Disk Abnormality."

b4) Join the master and proxy again with the `sdxproxy Join` command.

```
# sdxproxy Join -c Class1 -m Volume1 -p Volume2
```

c) Procedures to swap all disks connected to the group.

c1) Cancel the relationship with master using the `sdxproxy Break` command.

```
# sdxproxy Break -c Class1 -p Volume2
```

c2) Exit all applications accessing the volume in the group. When using the volume as a file system, execute `umount`.

c3) Stop all volumes in the group.

```
# sdxvolume -F -c Class1 -v Volume2
```

c4) Check the volume configuration of the group (volume name, size, etc.) with the `sdxinfo` command, and keep a note of it.

c5) Remove all volumes in the group.

```
# sdxvolume -R -c Class1 -v Volume2
```

c6) Follow the procedures and swap the disks. For details, see the "D.8 `sdxswap` - Swap disk", or section "5.3.3 In Case of Disk Abnormality."

c7) Create the volume that you removed in procedure (c5) again.

```
# sdxvolume -M -c Class1 -g Group2 -v Volume2 -s size
```

c8) Stop the volume you created in procedure (c7).

```
# sdxvolume -F -c Class1 -v Volume2
```

c9) Join the master volume and the proxy volume again, with the sdxproxy Join command.

```
# sdxproxy Join -c Class1 -m Volume1 -p Volume2
```

d) Procedures to swap disks connected to the proxy group.

d1) Cancel the relationship with master using the sdxproxy Break command.

```
# sdxproxy Break -c Class1 -p Group2
```

d2) Exit all applications accessing the volume in the group. When using the volume as a file system, execute unmount.

d3) Stop all volumes in the group.

```
# sdxvolume -F -c Class1 -v Volume2
```

d4) Remove all volumes in the group.

```
# sdxvolume -R -c Class1 -v Volume2
```

d5) Follow the procedures and swap the disks. For details, see the "D.8 sdxswap - Swap disk", or section "5.3.3 In Case of Disk Abnormality."

d6) Join the master group and the proxy group again, with the sdxproxy Join command.

```
# sdxproxy Join -c Class1 -m Group1 -p Group2 -a Volume1=Volume2:on
```

(6) Volume is in STOP status.

Explanation

Usually, volume is automatically started when the system is booted and becomes ACTIVE. The status will change to STOP, when you stop the volume with the sdxvolume -F command, or perform [Stop Volume] from Operation Management View. Accessing a volume in STOP status will result in I/O error.

Resolution

Start the volume with the sdxvolume -N command, or perform [Start Volume] from Operation Management View as necessary.

(7) I/O error occurs although mirror volume is in ACTIVE status.

Explanation

A mirror volume consists of multiple slices, and in an event of an I/O error, the crashed slice will be detached. Therefore, accessing the volume will complete normally.

However, when an I/O error occurs when only one slice is ACTIVE amongst those configuring the volume, accessing the volume will result in an error. At such time, the status of the slice and the volume remains ACTIVE.

Probable situations resulting in such a problem will be described using a two-way multiplex mirroring configuration, where two disks or two lower level groups are connected to a group. As an example, means to circumvent such problems will also be described.

(Situation 1)

One of the slices was detached with the sdxslice -M in order to backup volume data. While accessing the volume, I/O error occurred with the other slice.

(Prevention 1)

Before executing the sdxslice -M command, connect a reserved disk and temporarily configure a three-way multiplex mirroring, or make the mirrored volume available for backup.

(Situation 2)

While restoring a slice with an I/O error, I/O error also occurred on another slice.

(Prevention 2)

By securing a spare disk within the class, effects due to delay in restoring the slice will be avoided to a certain degree

Resolution

1) Identify the cause of I/O error occurrence in the last ACTIVE slice, by referring to the disk driver log message.

Resolutions are described below assuming the following three circumstances:

- a) Error occurred due to a disk component failure. Will attempt recovery using backup data.
- b) Error occurred due to a disk component failure. Will attempt data recovery from a slice in INVALID status.
- c) Error occurred due to a failed or a defective non-disk component failure.

2a) When the error was caused by a disk component failure, no slice with valid data exists. Restore data from the backup data following the procedures given below.

3a) Exit the application accessing the volume. When the volume is used as a file system, execute the unmount command.
When I/O error occurs on the unmount command, execute -f option of the unmount command.

4a) Stop the volume with the sdxvolume command.

```
# sdxvolume -F -c Class1 -v Volume1
```

5a) If there is a TEMP status slice within the volume, attempt recovery following the procedures given in "F.1.1 Slice Status Abnormality".

6a) If there is a NOUSE status slice within the volume, attempt recovery following the procedures given in "F.1.1 Slice Status Abnormality".

7a) Record the volume size which can be checked as follows.

```
# sdxinfo -V -o Volume1
```

OBJ	NAME	CLASS	GROUP	SKIP	JRM	1STBLK	LASTBLK	BLOCKS	STATUS
volume	*	Class1	Group1	*	*	0	32767	32768	PRIVATE
volume	Volume1	Class1	Group1	off	on	32768	4161535	4128768	STOP

In this example, the volume size would be 4128768 blocks given in Volume1 BLOCKS field.

8a) Remove the volume with the sdxvolume command.

```
# sdxvolume -R -c Class1 -v Volume1
```

9a) Swap disks following the procedures given in "5.3.3 In Case of Disk Abnormality," or "D.8 sdxswap - Swap disk."

10a) Create a volume with the sdxvolume command again. For the *number_of_sectors*, use the size recorded in (7a), in this example, 4128768.

```
# sdxvolume -M -c Class1 -g Group1 -v Volume1 -s number_of_sectors
```

11a) Finally, restore the backup data to Volume1.

2b) When the error was caused by a disk failure, and no backup data exists, or even if it did, if the data is too old, restore data from the detached INVALID status slice, following the procedures given below.

3b) Exit the application accessing the volume. When the volume is used as a file system, execute the unmount command.
When I/O error occurs on the unmount command, execute `-f` option of the unmount command.

4b) Stop the volume with the `sdxvolume` command.

```
# sdxvolume -F -c Class1 -v Volume1
```

5b) If there is a TEMP status slice within the volume, attempt recovery following the procedures given in "F.1.1 Slice Status Abnormality".

6b) If there is a NOUSE status slice within the volume, attempt recovery following the procedures given in "F.1.1 Slice Status Abnormality".

7b) Determine the original mirror slice after the volume is recovered. Then, execute the `sdxfix` command.

(Example 1)

```
# sdxfix -V -c Class1 -d Disk2 -v Volume1
```

In this example, data is recovered after a mirror slice in disk Disk2 which is connected to the highest level mirror group.

(Example 2)

```
# sdxfix -V -c Class1 -g Group2 -v Volume1
```

In this example, data is recovered after a mirror slice in lower level group Group2 which is connected to the highest level mirror group.

8b) Start the volume.

```
# sdxvolume -N -c Class1 -v Volume1 -e nosync
```

9b) Create backup of Volume1 and regain data integrity by running `fsck` as necessary.

10b) Lastly, swap disks following the procedures given in "5.3.3 In Case of Disk Abnormality," or "D.8 sdxswap - Swap disk."

2c) When the error was caused by a failed or defective non-disk component, slice with valid data exists within the disk. Shut down the system and recover the failed component and boot your system again.

Mirroring status will be recovered by automatic synchronization copying.

(8) An I/O error occurs on the single volume.

Explanation

Since single volume consists of only one slice, accessing the volume at the time of an I/O error will result in an error. However, the status of slice and volume will remain ACTIVE.

Resolution

1) Identify the cause of I/O error occurrence by referring to the disk driver log message. How to resolve the problem is described in two cases: (a) a disk component failure occurs, and you wish to restore the disk by using the backup data, (b) a non-disk component failure occurs.

2a) In the event of a disk component failure, there will be no slice with valid data. Follow the procedures below and restore the data using the backup data. In this example, Disk1 (c1t11d0) has a failure.

```
# sdxinfo -D -o Disk1
```

OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	DEVCONNECT	STATUS
disk	Disk1	single	Class1	*	c1t11d0	8493876	node1	ENABLE

3a) Search the volumes within the faulty disk using the sdxinfo command. And record the volume size,

```
# sdxinfo -V -o Disk1
```

OBJ	NAME	CLASS	GROUP	SKIP	JRM	1STBLK	LASTBLK	BLOCKS	STATUS
volume	*	Class1	Disk1	*	*	0	32767	32768	PRIVATE
volume	Volume1	Class1	Disk1	off	on	32768	65535	32768	ACTIVE
volume	Volume2	Class1	Disk1	off	on	65536	4194303	4128768	ACTIVE
volume	*	Class1	Disk1	*	*	4194304	8421375	4227072	FREE

In this example, Volume1 and Volume2 are within the faulty Disk1. The size of Volume1 would be 32,768 blocks as shown in the BLOCKS field. The size of Volume2 would be 4,128,768 blocks as shown in the BLOCKS field.

4a) Exit the application accessing the volume. When the volume is used as a file system, execute the unmount command. When I/O error occurs on the unmount command, execute -f option of the unmount command.

5a) Stop the volume with the sdxvolume command.

```
# sdxvolume -F -c Class1 -v Volume1,Volume2
```

6a) Remove the volumes with the sdxvolume command.

```
# sdxvolume -R -c Class1 -v Volume1
# sdxvolume -R -c Class1 -v Volume2
```

7a-1) Before swapping the disks, execute the following command.

```
# sdxswap -O -c Class1 -d Disk1
```

Caution: If the disk is the only remaining disk in the disk class, it will result in an error as shown below.

In such case, follow the procedures 7a-2, 8a-2 and 9a-2.

SDX:sdxswap: ERROR: Disk1: The last ENABLE disk in class cannot be swapped

8a-1) Swap the disks.

9a-1) After swapping the disks, execute the following command.

```
# sdxswap -I -c Class1 -d Disk1
```

7a-2) Before swapping the disks, execute the following command.

Caution: If no error occurs after performing 7a-1, you do not have to perform procedures in 7a-2, 8a-2 and 9a-2.

```
# sdxdisk -R -c Class1 -d Disk1
```

8a-2) Swap the disks.

9a-2) After swapping the disks, execute the following command.

```
# sdxdisk -M -c Class1 -d c1t11d0=Disk1:single
```

10a) Create volumes with the sdxvolume command again. For the number of sectors, use the size recorded in (3a), in this example.

```
# sdxvolume -M -c Class1 -d Disk1 -v Volume1 -s 32768  
# sdxvolume -M -c Class1 -d Disk1 -v Volume2 -s 4128768
```

11a) Finally, restore the backup data to Volume1 and Volume2.

2b) In the event of a non-disk component failure, reboot your system after shutting down the system and recovering the failed component. There is no need to restore data, since slice data is valid.

(9) An I/O error occurs on a stripe volume or a volume in a concatenation group.

Explanation

Since a stripe volume or a volume within a concatenation group consists of only one slice, accessing the volume at the time of an I/O error will also result in an error. However, the status of slice and volume will remain ACTIVE.

Resolution

1) Identify the cause of I/O error occurrence by referring to the disk driver log message. You can confirm the error status of the disk related to volume and the physical disk name as shown below.

```
# sdxinfo -D -o Volume1 -e long
```

OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	FREEBLKS	DEVCONNECT	STATUS	E
disk	Disk1	concat	Class1	Group2	c1t1d0	17596416	*	node1	ENABLE	0
disk	Disk2	concat	Class1	Group2	c1t2d0	17596416	*	node1	ENABLE	1
disk	Disk3	concat	Class1	Group3	c2t3d0	17682084	*	node1	ENABLE	0
disk	Disk4	concat	Class1	Group3	c2t4d0	17682084	*	node1	ENABLE	0

In this example, I/O error occurs on Disk2, as shown in the E field. The physical disk name corresponding to Disk2 is c1t2d0, as shown in the DEVNAM field.

How to resolve the problem is described in two cases: (a) a disk component failure occurs, and you wish to restore the disk by using the backup data, (b) a non-disk component failure occurs.

2a) In case of a disk component failure, there will be no slices with valid data. Follow the procedures below and restore the data using the backup data.

3a) Record the configuration information of the group that was related to the failed disk using the `sdxinfo` command.

```
# sdxinfo -G -o Disk2 -e long
```

OBJ	NAME	CLASS	DISKS	BLKS	FREEBLKS	SPARE	MASTER	TYPE	WIDTH
group	Group1	Class1	Group2:Group3	70189056	65961984	*	*	stripe	32
group	Group2	Class1	Disk1:Disk2	35127296	*	*	*	concat	*
group	Group3	Class1	Disk3:Disk4	35127296	*	*	*	concat	*

In this example, lower level concatenation groups Group2 and Group3 are connected to the highest level stripe group Group1. Disks Disk1 and Disk2 are connected to Group2, and disks Disk3 and Disk4 are connected to Group3. The stripe width for Group1 is 32 blocks.

4a) Search the volumes that exists in the highest level group that is related to the faulty disk using the `sdxinfo` command.

```
# sdxinfo -V -o Disk2
```

OBJ	NAME	CLASS	GROUP	SKIP	JRM	1STBLK	LASTBLK	BLOCKS	STATUS
volume	*	Class1	Group1	*	*	0	65535	65536	PRIVATE
volume	Volume1	Class1	Group1	*	*	65536	98303	32768	ACTIVE
volume	Volume2	Class1	Group1	*	*	98304	4227071	4128768	ACTIVE
volume	*	Class1	Group1	*	*	4227072	70189055	65961984	FREE

In this example, Volume1 and Volume2 exist in the highest level group Group1, that is related to the faulty disk Disk2. The size of Volume1 is 32768 blocks, and the size of Volume2 is 4128768 blocks as shown in the `BLOCKS` field.

5a) Exit the application accessing the volume. When the volume is used as a file system, execute `umount` command. When I/O error occurs on `umount` command, execute `-f` option of `umount` command.

6a) Stop the volume with the `sdxvolume` command.

```
# sdxvolume -F -c Class1 -v Volume1,Volume2
```

7a) Remove the volumes with the `sdxvolume` command.

```
# sdxvolume -R -c Class1 -v Volume1
# sdxvolume -R -c Class1 -v Volume2
```

8a) Disconnect the faulty disk from the group. If the group is in a hierarchical structure, disconnect from the higher group in a descending order.

```
# sdxgroup -D -c Class1 -h Group1 -l Group2
# sdxdisk -D -c Class1 -g Group2 -d Disk2
```

In this example, the faulty disk Disk2 is connected to Group2, and Group2 is connected to Group1. Therefore, you should disconnect Group2 first, and then Disk2.

9a-1) Before swapping the disks, execute the following command.

```
# sdxswap -O -c Class1 -d Disk2
```

Caution: When the disk is the last disk in disk class, it will result in the following error. In such case, follow the procedures 9a-2, 10a-2 and 11a-2.

SDX:sdxswap: ERROR: Disk2: The last ENABLE disk in class cannot be swapped

10a-1) Swap the disks.

11a-1) After swapping the disks, execute the following command.

```
# sdxswap -I -c Class1 -d Disk2
```

9a-2) Before swapping the disks, execute the following command.

Caution: If no error occurs after performing 9a-1, you do not have to perform procedures in 9a-2, 10a-2 and 11a-2.

```
# sdxdisk -R -c Class1 -d Disk2
```

10a-2) Swap the disks.

11a-2) After swapping the disks, execute the following command.

```
# sdxdisk -M -c Class1 -d c1t2d0=Disk2
```

12a) Connect the swapped disk to the group, referring to the group information recorded in (3a). If the groups were in a hierarchical structure, connect the groups in an ascending order.

```
# sdxdisk -C -c Class1 -g Group2 -d Disk2  
# sdxgroup -C -c Class1 -h Group1 -l Group2 -a type=stripe,width=32
```

13a) Create volumes with the sdxvolume command again. For the number of sectors, use the size recorded in (4a), in this example, 32768 and 4128768.

```
# sdxvolume -M -c Class1 -g Group1 -v Volume1 -s 32768 -a pslice=off  
# sdxvolume -M -c Class1 -g Group1 -v Volume2 -s 4128768 -a pslice=off
```

14a) Finally, restore the backup data to Volume1 and Volume2.

2b) In case of a non-disk component failure, reboot your system after shutting down the system and recovering the failed component. There is no need to restore data, since slice data is valid.

F.1.4 Class Status Abnormality

(1) Class becomes closed status during operation.

Explanation

The class becomes closed when the number of configuration databases which store information on object configuration and object status within a class is insufficient, or when communication error between nodes occurs in a cluster environment.

All objects within a closed class are inaccessible.

Insufficient number of configuration databases will occur under the following conditions:

1. When there are no disks that can be accessed normally, if there are two or less disks in ENABLE status.
2. When there are one or less disks that can be accessed normally, if there are three to five disks in ENABLE status.
3. When there are two or less disks that can be accessed normally, if there are more than six disks in ENABLE status.

However, in case of root class, the class will not be closed unless there are no accessible disks.

GDS configuration databases cannot be stored in BCV devices and target (R2) devices since the devices are overwritten by data in copy source disks. Therefore, GDS does not regard BCV devices and target (R2) devices as "disks that can be accessed normally" described in the above conditions.

Resolution

- 1) You can check whether a class was closed during operation as follows. Do not reboot the system or restart sdxservd daemon, as it will make the checking impossible.

# /etc/opt/FJSVsdx/bin/sdxdcdown							
CLASS	DOWN	REASON	NDK	NEN	NDB	NLDB	DEVNAM
Class1	no	-	10	10	8	0	c1t1d0:c1t2d0:c1t3d0:c1t4d0:c2t1d0:c2t2d0:c2t3d0:c2t4d0
Class2	yes	Comm	10	10	8	0	c3t1d0:c3t2d0:c3t3d0:c3t4d0:c4t1d0:c4t2d0:c4t3d0:c4t4d0
Class3	yes	FewDB	10	10	1	7	c5t1d0
Class4	yes	NoDB	10	10	0	8	-

In this example, Class2, Class3, and Class4 with "yes" in the DOWN field are closed. The cause shown in the REASON field are as follows.

(Cause 1)

Comm Communication failure between nodes.

(Cause 2)

FewDB Insufficient number of valid configuration databases.

(Cause 3)

NoDB No valid configuration database.

- 2) Depending on the details of the cause, recovery may be difficult.

First, collect the investigation material.

For information on how to collect the investigation material, see "F.2 Collecting Investigation Material".

Resolutions are described for the following two cases:

a) closed due to a communication error

b) closed due to an insufficient number of configuration databases

3a) In case of Cause 1, contact your local customer support.

3b) In case of Cause 2 or 3, all (or majority) of the disks registered with class have abnormalities.

You can check the disks registered with class follows.

# sdxinfo -D -c Class3								
OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	DEVCONNECT	STATUS
disk	Disk31	mirror	Class3	Group1	c1t1d0	8847360	*	ENABLE
disk	Disk32	mirror	Class3	Group1	c2t1d0	8847360	*	ENABLE
disk	Disk33	mirror	Class3	Group2	c1t2d0	8847360	*	ENABLE
disk	Disk34	mirror	Class3	Group2	c2t2d0	8847360	*	ENABLE
disk	Disk35	mirror	Class3	Group3	c1t3d0	17793024	*	ENABLE
disk	Disk36	mirror	Class3	Group3	c2t3d0	17793024	*	ENABLE
disk	Disk37	mirror	Class3	Group4	c1t4d0	17793024	*	ENABLE
disk	Disk38	mirror	Class3	Group4	c2t4d0	17793024	*	ENABLE
disk	Disk39	spare	Class3	Group1	c1t5d0	17793024	*	ENABLE
disk	Disk40	spare	Class3	*	c2t5d0	17727488	*	ENABLE

In this example, ten disks from Disk31 to Disk40 are registered with Class3.

Physical disk names are shown in the DEVNAM field. Identify the cause of abnormality with these physical disks by referring to disk driver log message.

The cause of abnormality could be either of the followings:

(Failure 1)

Failed or defective non-disk component.

(Failure 2)

Failed disk component.

4b) Open the GDS configuration parameter file with an editor.

```
# vi /etc/opt/FJSVsdx/sdx.cf
```

Add the following line at the end of the file.

```
SDX_DB_FAIL_NUM=0
```

5b) Shut down the system.

6b) In case of Failure 1, recover the non-disk component (I/O adapter, I/O cable, controller, power supply, fan, etc).

7b) Boot the system.

8b) Confirm that objects within the class are accessible.

```
# sdxinfo -c Class3
```

If nothing is displayed, recovery was unsuccessful. You will have to contact your local customer support. If information is displayed normally, proceed with the following procedures.

9b) In case of Failure 2 when a disk component has failed, follow the procedures in " 5.3.3 In Case of Disk Abnormality," or "D.8 sdxswap - Swap disk", and swap the disks.

10b) After completing the recovery for both Failure 1 or 2, confirm the number of valid configuration databases as described below.

# /etc/opt/FJSVsdx/bin/sdxdcdown							
CLASS	DOWN	REASON	NDK	NEN	NDB	NLDB	DEVNAM
Class1	no	-	10	10	8	0	c1t1d0:c1t2d0:c1t3d0:c1t4d0:c2t1d0:c2t2d0:c2t3d0:c2t4d0
Class2	no	-	10	10	8	0	c3t1d0:c3t2d0:c3t3d0:c3t4d0:c4t1d0:c4t2d0:c4t3d0:c4t4d0
Class3	no	-	10	10	8	0	c5t1d0:c5t2d0:c5t3d0:c5t4d0:c6t1d0:c6t2d0:c6t3d0:c6t4d0
Class4	no	-	10	10	8	0	c7t1d0:c7t2d0:c7t3d0:c7t4d0:c8t1d0:c8t2d0:c8t3d0:c8t4d0

NLDB field gives the insufficient number of configuration databases. If this value is "0," the problem is resolved. If this value is "1" or more, there are still disks that have not been recovered. In the above example, all NLDB fields display "0," indicating the successful recovery.

11b) Open the GDS configuration parameter file with an editor.

```
# vi /etc/opt/FJSVsdx/sdx.cf
```

Remove the following line added at the end in procedure (5b).

```
SDX_DB_FAIL_NUM=0
```

12b) Reboot the system.

If you cannot perform the recovery with the described procedures, contact your local customer support.

F.1.5 System Disk Abnormality

Resolutions for system disk troubles with operating file systems such as /(root), /usr or /var are explained.

(1) Create backup of root file system.

Explanation

Mirroring the system disk will protect data in case one of the physical disks crashes. However, to restore data damaged by critical failures caused by multiple breakdown or misoperation, you must create backup data in advance.

Resolution

1) Exit all operating applications.

2) Shut down the system.

```
# shutdown -g30 i0 -y
```

3) Boot the system in single user mode.

```
ok boot -s
```

4) Check the volume name which is mounted with the file system you want to backup.

```
# mount
/ on /dev/sfdsk/System/dsk/root
  read/write/setuid/largefiles on Mon Aug  9 04:02:38 1999
/usr on /dev/sfdsk/System/dsk/usr
  read/write/setuid/largefiles on Mon Aug  9 04:02:38 1999
/var on /dev/sfdsk/System/dsk/var
  read/write/setuid/largefiles on Mon Aug  9 04:02:38 1999
```

In this example, file systems /(root), /usr, and /var is mounted on /dev/sfdsk/System/dsk/root, usr and var, respectively.

5) Create backup of file system data to a media, such as tapes.

```
# ufsdump [options] /dev/sfdsk/System/rdsk/root
```

In this example, backup of root file system data is created using the ufsdump(1M) command. For details, see manual attached to Solaris(TM) Operating Environment.

(2) System can be booted, but the system disk data is invalid.

Explanation

For some reason, the system disk data is invalid. You must restore data from the backup data created in advance.

Resolution

1) Exit all operating applications.

When absolute safety is required, create backup of system disk in advance.

For procedures on backup, see "Create Backup of Root File System".

2) Disconnect all disks from group, except the one disk you will use to restore data.

The command shown below is an example for when Disk Root1 and Root2 is mirrored, and Root1 is used for restoring data.

```
# sdxdisk -D -c System -g Group1 -d Root2
```

3) Shutdown the system.

```
# shutdown -g30 -i0 -y
```

4) Since file systems such as /(root), /usr, or /var cannot be unmounted during system operation, boot from a CD-ROM. You can boot the system by using the following OpenBoot ok prompt.

```
ok boot cdrom -s
```



Note

When booting from a CD-ROM device, use the same version with the OS you are trying to restore.



Note

Avoid booting from a CD-ROM device, since performing an illegal operation may disrupt the mirroring environment.

When there is no other alternatives, ensure to follow the procedures in this manual.

When performing operations not provided in the manual, such as mounting the physical slice at the time of recovery, mounting as read-only is recommended.

5) Create a special file for tape unit.

```
# tapes
```

6) Check the physical slice names where you want to restore data of file systems /(root), /usr or /var, and make a note of it.

```
# eeprom nvramrc | grep sdx-root-slice
devalias sdx-root-slice-1 /devices/pci@1f,4000/scsi@3/sd@0,0:a,raw
# eeprom nvramrc | grep sdx-usr-slice
devalias sdx-usr-slice-1 /devices/pci@1f,4000/scsi@3/sd@0,0:g,raw
# eeprom nvramrc | grep sdx-var-slice
devalias sdx-var-slice-1 /devices/pci@1f,4000/scsi@3/sd@0,0:b,raw
```

In this example, file system type is always ufs.

In regards to physical slice name:

/root is /devices/pci@1f,4000/scsi@3/sd@0,0:a,raw,

/usr is /devices/pci@1f,4000/scsi@3/sd@0,0:g,raw, and

/var is /devices/pci@1f,4000/scsi@3/sd@0,0:b,raw.



Note

Configurations without file systems /usr and /var will be displayed without the physical slice names as shown below.

```
# eeprom nvramrc | grep sdx-usr-slice
# eeprom nvramrc | grep sdx-var-slice
```

7) Restore the backup data created on a media, such as tapes.

In the example below, each file system data is restored using the ufsrestore(1M) command.

```
# newfs /devices/pci@1f,4000/scsi@3/sd@0,0:a,raw
# mount -F ufs /devices/pci@1f,4000/scsi@3/sd@0,0:a /mnt
# cd /mnt
# ufsrestore [options]
# cd /
# umount /mnt

# newfs /devices/pci@1f,4000/scsi@3/sd@0,0:g,raw
# mount -F ufs /devices/pci@1f,4000/scsi@3/sd@0,0:g /mnt
# cd /mnt
# ufsrestore [options]
# cd /
# umount /mnt

# newfs /devices/pci@1f,4000/scsi@3/sd@0,0:b,raw
# mount -F ufs /devices/pci@1f,4000/scsi@3/sd@0,0:b /mnt
# cd /mnt
# ufsrestore [options]
# cd /
# umount /mnt
```



Note

For mount (1M), do not end the physical slice name with , raw.

8) Install the boot block to root partition of system disk.

```
# installboot /usr/platform/`uname -i`/lib/fs/ufs/bootblk ¥
/devices/pci@1f,4000/scsi@3/sd@0,0:a,raw
```

9) Reboot the system.

```
# shutdown -g0 -i6 -y
```

10) Reconnect the disks that were disconnected in procedure 2).

```
# sdxdisk -C -c System -g Group1 -d Root2
```

Synchronization copying will automatically take place. After its completion, mirroring is recovered.

(3) System cannot be booted. (Failure of original boot disk.)

Explanation

Even when you have the system disk mirrored, you may experience unsuccessful booting. For example, you may encounter such a problem when the file accessed during the booting process in the boot disk is physically damaged, or when a disk is swapped improperly. If the console message of the unsuccessful boot disk indicates such is the case, try booting from the other boot disk, i.e. the mirror disk.

Resolution

1) Confirm the name of the mirror disk. Use the OpenBoot ok prompt shown below.

```
ok printenv boot-device  
/pci@1f,4000/scsi@3/disk@0,0:a /pci@1f,4000/scsi@3/disk@1,0:a
```

In this example, the original disk name is /pci@1f,4000/scsi@3/disk@0,0:a, which is indicated on the left. The mirror disk is /pci@1f,4000/scsi@3/disk@1,0:a. When mirroring is configured for 3-way multiplex mirroring or more, there would be more than one mirror disk.

2) Boot from the mirror disk.

```
ok boot /pci@1f,4000/scsi@3/disk@1,0:a
```

3) If booting is successful, find out what caused the original disk to fail, and recover by swapping the disk.

For details on disk swapping, see "D.8 sdxswap - Swap disk", or resolutions described in "5.3.3 In Case of Disk Abnormality."

When the above procedures cannot resolve the problem, it is possible that there are causes other than data abnormality of the original disk.

Perhaps data of the mirror disk may be abnormal. In such case, see "System cannot be booted (Boot disk data damage)"

(4) System cannot be booted (Boot disk data damage)

Explanation

For some reason, the system disk data is invalid. You must restore data from the backup data created in advance. If restoring data from the mirror disk was unsuccessful, follow the procedures shown below.

Resolution

1) Boot from the CD-ROM.

```
ok boot cdrom -s
```



Note

When booting from a CD-ROM device, use the same version with the OS your are trying to restore.



Note

Avoid booting from a CD-ROM device, since performing an illegal operation may disrupt the mirroring environment.

When there is no other alternatives, ensure to follow the procedures in this manual.

When performing operations not provided in the manual, such as mounting the physical slice at the time of recovery, mounting as read-only is recommended.

2) Create a special file for tape unit.

```
# tapes
```

3) Check the physical slice name where the you want to restore data of file systems /(root), /usr or /var, and make a note of it.

```
# eeprom nvramrc | grep sdx-root-slice
devalias sdx-root-slice-1 /devices/pci@1f,4000/scsi@3/sd@0,0:a,raw
devalias sdx-root-slice-2 /devices/pci@1e,4000/scsi@3/sd@0,0:a,raw
# eeprom nvramrc | grep sdx-usr-slice
devalias sdx-usr-slice-1 /devices/pci@1f,4000/scsi@3/sd@0,0:g,raw
devalias sdx-usr-slice-2 /devices/pci@1e,4000/scsi@3/sd@0,0:g,raw
# eeprom nvramrc | grep sdx-var-slice
devalias sdx-var-slice-1 /devices/pci@1f,4000/scsi@3/sd@0,0:b,raw
devalias sdx-var-slice-2 /devices/pci@1e,4000/scsi@3/sd@0,0:b,raw
```

In this example, file system type is always ufs.

In regards to physical slice name:

/(root) is /devices/pci@1f,4000/scsi@3/sd@0,0:a,raw(original), and

/devices/pci@1e,4000/scsi@3/sd@0,0:a,raw(mirror);

/usr is /devices/pci@1f,4000/scsi@3/sd@0,0:g,raw(original), and

/devices/pci@1e,4000/scsi@3/sd@0,0:g,raw(mirror);

/var is /devices/pci@1f,4000/scsi@3/sd@0,0:b,raw(original), and

/devices/pci@1e,4000/scsi@3/sd@0,0:b,raw(mirrors).

When mirroring is configured for 3-way multiplex mirroring or more, there would be more than one mirror slice.



Note

Configurations without file systems /usr and /var will be displayed without the physical slice names, as shown below.

```
# eeprom nvramrc | grep sdx-usr-slice
# eeprom nvramrc | grep sdx-var-slice
```

4) Restore the backup data on a media, such as tapes to original slices.

In the example below, each file system data is restored using the ufsrestore(1M) command.

```
# newfs /devices/pci@1f,4000/scsi@3/sd@0,0:a,raw
# mount -F ufs /devices/pci@1f,4000/scsi@3/sd@0,0:a /mnt
# cd /mnt
# ufsrestore [options]
# cd /
# umount /mnt

# newfs /devices/pci@1f,4000/scsi@3/sd@0,0:g,raw
# mount -F ufs /devices/pci@1f,4000/scsi@3/sd@0,0:g /mnt
# cd /mnt
# ufsrestore [options]
# cd /
# umount /mnt

# newfs /devices/pci@1f,4000/scsi@3/sd@0,0:b,raw
# mount -F ufs /devices/pci@1f,4000/scsi@3/sd@0,0:b /mnt
# cd /mnt
# ufsrestore [options]
# cd /
# umount /mnt
```

If restoring data to original slice is unsuccessful due to I/O error or such, try restoring it to mirror slice.



Note

For mount (1M), do not end the physical slice name with , raw.

5) Install the boot block to root partition of system disk.

```
# installboot /usr/platform/`uname -i`/lib/fs/ufs/bootblk ¥  
/devices/pci@1f,4000/scsi@3/sd@0,0:a,raw
```

6) Once restoring data from original slice is complete, copy data to mirror slice.

```
# dd if=/devices/pci@1f,4000/scsi@3/sd@0,0:a,raw ¥  
of=/devices/pci@1e,4000/scsi@3/sd@0,0:a,raw bs=1024k  
# dd if=/devices/pci@1f,4000/scsi@3/sd@0,0:g,raw ¥  
of=/devices/pci@1e,4000/scsi@3/sd@0,0:g,raw bs=1024k  
# dd if=/devices/pci@1f,4000/scsi@3/sd@0,0:b,raw ¥  
of=/devices/pci@1e,4000/scsi@3/sd@0,0:b,raw bs=1024k
```

When mirroring is configured for 3-way multiplex mirroring or more, there would be more than one mirror slice. You must copy data to all mirror slices.



Information

To reduce the copying time, specify the I/O block size of the dd (1M) command with bs=n option. Specifying the cylinder size value for I/O block size is recommended.

Cylinder size value may be calculated using the output of the prtvtoc (1M) command with the expression below.

(bytes / sector) X (sectors / cylinder)



Note

When physical slice contains the first sector (VTOC area), perform copying with the following procedures so that VTOC area of the mirror slice is not over written.

6-1) Check if the physical slice contains the first sector of the disk.

In the prtvtoc (1M) output, physical slice with 0 as its first sector contains the disk's first sector.

In the example below, physical slice (:a) of partition 0 contains the first sector.

```
# prtvtoc /devices/pci@1f,4000/scsi@3/sd@0,0:a,raw  
~  
*  
* Partition Tag Flags First Sector Last  
Sector Count Sector  
0 2 00 0 2097900 2097899  
~
```

6-2) Determine the I/O block size of the dd (1M) command. (See above "Information.")

6-3) Copy the first block, skipping its first sector.

Specify "I/O block size /512-1" for count option.

```
#dd if=/devices/pci@1f,4000/scsi@3/sd@0,0:a,raw  
of=/devices/pci@1e,4000/scsi@3/sd@0,0:a,raw iseek=1 oseek=1 count=2047
```

6-4) Copy the remaining blocks.

```
# dd if=/devices/pci@1f,4000/scsi@3/sd@0,0:a,raw  
of=/devices/pci@1e,4000/scsi@3/sd@0,0:a,raw iseek=1 oseek=1 bs=1024k
```

7) Reboot the system.

```
# shutdown -g0 -i6 -y
```

F.1.6 GDS Management View Abnormality

In this section, how to resolve problems related to Web-Based Admin View and GDS Operation Management View is explained. For problems described in (1) to (18), see "Web-Based Admin View Operation Guide" and follow the troubleshooting procedures.

(1) Web-Based Admin View top menu is not displayed. Or, even if it did, the screen is not displayed properly, or messages such as 0001 to 0050 are displayed.

(2) The screen does not change even after confirming 0002, 0003, or 0007 message. Or, the browser freezes when trying to reconnect to a previously-connected management server.

(3) The following message may be displayed on the management server or the Java console screen of the monitor node.

```
FJSVwvbs:OutOfMemoryError
```

(4) When trying to access to management server from browser, message 0005 appears and top menu is not displayed, or the message itself may not appear.

(5) Web-Based Admin View stopped operating after changing the IP address of the public LAN.

(6) When using the PRIMECLUSTER Operation Management View, existing cluster nodes are not displayed on the each product's node list screen.

(7) The screen does not change even after confirming the dialog message.

(8) The following message is displayed on the browser status bar, and top menu will not be displayed.

```
Applet com.fujitsu.webview.base.WebViewMain
```

```
error:java.lang.NoClassDefFoundError: com/sun/java/swing/JApplet
```

Or, Operation Management View for each product will not be displayed.

(9) It takes more than a few minutes before the Web-Based Admin View top menu is displayed. Or, even if it did, it freezes or message "0001 Cannot call Management server" may be displayed.

(10) When using a Java Plug-in on a client browser, and using a mouse on the Web-Based Admin View top menu or Operation Management View for each product, screen has problems such as displaying small flecks in the area where the mouse has been used.

(11) When starting or restarting the Web-Based Admin View, or rebooting the node, messages "FJSVwvbs:webview.cnf abnormal" or "Node process abnormal end." is displayed, and Web-Based Admin View does not start.

(12) When starting or restarting the Web-Based Admin View, message "WARNING:unrelated 'httpip' property." is displayed.

(13) After installing the manual package, the title menu for the manual is not displayed. Or, even after removing the manual package, the title menu for the manual is still displayed.

(14) When trying to start the menu for products that are the target of management from the Web-Based Admin View top menu, message "0023 This application must be started at the following URL" is displayed.

(15) Start button for products that are the target of management disappeared from the Web-Based Admin View top menu.

(16) Even after disconnecting from the management server, messages 0002, 0003 or 0007 is not displayed on the client.

(17) When trying to access the management server from the browser, message 0006 appears, and top menu is not displayed.

(18) When entering the user ID and the password for authentication, error message 0016 appears, and authentication fails.

(19) GDS Management View cannot be started.

Resolution

a. When GDS is not installed.

Install GDS.

b. When GDS is installed.

Check for the following possible causes:

- Wrong URL is specified.
- Wrong method is used for specifying host name.
- Environment cannot use Swing.
- Different browser type and level than recommended.

(20) Disk information such as mphd and mplb will not be displayed on GDS Operation Management View. Or, disk sizes for disk units are not displayed correctly.

Resolution

Click [Update Physical Disk Information] on [Operation] menu, and check if the disk information is displayed correctly.

F.1.7 EMC Symmetrix Abnormality

(1) SRDF is not utilized in the process of synchronization copying from a master to a proxy.

Explanation

The method handling copying from a master to a proxy can be confirmed in the CPTYPE field that is displayed with the `sdxinfo` command. When the CPTYPE field displays `soft` without displaying SRDF, instead of SRDF, the GDS soft copy function is performing the copy operation.

# sdxinfo -S -e long -o pv1												
OBJ	NAME	CLASS	GROUP	DISK	VOLUME	JRM	MODE	STATUS	COPY	CURBLKS	COPYBLKS	DLY
slice	*	Class1	Group1	Disk1	pv1	*	*	COPY	run	93952	10027008	0
CPTYPE		CPSOURCE										
-----		-----										
soft		*										

The following five possible causes can invite a situation that SRDF is not utilized.

(Cause a)

The configuration of the master or proxy group is inappropriate.

The source (R1) and target (R2) devices as the SRDF pair must be connected with the master and proxy groups respectively. With the proxy group, the target (R2) device only must be connected.

(Cause b)

A pair of volumes, instead of a pair of groups, was related as the master and proxy.

(Cause c)

When joining the master and proxy groups, the specific was set to create a proxy volume to the proxy group with the physical slice attribute that is different from that of the master volume.

(Cause d)

The master data was restored using the proxy with the `sdxproxy RejoinRestore` command.

When the `sdxproxy RejoinRestore` command performs master restoration, the SRDF pair is cancelled and SRDF becomes no longer available.

(Cause e)

With the `sdxproxy Join` command or the `sdxproxy Rejoin` command, the `-e` option was specified to utilize the soft copy function.

When the `-e softcopy` option is specified with the `sdxproxy Join` command, SRDF is not utilized.

When the `-e softcopy` option is specified with the `sdxproxy Rejoin` command, the SRDF pair is cancelled and SRDF becomes no longer available.

For details see "A.2.19 Snapshot by Cooperation with EMC TimeFinder or EMC SRDF".

Resolution

- 1) Check whether the possible cause is (Cause a).
 - 1.1) Identify the physical disk names of the disks connected with the master and proxy groups.

```
#sdxinfo -D -o mg1
```

OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	DEVCONNECT	STATUS
disk	Disk1	mirror	Class1	mg1	emcpower10	17596416	node1:node2	ENABLE
disk	Disk2	mirror	Class1	mg1	emcpower11	17682084	node1:node2	ENABLE

```
# sdxinfo -D -o pg1
```

OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	DEVCONNECT	STATUS
disk	Disk3	mirror	Class1	pg1	emcpower20	17596416	node1:node2	ENABLE

In this example, the physical disk emcpower10 and emcpower11 have been connected with the master group mg1, and the physical disk emcpower20 has been connected with the proxy group pg1.

If a lower level group is, or multiple disks or lower level groups are, connected with the proxy group, the possible cause is (Cause a).

1.2) Identify the Symmetrix IDs and the Symmetrix device names of the disks connected with the master and proxy groups using the sympd command provided by SYMCLI.

```
# sympd list
```

Symmetrix ID: 000183600262

Device Name	Directors			Device			
Physical	Sym	SA :P	DA :IT	Config	Attribute	Sts	Cap (MB)
...							
/dev/rdisk/emcpower10c	000	03A:0	02B:C0	RDF1	Grp'd	RW	8592
/dev/rdisk/emcpower11c	001	03A:0	02B:D0	Unprotected	N/Grp'd	RW	8632
/dev/rdisk/emcpower12c	002	03A:0	01B:D0	RDF2	Grp'd	RW	8632
...							

Symmetrix ID: 000282600920

Device Name	Directors			Device			
Physical	Sym	SA :P	DA :IT	Config	Attribute	Sts	Cap (MB)
...							
/dev/rdisk/emcpower20c	005	03A:0	02B:C0	RDF2	Grp'd	RW	8592
/dev/rdisk/emcpower21c	006	03A:0	01B:D0	RDF1	Grp'd	RW	8632
/dev/rdisk/emcpower22c	007	03A:0	01A:D0	RDF1	Grp'd	RW	8632
...							

In this example, as to the physical disk emcpower10c and emcpower11c, the Symmetrix ID is 000183600262 and the Symmetrix device names are 000 and 001 respectively. As to the physical disk emcpower20c, the Symmetrix ID is 000282600920 and the Symmetrix device name is 005.



For details on the sympd command, see the SYMCLI manual.

1.3) Check the source (R1) and target (R2) devices as the SRDF pair whether they are connected with the master and proxy groups respectively using the symrdf command provided by SYMCLI.

The SymDev field and the RDev field of the symrdf list command output display the Symmetrix device names of the source (R1) device and the corresponding target (R2) device respectively.

```
# symrdf list

Symmetrix ID: 000183600262

Local Device View
-----
Sym      RDF      STATUS      M O D E S      R1 Inv      R2 Inv      RDF S T A T E S
Dev  RDev  Typ:G  SA  RA  LNK  Mode  Dom  ACp  Tracks  Tracks  Dev  RDev  Pair
-----
...
000 005  R1:1  RW  RW  NR  SYN  DIS  OFF   0       0  RW  RW  Synchronized
002 006  R2:2  RW  RW  NR  SYN  DIS  OFF   0       0  RW  RW  Split
...

Symmetrix ID: 000282600920

Local Device View
-----
Sym      RDF      STATUS      M O D E S      R1 Inv      R2 Inv      RDF S T A T E S
Dev  RDev  Typ:G  SA  RA  LNK  Mode  Dom  ACp  Tracks  Tracks  Dev  Rdev  Pair
-----
...
005 000  R2:1  RW  RW  NR  SYN  DIS  OFF   0       0  RW  RW  Synchronized
006 002  R1:2  RW  RW  NR  SYN  DIS  OFF   0       0  RW  RW  Split
...

```

In this example, the source (R1) device 000 in the Symmetrix frame 000183600262 and the target (R2) device 005 in the Symmetrix frame 000282600920 are the SRDF pair. The master and the proxy have been configured properly, and the possible cause is not (Cause e).



For details on the symrdf command, see the SYMCLI manual.

1.4) If the possible cause is (Cause e), release the master-proxy relation once, and properly re-configure.

For example, according to the symrdf command results in procedure 1.3), you will see that the source (R1) device 006 in the Symmetrix frame 000282600920 and the target (R2) device in the Symmetrix frame 000183600262 are also the SRDF pair. The symrdf command results in procedure 1.2) are indicating that the physical disk names of the Symmetrix devices 006 and 002 are emcpower21 and emcpower12 respectively. The proxy configuration associated with SRDF can be created by registering emcpower21 and emcpower12 with a class, connecting them with respective groups, and relating the groups as the master and proxy.

In addition, the registration of the source (R1) device with a device group can be done after registering the disk with a class, but it must be carried out before joining the master and the proxy.

2) Check whether the possible cause is (Cause b).

```
# sdxinfo -GV -e long -o pv1
```

OBJ	NAME	CLASS	DISKS	BLKS	FREEBLKS	SPARE	MASTER	TYPE	WIDTH
group	pg1	Class1	Disk3	17596416	7503872	0	*	mirror	*

OBJ	NAME	TYPE	CLASS	GROUP	DISK	MASTER	PROXY	SKIP	JRM	MODE	CMODE	LOCK
volume	*	mirror	Class1	pg1	*	*	*	*	*	*	*	*
volume	pv1	mirror	Class1	pg1	*	mv1	*	off	on	ro	*	on
volume	*	mirror	Class1	pg1	*	*	*	*	*	*	*	*

1STBLK	LASTBLK	BLOCKS	STATUS	PSLICE	SNUM
0	65535	65536	PRIVATE	*	*
65536	10092543	10027008	STOP	on	1
10092544	17596415	7503872	FREE	*	*

In this example, the MASTER field of the group pg1 to which the proxy volume pv1 belongs displays the asterisk (*) and the MASTER field of pv1 displays the master volume name mv1, and the possible cause is (Cause b).

If the possible cause is (Cause b), the action as follows should be implemented.

2.1) Release the relation between the master and proxy volumes.

```
# sdxproxy Break -c Class1 -p pv1
```

2.2) Join a pair of groups instead of a pair of volumes as the master and proxy. In this example, the master group name is mg1.

```
# sdxproxy Join -c Class1 -m mg1 -p pg1 -a mv1=pv1:on
```

3) Check whether the possible cause is (Cause c).

```
# sdxinfo -V -e long -o mg1
```

OBJ	NAME	TYPE	CLASS	GROUP	DISK	MASTER	PROXY	SKIP	JRM	MODE	CMODE	LOCK
volume	*	mirror	Class1	mg1	*	*	*	*	*	*	*	*
volume	mv1	mirror	Class1	mg1	*	*	*	off	on	rw	rw	off
volume	*	mirror	Class1	mg1	*	*	*	*	*	*	*	*

1STBLK	LASTBLK	BLOCKS	STATUS	PSLICE	SNUM
0	65535	65536	PRIVATE	*	*
65536	10092543	10027008	ACTIVE	on	1
10092544	17596415	7503872	FREE	*	*


```
# sdxinfo -V -e long -o pg1
```

OBJ	NAME	TYPE	CLASS	GROUP	DISK	MASTER	PROXY	SKIP	JRM	MODE	CMODE	LOCK
volume	*	mirror	Class1	pg1	*	*	*	*	*	*	*	*
volume	pv1	mirror	Class1	pg1	*	mv1	*	off	on	ro	*	on
volume	*	mirror	Class1	pg1	*	*	*	*	*	*	*	*

1STBLK	LASTBLK	BLOCKS	STATUS	PSLICE	SNUM
0	65535	65536	PRIVATE	*	*
65536	10092543	10027008	STOP	off	1
10092544	17596415	7503872	FREE	*	*

In this example, the attributes of the physical slices of the master volume mv1 and the proxy volume pv1 are on and off respectively, and the possible cause is (Cause c).

If the possible cause is (Cause c), the action as follows should be implemented.

3.1) Release the master-proxy relation.

```
# sdxproxy Break -c Class1 -p pg1
```

3.2) Relate the master and the proxy again. With the -a option for the sdxproxy Join command, if the physical slice attribute is not specified, the physical slice attribute of a created proxy volume conforms to that of the master volume.

```
# sdxproxy Join -c Class1 -m mg1 -p pg1 -a mv1=pv1:on
```

4) Whether the possible cause is (Cause d) or (Cause e) can be determined by checking the command action history referencing the GDS log file /etc/opt/FJSVsdx/msglog/sdxservd.log.

If the possible cause is (Cause d) or (Cause e), the master and proxy should be disjoined once and be rejoined.

(2) The joined proxy volume is in the INVALID status.

Explanation

TimeFinder and SRDF may stop operating because the exclusive lock on a Symmetrix device is not released due to an event such as a node down. In this case, copying by TimeFinder or SRDF does not operate normally, and the proxy volume status may become INVALID.

The exclusive locks for Symmetrix devices are divided into two types. Whether respective exclusive locks are left not released can be checked using the symdev and symcfg commands provided by SYMCLI.

```
# symdev -lock list  
# symcfg -lock list
```

If an exclusive lock has remained, the lock number is displayed.

Resolution

Release the exclusive lock on the Symmetrix device with the following method.

As to an exclusive lock indicated by the symdev -lock list command, use the symdev command to release it. Substitute the lock number displayed with the symdev -lock list command for *number*.

```
# symdev -lock number release
```

As to an exclusive lock indicated by the symcfg -lock list command, use the symcfg command to release it. Substitute the lock number displayed with the symcfg -lock list command for *number*.

```
# symcfg -lockn number release
```



See

For details on the symdev command and the symcfg command, see the SYMCLI manual.

F.2 Collecting Investigation Material

When a problem occurs with GDS that can not be resolved by referring to this Troubleshooting appendix, collect the investigation material and contact your local customer support.



Note

- Investigation material must be collected promptly. With the passage of time, necessary information may be lost.
- When using GDS with a cluster system, collect investigation material from all nodes.

F.2.1 Collecting with fjsnap Command

The `fjsnap` command, a system information collection tool, is provided in the FJSVsnap package of the Enhanced Support Facility. When a problem occurs with GDS, this command allows you to collect investigation material to identify the cause.

How to execute the `fjsnap` command is as follows.

1. Log in with super user authority.
2. Execute the `fjsnap` command as follows.

```
# /opt/FJSVsnap/bin/fjsnap -h output
```

or

```
# /opt/FJSVsnap/bin/fjsnap -a output
```

output indicates either the special file name (example: `/dev/rmt/0`) of the output media which stores the investigation material collected with the `fjsnap` command, or the output file name.



See

For details on the `fjsnap` command, see the `README.ja` file and `README` file included in the FJSVsnap package, and Enhanced Support Facility user guide.

F.2.2 Collecting Initial Investigation Material with sdxsnap.sh Script

Since the amount of information collected with the `fjsnap` command is rather large, it is often required to be sent to us by postal mail, and not by e-mail, which could take a few days before our engineer can start investigating.

In order to promptly start the initial investigation, you can only collect the essential information with the `sdxsnap.sh` script. If you wish to reduce the time lapse before the actual investigation starts, collect information using both the `sdxsnap.sh` script and the `fjsnap` command, and e-mail this initial investigation material first.

Execute the `sdxsnap.sh` script by following the procedures below.

1. Login with the super user authority.
2. Execute the `sdxsnap.sh` script

```
# /etc/opt/FJSVsdx/bin/sdxsnap.sh -s [-d dir]  
# cd dir  
# tar cvf snap.tar SDXSNAP
```

Use the `-d` option to save the material in the `dir` directory. If the `-d` option is not used, the material will be saved in the `/tmp` directory.

Send us the `snap.tar` file, which will be used as initial investigation material.

Appendix G Frequently Asked Questions (FAQ)

In this section, frequently asked questions (FAQ) are compiled.

G.1 Operation Design

Questions regarding system operation design are explained.

Q: Is operation possible without using GUI (GDS Operation Management View)?

A

Yes. Two types of operation management interfaces are provided: GUI (Graphical User Interface) and CLI (Command Line Interface).

Q: Can I use a GUI (GDS Operation Management View) to manipulate proxy objects?

A

The current version does not support a GUI (Graphical User Interface) for manipulating proxy objects. Use the CLI (Command Line Interface) in proxy object manipulation.

Q: Can disks with different performance specifications, such as the size and the revolution speed, be mirrored?

A

Yes. However, when mirroring disks of different sizes, you will only be able to use the capacity of the smaller disk. For example, if a 4 GB disk and a 9 GB disk are mirrored, only 4 GB out of the 9 GB disk will be accessible. Also, when mirroring disks with different specifications such as the revolution speed, the read performance becomes unbalanced and the write performance will depend on the slower disk performance. Therefore, it is recommended to mirror disks with the same specifications. See "A.3.2 Guidelines for Mirroring" also.

Q: Do disks have to have the equivalent attributes, such as the size and the revolution speed, when I relate them as the master and proxy?

A

You can relate disks as the master and proxy even if they have different attributes, such as the size and the revolution speed. However, the read performance becomes unbalanced and the write performance will depend on the slower disk performance.

Q: Can I combine mirroring and striping?

A

You can mirror stripe groups. However, you cannot stripe across multiple mirror groups, using them as stripe columns.

Q: Do I need to specify certain settings in order to perform a hot swap?

A

No. As long as the disk you are using supports hot swap, no special settings are required. For procedures, see "5.3.3 In Case of Disk Abnormality" or "D.8 sdxswap – Swap disk."

Q: What is the purpose of single volume?

A

GDS is a software which improves availability and manageability of the disk data. Single volume is used when users need to improve manageability, rather than availability. For details, see "Chapter 1 Functions" in the manual.

For example, if availability is already realized by using disk arrays, the user will benefit by using the physical disk as a single volume excluded from mirroring, to improve its manageability.

Using GDS to manage the volume as a single volume not only ensures consistent management, but also allows you to include the volume in the mirroring configuration without stopping the business application, when the need arises.

Q: Can the messages output by GDS driver, daemon and command be displayed in a language other than English?

A

No. The messages are in English only.

Q: Are there any restrictions on the LUN (Logical Unit) configuration within the disk array when managing the disk array using GDS?

A

There are no restrictions on LUN RAID level or size.

Q: Are the mirrored disks differentiated as "primary disk" and "secondary" disks?

A

No. Disks mirrored with GDS are handled in the same way. Both the user and the application do not differentiate between the two.

However, the only exception is the system disk, since the primary disk needs to be initially booted with the OpenBoot. For details, see "F.1.5 System Disk Abnormality."

Q: When reading the mirror volume, which one of the slices in the volume will the IO request be issued to?

A

When a read is issued to a volume, the last access block number of each disk in the volume, will be checked. The slice with the value closest to the read block number will be read. In other words, disk with the shortest seek distance will be read first. When a write is issued to a volume, it will write to all slices in the volume (except from slice in ACTIVE status), and will return the write result after completing the write request on all slices.

Q: When mirroring a disk array that supports the hot spare function, which is preferable; using the spare disk of the disk array, or that of GDS's?

A

Generally, using the spare disk of the disk array is recommended.

Q: Which one of the GDS messages should be monitored?

A

Generally, the following messages given in the "Appendix E GDS Messages" should be monitored.

- "E.2 Driver" messages of PANIC and WARNING level.
- "E.2.4 Internal error messages"
- "E.3 Daemon" messages of HALT, ERROR and WARNING level.
- "E.3.5 Internal error messages"

For details, see "Appendix E GDS Messages."

Q: When configuring a 24-hour operation system, would files of incremental size such as message log file have any effect?

A

Since GDS's message log file is managed so that it will not exceed a certain size, there are no effects that need to be considered.

G.2 Environment Configuration

Questions regarding environment configuration are explained.

Q: How do I make a backup of a system disk that is managed with GDS?

A

For procedures on how to make backups of file systems `/`(root), `/usr` and `/var`, which cannot be unmounted while the system is booted, see "F.1.5 System Disk Abnormality."

Q: How do I make a backup of a disk (other than a system disk) that is managed with GDS?

A

The backup of GDS's logical volume can be made in the same way as a physical slice. However, make sure you make a backup using the logical volume. For example, do not boot from the CD-ROM and make a backup using the physical slice, since correct data may not be saved.

You can also efficiently create a backup with the mirroring function. For details, see "1.3.7

Snapshot by Detaching Slice", "D.5 sdxslice Slice operations" and "5.3.2 Backup."

You can also configure backup administration using proxy volumes minimizing the operation interrupting duration. For details, see "1.3.8 Proxy volume", "D.13 sdxproxy - Proxy object operations", "D.15 Snapshot Creation Using Command" and "D.16 Online Backup Making Using Command."

Q: The following message is output on the console after mirroring the system disk with GDS. Is there a problem? What is the resolution?

```
/usr/sbin/pmconfig: /etc/power.conf line (30) failed to convert mount point /dev/sfdsk/System/dsk/root to prom name: Invalid argument
```

A

There is no problem, and no resolution is necessary. The message says that by mirroring the system disk, the Power Management function and the suspend-resume function provided by the Solaris(TM) Operating Environment will become unavailable. Apart from these functions being unavailable, there is no effect.

For details on the Power Management function and the suspend-resume function, and whether the device you are using supports these functions, see the manual provided with Solaris(TM) Operating Environment.



Note

The Power Management system provided by Solaris(TM) Operating Environment is different from the automatic power control provided by Enhanced Support Facility (ESF). The above restrictions are unrelated to ESF's automatic power control.

Q: I cannot execute the format(1M) command or prtvtoc(1M) on a physical disk managed by GDS (for example c0t0d0s0). How can I solve this problem?

A

Since access to physical disk is restrained by GDS, you cannot perform operation. For details, see "A.2.3 Restraining Access to Physical Special File."

If all you need to do with the commands format(1M) or prtvtoc(1M) is to refer to the disk contents, you may do so by booting from the CD-ROM, and specifying the physical disk. However, you must exercise care since updating the data on the physical disk will damage the mirrored data. For details, see "A.2.4 Booting from a CD-ROM Device."

G.3 Operation

Questions regarding system operation are explained.

Q: Do I need to exit all applications to perform synchronization copying on the volume? If it is not necessary, what effect should I be aware of?

A

You do not need to exit the applications, but it could have some affect on the performance. If the application response is seriously affected, you can either suspend the synchronization copying, or adjust the load caused by the copying.

Interrupting the copying:

```
# sdxcopy -I -c class name -v volume name
```

Resuming the copying:

```
# sdxcopy -B -c class name -v volume name
```

Adjusting the copying load:

```
# sdxcopy -P -c class name -v volume name
```

Checking the copying process:

```
# sdxinfo -S -c class name -o volume name -e long
```

For details, see the command reference.

Q: How do I change the volume size?

A

Follow the procedures described below.

- 1) Exit the application that is accessing the volume
- 2) Backup the volume data as necessary.
- 3) Remove the volume.
- 4) Check if you can create a volume with the necessary size in the group in which the volume existed. You can check this as follows.

```
# sdxinfo -V -c class name -o group name
```

The volume which indicates FREE in the STATUS field will be available. The BLOCKS field indicates the volume size. When there is more than one available volume, the maximum size you can create would be the size of the largest available volume.
- 5) When you can not secure sufficient size within the group by performing procedure (4), backup a different volume and remove it, or create another group.
- 6) Specify a new size and create the volume. When creating the volume, specifying the same name as the volume you removed in procedure (3) will reduce the amount of information you have to input when changing the device definition information in the following procedure.
- 7) Update the device definition information such as /etc/vfstab, as necessary.
- 8) Restore the backup data you collected in procedure (2).

Q: I mistakenly updated the /etc/system file, and lost the definition for system disk mirroring. Now, I cannot boot the system. How do I restore the file?

A

The file contents at the point when the system disk mirroring was configured is saved in the following file.

/etc/opt/FJSVsdx/bkup/etc/system

Try booting the system by following the procedures described below.

1) From the OpenBoot ok prompt, use the boot -a command and boot the system as instructed.

2) For the name of system file, type as below.

/etc/opt/FJSVsdx/bkup/etc/system

3) For the physical name of root device, type as below.

/pseudo/sfdsk@0:2,blk

For details on the boot(1M) command, see the Solaris (TM) Operating Environment manual.

Q: I cannot boot the system. How do I restore it?

A

See "F.1.5 System Disk Abnormality"

Glossary

[A]

Advanced Copy function

Hard copying function of ETERNUS3000 and GR700 series storage system. A generic term for EC (Equivalent Copy) and OPC (One Point Copy). For details, see the ETERNUS3000 manual or GR700 series manual.

array disk

See "disk array."

alternate volume

Replication volume of master volume. Also referred to as "substitute volume."
The term "proxy" can refer to both proxy volume and proxy group.

attach slice

To add the detached slice to mirror volume making it accessible again as a component consisting the mirror volume.

available size of a disk

The capacity on an SDX disk that is available for creating volumes. It equals the physical disk size rounded to a cylinder boundary, minus the private slice size.

available size of a group

The capacity within the disk group that is available as volumes.

[B]

background copy

The actual copying that takes place between master and proxy, after virtual completion of creating a snapshot or restoring the master when performing instant snapshot or instant restoring. When performing normal copy, the volume to which the data is being copied cannot be activated. However, when performing background copy, the volume to which the data is being copied can be activated for users to access the valid data.

block

A unit for disk capacity or data transmission. Usually, one block is 512 bytes. Sector is used as a synonym.

break

To cancel the master-proxy relationship, and make the objects independent again.

[C]

class

See "disk class."

close down class

To make all objects within a class unavailable for reference.

close down volume

To make the volume unavailable (invalid).

cluster control

A component contained in PRIMECLUSTER responsible for detecting various abnormalities and state transition. For details, refer to "PRIMECLUSTER Installation / Administration Guide."

cluster system

A highly available and reliable system with multiple nodes connected via a special network. For details, refer to "PRIMECLUSTER Installation / Administration Guide."

concatenate

To connect more than one physical disks. It is helpful when using multiple disks as one virtual mass disk.

concatenation

To concatenate.

concatenation group

A disk group whose type attribute is "concat." It is also the set of disks that will be concatenated.

Configuration database

Database storing the configuration data of the objects managed by GDS. The configuration database is stored on a private slice.

connect disk

To add a SDX disk to a disk group. User can also connect an undefined disk. Spare disk will be connected automatically by GDS when required.

connect group

To add a group to another group.

The former group is called a "lower level group," and the latter, a "higher level group."

For example, when more than one stripe group is connected to a mirror group that is the highest level group, the mirror volume that is created within the highest level group will be a virtual device that will be both mirrored and striped.

[D]**detach slice**

To temporarily detach a mirror slice from a mirror volume, and making it inaccessible as a component consisting the mirror volume.

Other than by user command, GDS will automatically detach slice to maintain synchronization when an I/O error of a mirrored slice is detected.

disconnect disk

To remove a disk from a disk group.

disconnect group

To remove the lower level group from a higher level group.

disk

See "SDX disk."

disk array

A disk unit which combines multiple physical disks to be handled as one virtual disk. Spreading the access across multiple physical disks greatly improves performance and reliability. There are several access methods which are referred to as RAID (Redundant Array of Inexpensive Disks) levels. Main RAID level includes RAID 0 (striping), RAID 1 (mirroring), and RAID 5 (striping with parity).

disk class

A set of SDX objects. The user can create one or more disk class, and all objects will be managed under a certain disk class.

Unless you are using cluster system or a large-scale system with many disks which may require to separate the disk class, use only one class. Usually, the user does not have to be aware of the disk class.

Disk class has the following attributes, and all objects belonging to a certain class share those attributes.

Type

The "root" type or "local" type or "shared" type. With "root" type disk class, the object is available only to current node including keep disk. "local" disk class, the object is available only to current node (does not include keep disk). With "shared" disk class, the object is available to more than one node (does not include keep disk).

Scope

When the disk class is "shared" type, it shows the node set that can use the objects.

Shared disk class is also a resource used in a cluster application. Therefore, if multiple applications operating under a cluster system use SDX objects, the objects must be in separate disk class.

Also simply referred to as "class."

disk group

A set of disks or lower level groups that will be mirrored, striped or concatenated as a unit.

Type attribute is one of the attributes of a disk group. There are three kinds of type attributes; mirror, stripe or concatenation. Disks and lower level groups connected to the same disk group will be mirrored, striped or concatenated depending on the type attribute of the disk group.

Disk group is also simply referred to as "group."

[E]**EC**

Copying process using the EC (Equivalent Copy) function.

EC function

An abbreviation for Equivalent Copy function. It is one of the Advanced Copy function provided by GR700 series storage system. It keeps duplicating disk data within a disk unit. By performing detaching control with a software, a replication of the disk will be created.

For details, see the GR700 series manual.

EC session

EC request from the host (server) to the GR700 series storage system.

For details, see the GR700 series manual.

emcpower disk

A virtual device that provides multiple access paths with the emcpower driver. Being able to use multiple paths allows to spread out the load when accessing the disk, and increases the availability at the time of path failure.

[F]**FDU/FU**

A disk file unit with multiple physical disks mounted on a single enclosure.

[G]**GDS Management View**

Software which allows various operations including GDS configuration setting and operation management using the Web screen.

GFS

Software which provides two types of file system, GFS Local File System and GFS Shared File System.

For details, see "GFS Local File System" and "GFS Shared File System."

GFS Local File System

Local file system that features high-performance file access with contiguous block allocation of the file data, as well as a high-speed consistency recovery function after a system down, which is effective for file system switchover.

For details, refer to "PRIMECLUSTER Global File Services Configuration and Administration Guide."

GFS Shared File System

Shared file system that features simultaneous accessibility that maintains consistency between multiple Solaris systems to which a shared disk device is connected and enables the continuation of processing on another node even if a certain node fails.

For details, refer to "PRIMECLUSTER Global File Services Configuration and Administration Guide."

GDS Management View

Software which allows various operations including GDS configuration setting and operation management using the Web screen.

Global File Services

See "GFS."

group

See "disk group."

GR700 series

Generic term for ETERNUS GR710, GR720, GR730 and GR740 storage system. For details, see the GR700 series manual.

[H]**hard copy**

To copy disk data created within the device using the disk array's hardware function so that no load is imposed on the server or the SAN.

higher level group

A group to which another group belongs.

hot spare

A feature realizing automatic mirroring recovery while maintaining data accessibility at times of mirror disk failure.

hot swap

A feature allowing user to exchange the faulty physical disk components without stopping the application in use. Also referred to as active exchange.

[I]**instant restore**

To instantly restore the master's data, by overwriting the master using the parted proxy data.

instant snapshot

A snapshot of a master instantly created by cooperation with the OPC (One Point Copy) function provided by ETERNUS3000 or GR700 series storage system.

instant snapshot created by OPC

See "instant snapshot."

[J]**JBOD**

An abbreviation for "Just a Bunch Of Disks." Multiple disk units are used as a single disk. See also, "disk array."

join

To relate a pair of volumes or a pair of groups as a master-proxy and make the two synchronized.

joined

A state when master and proxy are joined. Proxy volume in joined status cannot be accessed.

JRM

See "just resynchronization mechanism."

just resynchronization copying

A high-speed resynchronization copying process. Available when just resynchronization mode is valid.

Just Resynchronization Mechanism

A mechanism realizing just resynchronization copying. Abbreviated as JRM.

[K]**keep disk**

A disk that retains the disk format and data when being registered with a class or being connected to a group.

[L]

local class

A disk class of a "local" type.

local objects

An object set that belongs to "local" type class.

logical slice

The smallest unit of object managed by GDS. One or more logical slices consist a logical volume. Each logical slice contains a data replication of the logical volume.

Mirror volume is configured by mirroring more than one logical slices. For each mirror volume, there will be one logical slice in each disk and each lower level group belonging to the mirror group.

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

Usually, the user does not have to be aware of the logical slice.

Note that logical slices consisting a mirror volume is called a "mirror slice," and a logical slice consisting a single volume is called a "single slice."

Logical slice is also simply referred to as "slice."

logical volume

A generic term for virtual devices, which the user can access directly. User can access a logical volume in the same way as accessing physical disk slice (partition).

Especially, a logical volume that is created within a single disk is called a "single volume," a logical volume that is created within a mirror group is called a "mirror volume," and a logical volume that is created within a stripe group is called a "stripe volume."

Logical volume is also simply referred to as "volume."

lower level group

A group that belongs to another group. A volume can not be created within the lower level group.

[M]

master

See "master object."

master group

The original group from which the proxy group is copied.

The term "master" can refer to both master group and master volume.

master object

A generic term for master volume and master group.

Also simply referred to as "master."

master volume

The original from which the proxy volume is copied.

Volumes within master group are also called master volumes.

The term "master" can refer to both master volume and master group.

mirror group

A disk group whose type attribute is "mirror." It is the set of disks and lower level groups, which will be mirrored each other.

mirror volume

A volume that is created within a mirror group. Its data will be made redundant by mirroring.

mirror slice

A logical slice that consists a mirror volume. Mirror volume is created by mirroring more than one logical slices. For each mirror volume, there will be one logical slice in each disk and each lower level group belonging to the mirror group.

mirroring

To maintain the copied data by using multiple mirror slices. A feature which allows you to continue access to a volume if one or more valid slices exist in the volume, even if some failures occur on the other slices.

MPHD disk

A virtual device that provides multiple access paths with the mphd driver. The user will be able to continue operation, even if an access path failure occurs.

MPLB disk

A virtual device that provides multiple access paths with the mplb driver. Being able to use multiple paths allows to spread out the load when accessing the disk, and increases the availability at the time of path failure.

[N]**node**

One computer (with one operating system). For details, refer to PRIMECLUSTER Installation / Administration Guide.

node identifier

A general name for each node within PRIMECLUSTER system.

node name

The name returned with the `uname -n` command

NVURM

An abbreviation for Non-Volatile Update Region Map. This is information of the updated volume region stored on a private slice for just resynchronization.

[O]**object**

See "SDX object."

online backup

To create volume backup while an application is accessing the volume.

online disk swapping

To replace a disk with another disk while an application is accessing the disk data.

online snapshot

To create a volume snapshot while an application is accessing the volume.

OPC

Copying process using the OPC (One Point Copy) function.

OPC function

An abbreviation for One Point Copy function. It is one of the Advanced Copy function provided by ETERNUS3000 and GR700 series storage system. Disk data within the disk unit at a certain point will be instantly copied to separate area in the disk unit. The logical copying process is completed immediately after the request to start copying is sent from the server. The actual copying process will be performed later in the background.

For details, see the ETERNUS3000 manual or GR700 series manual.

OPC session

OPC request from the host (server) to the ETERNUS3000 or GR700 series storage system.

For details, see the ETERNUS3000 manual or GR700 series manual.

[P]**part**

To temporarily separate proxy from master so that it may be accessed independently. The parted proxy becomes a snapshot with master data of the moment it has been parted. Although parted, the master-proxy relationship is maintained.

parted

The status of a proxy isolated from the master. When a proxy volume is parted, access to the proxy that is independent of access to the master is available.

physical disk

A storage media possible of handling random input and output of data. Mainly used to refer to a set of physical slices.

physical slice

A part of physical disk and a series of physically continued blocks. The starting sector number and the number of slices occupied by the physical slice is registered with the VTOC.

PRIMECLUSTER

Software to configure the Cluster System. For details, refer to "PRIMECLUSTER Installation / Administration Guide."

PRIMECLUSTER system

A cluster system configured by using PRIMECLUSTER. For details, refer to "PRIMECLUSTER Installation / Administration Guide."

private slice

A special slice used by GDS for control. Part of the SDX disk. Cannot be used as a logical slice.

proxy

See "proxy object."

proxy group

Replication group of master group.

It has the same volume configuration as the master group, and each volume is the proxy volume of the master volume that is within the master group.

The term "proxy" can refer to both proxy group and proxy volume.

proxy object

A generic term for proxy volume and proxy group.
Also simply referred to as "proxy."

[R]**REC**

Copying process by REC (Remote Equivalent Copy) function.

REC function

An abbreviation Remote Equivalent Copy function. It is one of the Advanced Copy function provided by GR700 series storage system. REC is EC function available between GR700 series boxes. For details, see the GR700 series manual.

register disk

To create a SDX disk by adding a physical disk to a certain disk class. Registered disks are managed by GDS.

rejoin

To re-join the parted proxy with master.

remove disk

To remove a SDX disk from disk class.

resource

The concept describing the elements managed by PRIMECLUSTER among hardware and software configuring the PRIMECLUSTER system. A PRIMECLUSTER system manages GDS disk classes as resources.
For details, refer to PRIMECLUSTER manual.

restore

To restore the master's data, by overwriting the master using the parted proxy data.

resynchronization copying

To copy data between mirror slices to recover mirroring state when rebooting after a system failure, after a cluster failover, or after re-attaching a detached slice to a volume.
Or, to copy data between master and proxy to match their data when rejoining a parted proxy to a master.

root class

A disk class of a "root" type.

root disk

A disk to which the root file system will be mounted. Also referred to as "boot disk."

root objects

An object set that belongs to "root" type class.

ROPC

Copying process by ROPC (Remote One Point Copy) function.

ROPC function

An abbreviation Remote One Point Copy function. It is one of the Advanced Copy function provided by GR700 series storage system. ROPC is OPC function available between GR700

series boxes. For details, see the GR700 series manual.

[S]

SAN

An abbreviation for "Storage Area Network." It is a network connecting multiple servers and storage devices. Usually, Fibre Channel is deployed for connection.

scope

A range of node set that can share an object when a disk class is shared.

SDX object

A generic term for disk class, group, SDX disk, volume, and slice. All of these objects are resources managed by GDS. SDX object may be simply referred to as "object."

SDX disk

A generic term for disks managed by GDS. SDX disk may be called single disk, keep disk, spare disk, or undefined disk, depending on its usage.
It may be simply referred to as "disk."

sector

See "block."

shared class

A disk class of a "shared" type.

shared objects

Object set which belongs to shared class. Each object may also be referred to as shared class, shared group and shared disk.

single disk

An SDX disk which does not belong to any group, and is capable of creating single volumes.

single node

A node that is not Cluster System configuration. For details, refer to "PRIMECLUSTER Installation / Administration Guide."

single slice

A logical slice that consists a single volume. Single volume consists of one single slice. It is also a series of physically continued block in a single disk.

single volume

A volume that is created within a single disk that does not belong to a group. Its data is not made redundant.

size of a group

The total of available size of a group and the size of one private slice.

slice

See "logical slice."

snapshot

A replica of a volume at a certain moment which is created on the other area.
GDS enables a snapshot to be created by detaching a slice or by using a proxy volume. When using proxy volumes, snapshots created by synchronization and instant snapshots created by

OPC are available.

snapshot by detaching slice

A snapshot of a mirror volume created on a slice by detaching the slice from the mirror volume temporarily.

snapshot by synchronization

A snapshot of a master created on its proxy by parting the master and proxy which are previously joined and synchronized.

soft copy

To copy disk data, using the GDS 's driver.

spare disk

A reserved disk to recover mirroring when a mirrored disk fails. When an I/O error is detected in a disk belonging to a mirror group or its lower level group, this spare disk will automatically be connected to recover normal mirroring.

SRDF

One of the EMC's products. An abbreviation for Symmetrix Remote Data Facility. SRDF consists of the microcode and hardware required to support Symmetrix remote mirroring. For details, see the SRDF manual.

start volume

To make volume accessible.

stop volume

To make volume inaccessible.

Storage Area Network

See "SAN."

storage device

A storage media, such as disk unit and tape unit, capable of data input and output.

stripe column

Disks and lower level groups belonging to a stripe group.

stripe group

A disk group whose type attribute is "stripe."

stripe volume

A volume that is created within a stripe group. Stripe volume can distribute the I/O load among multiple disks with its striping function. Its data will not be redundant.

stripe width

The size of divided data when striping.

striping

To divide and allocate data, alternately and evenly to two or more slices. Striping distributes the I/O load across multiple physical disks.

substitute volume

See "proxy volume."

swap

To swap the slices comprising master and proxy that are joined and synchronized.

Symmetrix

Storage devices from EMC. For details, refer to a guide to Symmetrix.

synchronization

A state when mirroring is operating normally.

Or, a state when the master and proxy data are the same.

synchronization copying

To copy data from a mirror volume to a mirror slice which is attached to the volume in order to recover a mirrored state or to increase the number of mirrored disks.

synchronized

A state when synchronization is maintained.

system disk

A disk that contains the Solaris(TM) Operating Environment running on the system. It refers to the entire disk containing slice operating as one of the following file systems (or swap area).

/, /usr, /var (or swap area)

[T]**TimeFinder**

EMC's software. For details, see TimeFinder manual.

the highest level group

A group that does not belong to any other group. Volumes can be created within the highest level group.

[U]**undefined disk**

An SDX disk that is neither a single disk, keep disk or a spare disk. Its usage is not specified yet.

[V]**volume**

See "logical volume."

VTOC

An abbreviation for "Volume Table Of Contents." The list of starting sector numbers of a physical slice, the numbers of sectors, etc. which exist on a physical disk.

It is stored in the first sector of the physical disk.

[W]**Web-Based Admin View**

Software for system operation and management using the Web screen.

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