



PRIMECLUSTER™

Global Disk Services

Configuration and Administration Guide 4.2

(Linux)

Appendix

June 2006

Appendix A General Notes

A.1 Rules

A.1.1 Object Name

Users can name the following objects:

- Classes
- Disks (excluding shadow disk)
- Groups
- Volumes

The object name can contain a maximum of 32 alphanumeric characters, including the hyphen (-) and the underscore character (_). However, in the event of the single disk name, the length limit is a maximum of 28 alphanumeric characters.

The object name cannot start with the hyphen (-) or the underscore character (_). Be sure to assign an alphanumeric character to the first character in the object name.

The class name must be unique within the entire system (for a cluster system, within the entire cluster system). For this reason, an error occurs if you try to create more than one disk class with the same name.

The other object names are unique within the class and an attempt to create more than one object with the same name in a class will result in an error.

GDS assigns slice names by combining the names of disks, groups, and volumes to which the slices belong. The slice naming conventions are as follows.

- When the slice is a mirror slice and exists in a disk that is connected directly to the highest level mirror group:
disk_name.volume_name
- When the slice is a mirror slice and exists in a lower level group that is connected directly to the highest level mirror group:
lower_level_group_name.volume_name
- When the slice belongs to a stripe volume:
the_highest_level_stripe_group_name.volume_name
- When the slice belongs to a volume created within the highest level concatenation group:
the_highest_level_concatenation_group_name.volume_name
- When the slice belongs to a switch volume:
active_disk_name.volume_name
- When the slice is a single slice:
single_disk_name.volume_name

You can designate each object uniquely in the entire system by the object name and the class name to which it belongs.



Note **Same Class Names [Linux2.6]**

Multiple single nodes on which classes with the same name exist can be changed over to a cluster system through installation of the cluster control facility. For details, see "A.2.24 Changing Over from Single Nodes to a Cluster System [Linux2.6]."



Note **Shadow Disk Name**

Name shadow disks according to the following rules:

- When the shadow disk is already registered with a class in another domain and managed as an SDX disk, assign the SDX disk name in said domain.
- When the shadow disk contains data copied from an SDX disk with the disk unit's copy function, assign the copy source SDX disk name.

A.1.2 Number of Classes

The number of root classes [Itanium] you can create for one node is limited to one. There is no limit to the number of local classes [Linux2.6] and shared classes.

Separate classes conforming to the following rules.

- Register system disks with a root class. [Itanium]
- It is recommended to register local disks other than system disks (disks used on one node) with local classes, but not with a root class, to differentiate the local disks from the system disks in management. [Linux2.6]
- Register shared disks in a cluster system (disks used from multiple nodes in the cluster) with shared classes.
- In a cluster system, register shared disks whose scopes (groups of sharing nodes) are different with separate shared classes.
- In a cluster system, for applications to use shared disks, create one or more shared classes with respect to each cluster application.
- Divide the class if the number of created disks or volumes in a class exceeds the limit. For the numbers of disks and volumes, see "A.1.3 Number of Disks" and "A.1.5 Number of Volumes."
- In a large-scale system to which numerous disk units are connected, separating classes based on physical configurations and data contents of disks may bring higher manageability.
- When a disk unit is expanded, unless the rules above apply, register expanded disks to existing classes. Creating new classes is not required.

Do not separate classes more than necessary. Keeping the number of classes to a minimum will offer the following advantages.

- If a class includes more disks, the probability that GDS configuration information stored on disks is lost due to disk failure will be lower. Therefore, not separating classes more than necessary to increase the number of disks within one class will raise system reliability.
- In a cluster system, if there are less shared classes, it takes shorter time to switch nodes.
- If there are fewer classes, less memory resources will be required.

A.1.3 Number of Disks

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

- To root class [Itanium], you can register up to 100 disks.
- To local class [Linux2.6] or shared class, you can register up to 1024 disks.

There are the following limits to the number of disks that can be connected to one group:

- To a mirror group, a maximum of 8 disks and lower level groups can be connected collectively. In other words, a maximum of eight-way multiplex mirroring is supported. However, be aware that the spare disk that will be automatically connected when a disk failure occurs is also included in the count.
- To a stripe group, a maximum of 64 disks and lower level groups can be connected collectively. In other words, a maximum of 64-column striping is supported.
- To a concatenation group, a maximum of 64 disks can be connected. In other words, a maximum of 64 disks can be concatenated.
- To a switch group, a maximum of 2 disks can be connected.

A.1.4 Number of Groups

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

- Within root class [Itanium], you can create up to 100 groups.
- Within local class [Linux2.6] or shared class, you can create up to 1024 groups.

A.1.5 Number of Volumes

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

- The number of volumes with the physical slice attributes set to "on" is limited to a maximum of 14.
- Volumes with the physical slice attributes set to "off" cannot be created.

There are the following limits to the number of volumes that can be created within a group or a single disk in the local class [Linux2.6] or shared class:

- You can create a maximum of 4 volumes with their physical slice attribute set to "on."
- You can create a total of 224 volumes with their physical slice attribute set to "on" or "off."
- You cannot create a volume with the physical slice attribute set to "on" in a stripe group or a concatenation group.
- You cannot create a volume with the physical slice attribute set to "on" in a shadow class regardless whether the physical slice is registered with the partition table.

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

- For the root class [Itanium], the number is limited to a maximum of 256.
- For a local class [Linux2.6], the number is limited to a maximum of 224.
However, when groups are nested, the nested group can contain a maximum of 224 volumes and lower level groups collectively.

A.1.6 Number of Keep Disks [Itanium]

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

A.1.7 Creating Group Hierarchy

The following nine kinds of group hierarchical structures, including a nonhierarchical structure, are available. However, groups that can be created in the root class are only non-nested mirror groups.

higher level group <-----> lower level group
mirror group (*1)
mirror group (*1) - stripe group (*3)
mirror group (*1) - stripe group (*3) - concatenation group (*7)
mirror group (*1) - concatenation group (*5)
stripe group (*2)
stripe group (*2) - concatenation group (*6)
concatenation group (*4)
concatenation group (*4) - switch group (*9) [Linux2.6]
switch group (*8)

Possible operations on groups at each hierarchical level that change the structure are as follows.

(*1) The highest level mirror group

- Disks, lower level stripe groups, and lower level concatenation groups can be connected or disconnected. However, disconnection is impossible if it can change the volume configuration or status.
- If no volume exists, the group itself can be removed.
- Volumes can be created or removed.

(*2) The highest level stripe group

- If no volume exists, disks and lower level concatenation groups can be connected or disconnected.
- If no volume exists, this type group can be connected to a mirror group.
- If no volume exists, the group itself can be removed.
- If more than one disks or lower level concatenation groups are connected to, volumes can be created or removed.

(*3) The lower level stripe group

- This type group can be disconnected from the highest level mirror group. However, disconnection is impossible if it can change the volume configuration or status.

(*4) The highest level concatenation group

- If no switch group is connected to, disks can be connected.
- If no disk is connected to, lower level switch groups can be connected.
- If no volume area exists on the disk that was connected last, that disk can be disconnected.
- If no volume area exists in the group that was connected last, that lower switch group can be disconnected.
- If no volume exists and if no lower level switch group is connected to, this group can be connected to a mirror group or a stripe group.
- If no volume exists, the group itself can be removed.
- Volumes can be created or removed.

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

- Disks can be connected.
- If more than one disk is connected and no volume area exists on the disk that was connected last, that disk can be disconnected.
- This type group can be disconnected from the highest level mirror group. However, disconnection is impossible if it can change the volume configuration or status.

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

- If more than one disk is connected and no volume exists within the highest level group, the disk that was connected last can be disconnected.

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

- None.

(*8) The highest level switch group

- Disks can be connected.
- If no volume or inactive disk exists, the active disk can be disconnected.
- The inactive disk can be disconnected.
- If no volume exists, the group itself can be removed.
- Volumes can be created or removed.

(*9) Lower level switch group [Linux2.6]

- Disks can be connected.
- The inactive disk can be disconnected.
- If no volume area exists and if this group is the switch group that was last connected to a higher level concatenation group, this group can be disconnected from that concatenation group.

A.1.8 Exception to Proxy Configuration

Although a pair of volumes or a pair of groups can be related as the master and the proxy, there are the following exceptions:

- Objects that belong to different classes cannot be related as the master and the proxy.
- Objects that belong to the same group cannot be related as the master and the proxy.
- Stripe type objects cannot be related as the master and the proxy. However, a mirror group to which a lower level stripe group is connected, or a mirror volume within that mirror group, can be related as the master and the proxy. For example, to stripe a master volume, connect a stripe group to a mirror group and create the master volume in the mirror group.
- Concatenation type objects cannot be related as the master and the proxy. However, a mirror group to which a lower level concatenation group is connected, or a mirror volume within that mirror group, can be related as the master and the proxy. For example, to concatenate disks to create a large-sized master group, connect a concatenation group to a mirror group and create the mirror group the master group.
- For a proxy that is already related to a master, another proxy cannot be related.
- A master to which a proxy is already related cannot be related as a proxy to another master.
- Shadow class objects cannot be related as the master and the proxy.

A.1.9 Number of Proxy Volumes

Although more than one proxy volume can be related to one master volume, the number of proxy volumes that can be related is limited to meet the following conditions:

- The total number of slices composing one master volume and slices composing any proxy volumes that are related to the master volume cannot exceed 32.

For example, if all master and proxy volumes consist of single volumes, a maximum of 31 proxy volumes can be related to one master volume.

A.1.10 Proxy Volume Size

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



Note

System Volume's Proxy Volume [Itanium]

When relating each individual proxy volume to a system volume, the proxy volume must be created in the group to which the keep disk with the same cylinder size as the system disk's cylinder size belongs. This is because the system volume size is the multiple of the cylinder size.

When relating each individual proxy group to a group to which system volumes belong, it is not necessary to consider the cylinder size. This is because the proxy group's cylinder size changes to the same size as the joined master group's cylinder size.

A.1.11 Proxy Group Size

For a local class [Linux2.6] or a shared class, the size of a proxy group must be larger than the last block number of volumes within the master group to which the proxy is related.

For a root class [Itanium], the size of the smallest physical disk that is directly connected to a proxy group must be larger than the last block number of volumes within the master group to which the proxy is related.

A.2 Important Points

A.2.1 Managing System Disks

Volumes in local classes and shared classes cannot be used as:

/ (root), /usr, /var, /boot, /boot/efi, swap areas

To manage disks for these uses with GDS, register the disks with the root class.



Note **Root Class**

With servers other than IPF servers, the root class cannot be created and system disk management cannot be performed.

A.2.2 Restraining Access to Physical Special File

After registration of disks with GDS is complete, accessing the disks using physical special files as below becomes impossible. GDS prevents unintentional access by the users to physical special files in order to protect mirroring statuses.

/dev/sdXn (for normal hard disk)

/dev/emcpowerXn (for emcpower devices)

X is a lowercase alphabetic character that indicates the device ID, and n is the slice number.

For example, if you execute the dd command to write to a disk using a physical special file, an error as below occurs.

```
# dd if=/dev/zero of=/dev/sda1
dd: /dev/sda1: open: Device busy
```

The disk list output by the parted(8) command does not display disks registered with GDS.

This access protection is canceled when a disk is removed from a class. It will also be canceled when an object is used for disk swap. The access protection function will turn back on if disk recovery operation is performed on the object.

Access to physical special files is prevented on the following nodes:

- For disks that are registered with disk classes
 - When registered with the PRIMECLUSTER's resource database
All of the nodes in the relevant domains
 - When not registered with the PRIMECLUSTER's resource database
Nodes where the relevant disk classes (root classes or local classes) reside
- For disks that are registered with shadow classes
Nodes where the relevant shadow classes reside

For the following disks, access to physical special files is not prevented. Be careful not to access these physical special files.

- Disks that are registered with classes in other domains
- Disks that are registered with root classes or local classes on other nodes and for which disk resources are not registered with the resource database
- Disks that are registered with shadow classes on other nodes

A.2.3 Booting from a CD-ROM Device

If the system cannot be booted, for example, it may be required recovering the system booting from a CD-ROM device.

However, there are the following concerns when booting from a CD-ROM device.

- An operation mistake can easily be caused since circumstances when booting with this method and when booting from a boot disk may alter the correlation between the disk and the device special file (`/dev/sd[a-z]*[1-4]*`).
- Only partial mirrored disks can be written in since access to the device special file of a physical slice is not restrained. Therefore, the mirroring state could collapse.

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

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

A.2.4 Initializing Disk

When physical disks are registered with classes other than shadow classes, all data contained in the physical disks will be lost since GDS automatically reformats the disks (excepting when the disks are registered with the root class as keep disks). Therefore, when registering physical disks that contain data with classes, back up disk data in advance, and restore the data after creating volumes.

If the device special file name of the physical slice on a physical disk registered with a class is set to `/etc/fstab` and so on, it is required to change it to the device special file name for the following volume.

Character device special file [Linux2.4]:

```
/dev/sfdsk/class_name/rdsk/volume_name
```

Block device special file:

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

If a disk restore operation is performed after swapping physical disks, data contained in the physical disk will also be lost.

When using GDS Management View, a confirmation window will be displayed if an operation that can cause data loss is attempted. However, when using a command, perform operations carefully because such a confirmation message is not displayed.

In order to register physical disks with classes other than shadow classes ensuring disk data, it is necessary to register the physical disks with the root class as keep disks.



Note

Root Class and Keep Disks

On a server other than an IPF server, keep disks cannot be registered with the root class.

A.2.5 Disk Size

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

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

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

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

- When the size of the disk first registered with the class is 10GB and below:
32MB
- When the size of the disk first registered with the class is over 10GB:
 $32\text{MB} + (0.1\% \text{ of the disk capacity to the cylinder boundary, rounded-up})$

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

When physical disks are registered with shadow classes, the sizes of private slices are determined based on the values stored in their private slices. For this reason, regardless of the order of registering physical disks with shadow classes, the sizes of private slices of the shadow classes become consistent with those of the corresponding disk classes. In other words, the order of registering physical disks with shadow classes is not a cause for concern.



Note

Cylinder Size

The cylinder size of a disk registered with a local class or a shared class is 32,768 blocks (= 16 MB). To calculate the sizes of disks to register with local classes or shared classes, assume that the cylinder size is 16 MB.



Note

Large Disks

1 TB or larger disks cannot be managed with GDS.

A.2.6 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 [Itanium]
A volume is created in the size calculated by rounding-up the size specified when creating the volume to a multiple of the keep disk's cylinder size.
(Example)
If the keep disk's cylinder size is 8 MB (= 16,384 blocks) and 20 MB (= 40,960 blocks) is specified as the size of a volume when it is created, the size is rounded to a multiple of 16,384 blocks, and a volume with 49,152 blocks (= 24 MB) is created.
- When creating a volume into a stripe group
A volume is created in the size calculated by rounding-up the size specified when creating the volume to a common multiple of (stripe width) * (stripe column number) and the cylinder size (32,768 blocks = 16 MB).
(Example)
If the stripe width is 32, the number of the stripe columns is 3 and 20 MB (= 40,960 blocks) is specified as the size of a volume when it is created, the size is rounded to a common multiple of 96 (= 32 * 3) blocks and 32,768 blocks, and a volume with 98,304 blocks (= 48 MB) is created.
- When creating other volumes
A volume is created in the size calculate by rounding-up the size specified when creating the volume to a multiple of the cylinder size (32,768 blocks = 16 MB).
(Example)
If 20 MB (= 40,960 blocks) is specified as the size of a volume when it is created, the size is rounded to a multiple of 32,768 blocks, and a volume with 65,536 blocks (= 32 MB) is created.



Note

Large Volumes

For 1 TB or larger volumes, mirroring and snapshot operations cannot be performed.

A.2.7 Hot Spare


- **Hot Spare for Hierarchized Mirror Groups**
If an I/O error occurs in a disk that is connected to a lower level group, a spare disk is automatically connected to the highest level mirror group, but not to the lower level group. For example, if a disk and a concatenation group are mirroring each other, and an I/O error occurs in one of the disks connected to the concatenation group, a spare disk is automatically connected in place of the concatenation group for mirroring.

A spare disk is selected independently of the disk case or controller number of a disk where an I/O error occurred.

In external hot spare mode (default), spare disks are selected randomly.

In internal hot spare mode, spare disks whose controller number is 0 and that do not belong to disk array units are selected.

Figure: Hot Spare for Hierarchized Mirror Groups



- **Number of Spare Disks**
There is no limit to the number of spare disks that can be registered with one class. Although there is no general rule in deciding the number of spare disks, it is recommended to assign 10% of disks and lower level groups for spare disks. In other words, one spare disk for every 10 disks or lower level groups combined is a good rule of thumb.
- **Spare Disk Size**
Spare disk automatic connection is restrained if there is not sufficient space on a spare disk to copy volumes in a mirror group. It is recommended to assign the largest size disk within the class for the spare disk.
- **Hot Spare for Proxy Volumes**
Spare disks are not connected to groups that include proxy volumes. It is recommended to create proxy volumes in groups other than those that include volumes used by primary services, or on single disks.
- **Shadow Classes**
Spare disks cannot be registered with shadow classes.
- **Disk Array Unit's Hot Spare Function**
If disk array units with hot spare functions are mirrored, it is recommended to use their own hot spare functions.
- **Spare Disk Failure**
If an I/O error occurs in a spare disk that was automatically connected to a mirror group, another spare disk will not automatically be connected in place of the failed spare disk.
- **Synchronization Copying Invoked by Hot Spare**
The synchronization copying with hot spare is run at lower speed as compared to similar copying with other events (such as volume creation and disk creation) in order to suppress the load imposed on the system. By default, delay time by 50 milliseconds is set. To change this delay time, use the `sdtparam` command.



[See](#)

For details, see "D.12 `sdtparam` - Configuration parameter operations."

- **Spare Disk Manual Connection**

In the within-case hot spare mode, if a disk case totally becomes inaccessible due to an I/O cable coming out or disk case going down, a spare disk that belongs to another disk case is not automatically connected. For example, if disk case 1 shown in [Chapter 1-1 Figure: Hot Spare in Internal Mode] of "1.2.2 Hot Spare" is down, a spare disk (disk 4) is not automatically connected in place of disk 1.

In such an event, follow the procedures below to manually recover the mirroring status by using a spare disk.

- 1) Change the spare disk to an undefined disk.



See

For disk type changing methods, see "Changing Disk Class Attributes" in "5.4.1 Class Configuration" when using the GDS Management View, or "D.7 sdxattr - Set objects attributes" when using the command.

- 2) Connect the disk in step 1 to the mirror group where the I/O error occurred.



See

For the disk connection methods, see "5.4.2 Group Configuration" when using the GDS Management View, or the description about the -C option in "D.2 sdxdisk - Disk operations" when using the command.

A.2.8 System Disk Mirroring [Itanium]

You can mirror the system disk containing the operating root file system by registering the system disk with root class. When registering a system disk, you must specify "keep" for its disk type. You cannot register system disks with local class and shared class.



Note

In the root class, the logical partitioning, disk concatenation, and disk striping functions are not available.

A.2.9 Keep Disk [Itanium]

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

- The partitioning method is not the GUID partition table (GPT).
- There is no unused slice number (not assigned to a slice) 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.2.5 Disk Size."

Additionally, when registering physical disks other than system disks (with running /, /usr, /var, /boot, /boot/efi, or swap areas) as keep disks, all slices on the physical disks must be in open status. For example, if the slice is mounted as a file system, unmount it in advance.

You cannot register a keep disk with local class or shared class.

A.2.10 Creating a Snapshot by Slice Detachment

A slice can be detached only from a mirror volume with a physical slice. Therefore, if a disk is not connected directly to a mirror group, it is impossible to perform snapshot creation by slice detachment. In addition, slices cannot be detached from shadow volumes and from mirror volumes in the root class.

A.2.11 The Difference between a Mirror Slice and a Proxy Volume

Although data matches on mirrored slices or synchronized master volumes and proxy volumes, the purposes of use are different.

Mirrored slices are equal one another, and their purpose is to maintain data redundancy in order to provide continuous access as long as any normal slice remains even if an error occurs in one of the slices.

However, even if the master volume and the proxy volume are synchronized, they are separate volumes and not equals. You may consider the master the primary volume, and the proxy the secondary volume. This means that you cannot continue accessing a master volume where all slices comprising the master volume are abnormal even if proxy volumes are normal. The purpose of proxy volumes is to create snapshots (saved copies from the master volume at certain moment) for a different service running concurrently with the primary service but not to improve the data redundancy of the master volume used in the primary service.

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

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



See

See "Figure: Difference between a Mirrored Slice and Synchronized Proxy Volume" in "1.5.1 Snapshot by Synchronization."

A.2.12 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.



See

- For details on the setting methods, see "5.2.2.4 Volume Configuration" and "D.4 sdxvolume - Volume operations."
- For details on the changing methods, see "5.4.3 Volume Configuration" and "D.7 sdxattr - Set objects attributes."
- When using GDS Management View, the mode ("on" or "off") of JRM for volumes can be checked using the volume information field in the Main Screen. For details, see "5.3.1.1 Confirming SDX Object Configuration."
- When using a command, the mode of JRM for volumes can be checked using the JRM field of the volume information displayed with the sdxinfo command. For details, see "D.6 sdxinfo - Display object configuration and status information."

JRM for Slices

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

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

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



See

- For details on the setting methods, see "D.5 sdxslice - Slice operations."
- For details on the changing methods, see "D.7 sdxattr - Set objects attributes."
- The mode of JRM for slices can be checked using the `JRM` field of the slice information displayed with the `sdxinfo` command with the `-e` long option. For details, see "D.6 sdxinfo - Display object configuration and status information."

JRM for Proxies

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

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

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

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



See

- For details on the setting methods, see "5.3.2.2 Backup (by Synchronization)" and "D.14 sdxproxy - Proxy object operations."
- For details on the changing methods, see "D.7 sdxattr - Set objects attributes."
- When using GDS Management View, the mode ("on" or "off") of JRM for proxies can be checked using the proxy volume information field in the Main Screen. For details, see "5.3.1.2 Viewing Proxy Object Configurations."
- When using a command, the mode of JRM for proxies can be checked using the `PJRM` field of the volume information displayed with the `sdxinfo` command with the `-e` long option. For details, see "D.6 sdxinfo - Display object configuration and status information."

A.2.13 Online Volume Expansion

- **Volume Configuration Limitations**

Online volume expansion is available for volumes in the following configurations.

- Single volume
- Mirror volume
 - Any mirroring multiplicity is supported.
 - Hierarchized groups are supported.

- **Online Mirror Volume Expansion**

For mirror volumes with mirroring multiplicity of two and higher, change the mirroring multiplicity to one, expand the volumes, and then execute synchronization copying for re-mirroring. See the outline of the operating procedures below. These operations are executable without stopping applications using the volumes.

- 1) Disconnect disks and lower level groups from the mirror group to change the mirroring multiplicity to one.
- 2) Expand the volume size with the `sdxvolume -S` command.

3) Reconnect the disks and lower level groups disconnected in step 1 with the mirror group.

If the mirror volumes are active, resynchronization copying is automatically performed after step 3 is done. If the mirror volumes are inactive, similar copying is automatically performed when they are started.

- **Stripe Type Volume and Concatenation Type Volume Expansion**

The capacity of stripe type volumes and concatenation type volumes cannot be expanded. To expand one of these types, back up data, recreate the volume, and then restore the data back to the volume. In configurations where a striping group or a concatenation group is connected to a mirror group (with any multiplicity), volume expansion applied using the striping or concatenation feature is possible.



[See](#)

For the methods of backing up and restoring, see "Chapter 6 Backing Up and Restoring."

- **Concatenation and Online Volume Expansion**

Even if there is no sufficient continuous free space after the last block of a volume, by concatenating unused disks, online volume expansion will be available. Online volume expansion is available for volumes that meet all of the following conditions:

- Volumes belong to a mirror group;
- To the mirror group, one or more concatenation groups are connected;
- To each of the concatenation groups, one or more disks are connected.

To use this function, create volumes in configuration conforming to these conditions in advance.

For example, if there is only one available disk, connect only the disk to a concatenation group, connect the concatenation group to a mirror group, and create volumes in the mirror group.

- **Expansion of Areas Used by Applications**

After volumes are expanded, applications such as file systems and databases need to recognize the expanded areas with methods specific to the applications.

If an application using a volume cannot recognize an expanded volume area, do not expand the volume. If such a volume is expanded, the application may no longer operate normally or volume data may be unavailable. A volume that contains the GFS Shared File System cannot be expanded.

A.2.14 Swapping Physical Disks

For the procedures for swapping disks, see “5.3.4 Disk Swap” when using GDS Management View, or see “D.8 sdxswap - Swap disk” when using a command.

This sub-section describes important points about disk swapping.

Physical Disk Size

Physical disk swapping cannot be performed using a physical disk whose size is smaller than the original disk size.

Physical Disks That Cannot Be Swapped

Physical disk swapping cannot be performed for a disk where the only valid slice (in ACTIVE or STOP status) within the volume exists.

For example, it is impossible to perform disk swapping if:

- The volume is in a group to which only one disk is connected.
- The volume is on a single disk that is not connected to a group.

In these situations, it will be possible to perform disk swapping by making any one of the following configuration changes.

- a. If the disk to be swapped is connected to a mirror group, add a new disk to the mirror group and complete volume synchronization copying normally.
- b. When performing disk swapping for a single disk, add the single disk and another unused disk to a new mirror group, and complete volume synchronization copying.
- c. Remove the existing volume from the disk to be swapped.
Before removing the volume, back up volume data if necessary.

When data on a disk to be swapped is valid, for example when conducting preventive maintenance, a or b is recommended. Here, if the disk unit supports hot swapping, disks can be swapped without stopping active applications.

Swapping Physical Disks Registered with Shadow Classes

When swapping disks of shadow disks registered with GDS Snapshot shadow classes, the shadow disks must be removed with the relevant GDS Snapshot command first. For details on GDS Snapshot commands, see “Appendix D Command Reference.”

The subsequent operations vary depending on whether the disks to be swapped are registered with disk classes.

- If the disk is registered with a disk class
After removing the shadow disk, perform disk swapping in the domain managing the disk class.
An error in a disk unit may not cause failures on both the related SDX and shadow objects. Even if only either of the objects fails, shadow disk removal and disk unit swapping are both necessary.
- If the disk is not registered with a disk class
In this situation, the disk to be swapped is a copy destination of the disk unit’s copy function.
It is not necessary to perform GDS operations described in “5.3.4 Disk Swap” and “D.8 sdxswap - Swap disk.” After removing the shadow disk, perform disk swapping referring to the manual of the disk unit’s copy function.

A.2.15 Object Operation When Using Proxy

If classes, groups, and slices are related to master volumes or proxy volumes, operations that change the class attributes, change the group configurations or attributes, or handle the slices cannot be performed. To perform such operations, cancel the relationship between the master and the proxy once. For objects that are related to the master and the proxy:

- Master volumes can be started or stopped.
- Proxy volumes that are parted from the master can be started or stopped.
- New volumes can be created within a group other than the master group or the proxy group and existing volumes other than the master or the proxy can be removed.
- The hot spare function is available for groups with master volumes. Spare disks are not connected to groups with proxy volumes.

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

- Attributes of master and proxy volumes can be changed with the `sdxattr -V` command.
- 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 parameters of the synchronization copying changed. However, synchronization copying between volumes or its parameters cannot be changed.
- Master volume can be restored with the `sdxfix -V` command.

For EC sessions, BCV pairs, and SRDF pairs between the master and the proxy, check the FUNC field displayed with the `sdxinfo` command. If the master and the proxy are in parted status, a copy session can be canceled with the `sdxproxy Cancel` command.

A.2.16 Using the Advanced Copy Function in a Proxy Configuration

In a proxy configuration, by working in cooperation with the Advanced Copy functions ((R)EC/(R)OPC) of Fujitsu Storage Systems ETERNUS, the copying between master and proxy can be performed without imposing loads on primary servers or a SAN. In this situation, disk array units carry out copying processes and the processes will continue running even if the server is rebooted.

When the Advanced Copy function is available, it is used for copying between the master and the proxy. However, in the following situations, a soft copy function (a copy function of a GDS driver operating on a server) is used.

- The Advanced Copy function is not used when:
 - Use of soft copy functions was specified explicitly using the `sdxproxy` command with the `-e softcopy` option
 - The master and the proxy belong to a root class. [Itanium]
 - The copy destination volumes are in a mirroring configuration.
Note, however, that even if the master as the copy destination is in a mirroring configuration, OPC is available for copying from the proxy when:
 - Executing [Operation]:[Proxy Operation]:[Restore] in GDS Management View and selecting "No" to "Rejoin" in the [Restore Master] dialog box. For details see "5.3.3 Restore."
 - Executing the `sdxproxy Restore` command. For details see "D.14 `sdxproxy` - Proxy object operations."
 - A lower level group is connected to a group to which master volumes or proxy volumes belong.
 - The number of concurrent EC or OPC sessions has reached the upper limit defined by the disk array unit.



See

The number of allowed concurrent sessions is either the upper limit within one physical disk (LU) or the upper limit within one disk array unit. For details see the handbook of the relevant disk array.

For details on the upper limit of supported concurrent sessions, see the disk array device handbook.

- Multiple proxy volumes are related to a master volume and the number of proxy volumes with EC sessions has reached the upper limit (16 volumes).
- Disks that constitute the master or the proxy were registered with a class before installing the Advanced Copy function on the disk array unit.
- Disks that constitute the master or the proxy were registered with a class before installing GDS Snapshot.

In these situations, operations based on OPC functions are impossible. For such operations see "A.2.17 Instant Snapshot by OPC."

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

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

If any EC sessions exist between a master and a proxy, OPC cannot be initiated between the master and another proxy.

When both EC and OPC is available as with ETERNUS6000 series, EC has precedence over OPC. Once an EC session is stopped, you cannot use the EC function afterwards. EC sessions are stopped in the following situations.

- EC sessions are stopped when:
 - Executing [Operation]:[Proxy Operation]:[Part] in GDS Management View and selecting "Yes" to "Instant Snapshot" in the [Part Proxy] dialog box
 - Canceling the EC sessions with the `sdxproxy Cancel` command
 - Joining a master and a proxy and using a soft copy function with the `sdxproxy Join -e softcopy` command
 - Creating instant snapshots with the `sdxproxy Part -e instant` command
 - Rejoining a master and a proxy and using a soft copy function with the `sdxproxy Rejoin -e softcopy` command

To make the EC function available after conducting these operations, break the relationship between the master and the proxy once and rejoin them.

To check the modes of the copying in execution, use either:

- The [Copy Type] field of the slice information field in GDS Management View
- The CPTYPE field displayed with the `sdxinfo` command

Additionally, the types and the statuses of sessions between the master and the proxy can be viewed in the FUNC field and the CPSTAT field displayed with the `sdxinfo` command.



Note **Advanced Copy Control**

When the Advanced Copy function is available, executing the `sdxproxy` command directs GDS to control the Advanced Copy on the master and the proxy. Do not use any other methods other than the `sdxproxy` command to apply such control on a master and a proxy.

A.2.17 Instant Snapshot by OPC

The following functions are based on OPC functions of Fujitsu Storage Systems ETERNUS.

- Instant snapshot by OPC
 - [Operation]:[Proxy Operation]:[Update] in GDS Management View
 - "Yes" to "Instant Snapshot" in the [Part Proxy] dialog box invoked through [Operation]:[Proxy Operation]:[Part] in GDS Management View
 - `sdxproxy Update` command
 - `sdxproxy Part -e instant` command
- Master restoration by OPC
 - "No" to "Rejoin" in the [Restore Master] dialog box invoked through [Operation]:[Proxy Operation]:[Restore] in GDS Management View
 - `sdxproxy Restore` command

These functions are unavailable with disk units other than the following types.

ETERNUS8000 series, ETERNUS6000 series,
ETERNUS4000 series, ETERNUS3000 series,
GR720, GR730, GR740

These functions are also unavailable under such conditions that prevent the use of Advanced Copy functions as described in "A.2.16 Using the Advanced Copy Function in a Proxy Configuration."



Note **Rebooting a Server While OPC Running**

Even if a server is rebooted while the copying between master and proxy is being processed by OPC, the OPC copying will continue running. However, if OPC sessions are not present when the server is up again, GDS will assume that the copying failed and copy destination volumes will be INVALID. For recovering this status, see "(4) Master volume is in INVALID status." and "(5) Proxy volume is in INVALID status." in "F.1.3 Volume Status Abnormality."

A.2.18 To Use EMC Symmetrix

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

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 EMC's SAN management software (Volume Logix, ESN Manager, SAN Manager and so on).

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

[Procedure]

1. The `syminq` command provided by SYMCLI is available for checking BCV, R2, GateKeeper, and CKD devices. Execute the `syminq` command, and describe all devices indicated as BCV, R2, GK, and CKD (`sdX`, `emcpowerX`) in the Excluded Device List.
2. The `syminq` command is unavailable for checking VCMDB devices. When using EMC's SAN management software (Volume Logix, ESN Manager, SAN Manager and so on), ask your EMC engineer or systems administrator who configured that SAN management software about the names of VCMDB devices and describe them in the Excluded Device List.
3. Describe all native devices (`sdX`) in the Excluded Device List.

emcpower0 (Target)
sdb (Nontarget)
sdc (Nontarget)
4. In addition to the devices specified in steps 1) through to 3), describe any other devices to be excluded from GDS management in the Excluded Device List.

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

The Excluded Device List should appear as follows.

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



Information **exdevtab.sh**

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

To use one of these script, open the script with an editor and modify the following two parameters (the `syminq` command and `powermt` command paths) according to the execution environment.

```
SYMINQ=/usr/symcli/bin/syminq
POWERMT=/etc/powermt
```

By executing `exdevtab.sh`, the native devices of `emcpower` devices and all BCV, GateKeeper, and CKD devices are included in the Excluded Device List. R2 devices and VCMDB devices will not be included in the list. If necessary, edit `exdevtab.sh` in advance or add any other disks to be excluded from GDS management to the list using the above steps 1 through to 4.



Note **Using EMC TimeFinder or EMC SRDF in a Proxy Configuration [Linux2.6]**

For using EMC TimeFinder or EMC SRDF in a proxy configuration, do not describe BCV and R2 devices connected with proxy groups in the Excluded Device List, but describe the native devices that compose those devices in the list. For details, see "A.2.19 Using EMC TimeFinder or EMC SRDF in a Proxy Configuration [Linux2.6]."



Note **For PRIMECLUSTER Systems**

- In a PRIMECLUSTER system, create Excluded Device Lists on all nodes that constitute the cluster.
- Devices that cannot be managed with GDS are also nontargets of resource registration, and do not describe these devices in the shared disk definition file. For details on resource registration and the shared disk definition file, see "Appendix H Shared Disk Unit Resource Registration."



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 Using EMC TimeFinder or EMC SRDF in a Proxy Configuration [Linux2.6]

In a proxy configuration, by using EMC's TimeFinder and SRDF, the copy functions of EMC's Symmetrix storage units can be used and the synchronization copying between master and proxy can be performed without imposing loads on primary servers or a SAN. In this situation, the storage units carry out copying processes and the processes will continue running even if the server is rebooted.

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

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

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

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

- GDS configuration databases cannot be stored in BCV devices and target (R2) devices since the devices are overwritten by data in copy source disks. Therefore, GDS does not regard BCV devices and target (R2) devices as "disks that can be accessed normally" described in "[Explanation]" of "(1) Class becomes closed status during operation" in "F.1.4 Class Status Abnormality."
- 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."
- The BCV, source (R1), and target (R2) devices used in a proxy configuration should not be operated using the SYMCLI commands.
- If the master and proxy are parted forcibly while copying by TimeFinder or SRDF is in process, a message informing device abnormality is submitted to EMC's customer support center.

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

- The soft copy function is specified to be used.
- A pair of volumes instead of a pair of groups is related as a master and a proxy.
- The configuration of the proxy group to which the data are copied is a mirroring configuration.
- The master and the proxy belong to a root class. [Itanium]
- A proxy volume with a different physical slice attribute from the master volume is created into a proxy group.
- A lower level group is connected to the master group or the proxy group.
- A disk of the same size as the disk to which the data are copied is not connected to the master group from which the data are copied.
- Disks that constitute the master or the proxy were registered with a class before installing GDS Snapshot.

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

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

Once the BCV pair or SRDF pair is cancelled, the TimeFinder function or SRDF function is not used. BCV pairs and SRDF pairs are cancelled when:

- Executing [Operation]:[Proxy Operation]:[Restore] in GDS Management View
- Using the `sdxproxy Cancel` command to cancel (break) the BCV or SRDF pair
- Using a soft copy function when joining a master and a proxy with the `sdxproxy Join -e softcopy` command
- Using a soft copy function when rejoining a master and a proxy with the `sdxproxy Rejoin -e softcopy` command
- Restoring data from a proxy back to a master with the `sdxproxy RejoinRestore` command

To make the TimeFinder and SRDF functions available after conducting these operations, break the relationship between the master and the proxy once, remove proxy volumes, and then rejoin the master and the proxy.

To check the modes of the copying in execution, use either:

- The [Copy Type] field of the slice information field in GDS Management View
- The CPTYPE field displayed with the `sdxinfo` command

Additionally, the types and the statuses of BCV pairs and SRDF pairs between the master and the proxy can be viewed in the FUNC field and the CPSTAT field displayed with the `sdxinfo` command.

A.2.20 Ensuring Consistency of Snapshot Data

If snapshots are created while an application is accessing the volume, the snapshots may result from incomplete volume data and the data consistency may not be ensured.

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

For example, if the volume (master) has been used as a file system such as GFS and ext3, before and after creating snapshots, unmount the system and remount it in order to ensure the snapshot data integrity.

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

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

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 normally use the volume in the read and write mode, and temporarily switch to the read only mode, set the access mode attribute to "rw", and use the `sdxvolume -N` command specifying the `-e mode=ro` option to activate the volume in the read only mode temporarily.

The "default access mode" for a shadow volume is ro (read only) and it cannot be changed, but the "current access mode" can be set to either ro (read only) or rw (read and write).



See

- For the method for setting the "default access modes" (access mode attribute values) of logical volumes, see "D.7 sdxattr - Set objects attributes."
- For the method for setting the "current access modes" of logical volumes, see "D.4 sdxvolume - Volume operations."
- For the method for setting the "current access mode" for shadow volumes, see "D.17 sdxshadowvolume - Shadow volume operations."
- The "default access modes" (access mode attribute values) and the "current access modes" of logical volumes and shadow volumes can be checked using the MODE field and the CMODE field displayed with the `sdxinfo -V` command respectively. For details, see "D.6 sdxinfo - Display object configuration and status information."

A.2.23 Operation in Cluster System

Understand and pay attention to the following points when changing the configuration in a cluster system.

- Before registering disks with a class, perform resource registration and register shared disk units with the PRIMECLUSTER resource database.



See

For details on resource registration, see "Appendix H Shared Disk Unit Resource Registration."

- Disks that have not been registered with the resource database yet cannot be registered with a shared class.
- Disks that have not been registered with the resource database yet can be registered with a root class [Itanium] or a local class [Linux2.6].
- When expanding the scope of a local class to which disks not registered with the resource database belong into a shared class, perform resource registration in advance, and register all disks that belong to the local class with the resource database.
- For disks to be registered with a shadow class, disk resource creation is not required.
- Do not register certain disks with the resource database in multiple cluster domains.
- To perform object operations in a cluster system, enable cluster control. If cluster control is "off", it is impossible to perform shared object operations. Additionally, root and local object operations may cause errors or inconsistency.

A.2.24 Changing Over from Single Nodes to a Cluster System [Linux2.6]

The following two methods are available for installing the cluster control facility on one or more single nodes where classes exist and changing over to a cluster system. Normally, the method "a" is used, but the method "b" can be used to skip backup and restore.

- a. Deleting classes and re-creating classes after changing over to a cluster system**
Back up volume data if necessary, and delete the root class and local classes. After installing and configuring the cluster control facility, re-create the classes and volumes, and restore volume data as needed.
- b. Changing over to a cluster system without deleting classes**
Change over from single nodes to a cluster system without deleting classes. However, you must understand and pay attention to the following points.

Important Point 1

A local class created on a single node cannot be used directly in a cluster system. When the cluster control facility is activated, the following error message is output to the system log and the GDS daemon log file, and the operation of the local class becomes unavailable.

ERROR: *class*: cannot operate in cluster environment, created when cluster control facility not ready

For details on resolution, see (1) in "F.1.10 Cluster System Related Error."

Important Point 2

Expanding the class scope after changing over from multiple single nodes to a cluster system, may output the following messages.

ERROR: *class*: class names must be unique within a domain

This error occurs when the name of a class created on a single node is the duplicate name of a class on another node. If this error occurs, rename either of the classes, and expand the class scope.

ERROR: *class*: volume minor numbers must be unique within a domain

This error occurs when the minor number of a volume created on a single node is the duplicate number of a volume on another node. If this error occurs, re-create either of the volumes, and expand the class scope.

The minor number of a volume can be viewed in the following manner.

```
# cd /dev/sfdsk/class/dsk
# ls -l
brw----- 1 root    root    253, 33 May  6 09:00 volume1
                ^^
```

Additionally, this error may occur when any lower level group exists in a class created on a single node or a class on another node. In this event, duplicate minor numbers cannot be checked with the method as shown above. Re-create all volumes and lower level groups in the class created on the single node and then expand the class scope.

A.2.25 Disk Switch

The disk switch function is available only when using an application that controls the disk switch function. Unless the application's manual instructs to creation of switch groups and switch volumes, do not create switch groups.

A.2.26 Shadow Volume

Rebooting a Node

The configuration information of a shadow volume is not saved on the private slice but managed in the memory. For this reason, the shadow volume configuration definitions are cleared when the node on which the shadow volume is defined is rebooted. However, the device special file remains. If such a device special file is left not deleted, issues as described below may occur.

Before intentional shutdowns, it is recommended to remove shadow volumes. If a shadow volume is removed with the `sdxshadowvolume -R` command, the device special file is also deleted. For details on the `sdxshadowvolume -R` command, see "D.17 sdxshadowvolume - Shadow volume operations."

When a node is shut down leaving the relevant shadow volume not removed, or if a node on which a shadow volume is defined is rebooted unexpectedly because of an event such as a panic and a power cutoff, the device special file for the shadow volume must be deleted in the following procedure.

[Procedure]

1. Check the system for existing classes.

In the following example, there are `RootClass`, `Class1`, and `Class2`.

```
# sdxinfo -C
OBJ      NAME      TYPE      SCOPE      SPARE
-----  -
class    RootClass  root      (local)    0
class    Class1     local     node1      0
class    Class2     shared    node1:node2 0
```

2. Find the directories containing the device special files of classes.

In the following example, `RootClass`, `Class1`, and `Class2` are the directories for the device special files of those existing classes, and `_adm` and `_diag` are the special files used by GDS. `Class3`, other than those directories, is the directory for the device special file of the extinct shadow class.

```
# cd /dev/sfdsk
# ls -F
Class1/   Class2/   Class3/   RootClass/  _adm@    _diag@
```

3. Delete the directory for the device special file of the extinct shadow class.

```
# rm -r Class3
```

Even if the device special file of an extinct shadow volume remains, no problem will arise if a shadow volume in the same configuration, of the same class name, and with the same volume name is re-created.

Otherwise, the following issues will occur. If a logical volume or a shadow volume is created in the situation that the device special file, `/dev/sfdsk/Shadow Class Name/[r]dsk/Shadow Volume Name`, of an extinct shadow volume remains, the minor number of the created volume may become the same as the minor number of `/dev/sfdsk/Shadow Class Name/[r]dsk/Shadow Volume Name`. In this situation, if `/dev/sfdsk/Shadow Class Name/[r]dsk/Shadow Volume Name` is accessed without recognition of extinction of the shadow volume, the newly created volume is accessed, and it can cause an application error and corruption of data on the newly created volume.

Accessing a Shadow Volume

Shadow volumes and the corresponding logical volumes are managed independently. For example, the change of the slice status in one volume is not updated in the slice status in the other volume. For this reason, you must note the following operational particulars when using shadow volumes.

Synchronization of Shadow Volumes

When a shadow volume is created in another domain (domain beta) for the disk area managed as a mirror volume in a certain domain (domain alpha), the mirror volume in domain alpha and the shadow volume in domain beta cannot be accessed simultaneously. If they are accessed simultaneously, the following issues arise.

- If an I/O error occurs in the slice configuring the mirror volume in domain alpha, that slice becomes INVALID and is detached from the mirror volume. However, GDS in domain beta does not detect this I/O error, and consequently the shadow slice is not made INVALID and is not detached from the shadow volume. Here, synchronization of the shadow volume is not ensured.
- Likewise, if an I/O error occurs in the shadow slice in domain beta, the slice in the corresponding mirror volume in domain alpha is not made INVALID and is not detached from the mirror volume. Here, synchronization of the mirror volume is not ensured. If an I/O error occurs on the shadow slice, working around, such as swapping the disks and resynchronization copying of the mirror volume, is required in domain alpha.

These particulars apply when the disk area for a mirror volume and a shadow volume are identical. A mirror volume and a shadow volume corresponding to a replica of the mirror volume (a temporarily detached slice, a proxy volume or a copy destination disk area for a disk unit's copy function) can be accessed simultaneously.

Just Resynchronization Mechanism (JRM) for Volumes

When a shadow volume is created in another domain (domain beta) for the disk area managed as a mirror volume in a certain domain (domain alpha) and accessed, the following must be set up for the mirror volume in domain alpha.

- Mirror volumes must be inactivated to prevent access to the mirror volume corresponding to the shadow volume.
- JRM for volumes must be enabled ("on") for the mirror volume corresponding to the shadow volume.

These settings are necessary for the following reasons.

If a node in domain alpha panics and resynchronization copying is conducted on the mirror volume in domain alpha while the shadow volume is accessed in domain beta, synchronization between the shadow volume and the mirror volume is no longer ensured. Though the settings as above, resynchronization copying is no longer conducted on the mirror volume in domain alpha even if a node in domain alpha panics.

The settings as above are necessary only for a mirror volume created for the disk area identical to the shadow volume's disk area. When a shadow volume corresponding to a replica of a mirror volume (a temporarily detached slice, a proxy volume or a copy destination disk area for a disk unit's copy function) is created, these settings are not necessary for the copy source mirror volume.



Information Resynchronization Copying after Panic

Resynchronization copying is not conducted after panic when JRM for volumes is enabled ("on") and that volume is not written in. Resynchronization copying occurs after panic when JRM for volumes is disabled ("off") and that volume is active.

Just Resynchronization Mechanism (JRM) for Slices

When a slice is temporarily detached from a mirror volume in a certain domain (domain alpha) and data is written from a shadow volume in another domain (domain beta) to the area of this volume or slice, JRM for slices must be disabled ("off") prior to reattaching the slice.

If JRM for slices is enabled ("on"), the following issue arises.

When JRM for slices is enabled ("on"), only the difference between the volume and the slice is copied by reattaching the slice. The difference information for the volume and the slice is managed by JRM for slices in domain alpha. However, JRM for slices in domain alpha does not recognize write events from domain beta, and the difference resulting from data being written from domain beta are not updated in the difference information. The difference resulting from write events from domain beta, therefore, are not copied when the slice is reattached while JRM for slices is "on" in domain alpha. As a result, synchronization of the volume is no longer ensured.

Just Resynchronization Mechanism (JRM) for Proxies

If a proxy volume is parted from the master in a certain domain (domain alpha) and data is written from a shadow volume in another domain (domain beta) to the area of this master or proxy, JRM for proxies must be disabled ("off") prior to rejoining the proxy. In addition, JRM for proxies must be disabled ("off") prior to restoring the master using the proxy.

If JRM for proxies is enabled ("on"), the following issues arise.

When JRM for proxies is enabled ("on"), only the difference between the master and the proxy is copied by rejoining or restoring. The difference information for the master and the proxy is managed by JRM for slices in domain alpha. However, JRM for proxies in domain alpha does not recognize write events from domain beta, and the difference resulting from data being written from domain beta are not updated in the difference information. The difference resulting from write events from domain beta, therefore, are not copied when the proxy is rejoined or the master is restored while JRM for proxies is "on" in domain alpha. As a result, synchronization between the master and the proxy is no longer ensured.

When one of a disk unit's copy function (EC, REC, TimeFinder, and SRDF) with a resynchronization feature based on equivalent copy capability is used for master-to-proxy copy processes, data written from domain beta is also updated in the difference information managed by these disk unit's copy functions. Under these circumstances, JRM for proxies do not have to be disabled ("off") prior to rejoining. Note, however, that JRM for proxies must be disabled ("off") prior to restoring since necessity of synchronization copying is determined based on the difference information managed by JRM for proxies. To ensure the data integrity, it is recommended to disable JRM for proxies prior to rejoining even when a disk unit's copy function is used.



Information A Copy Function of a Disk Unit with Resynchronization Feature Based on Equivalent Copy

When just resynchronization copying from a master to a proxy is conducted with one of a disk unit's copy functions (EC, REC, TimeFinder, SRDF) with a resynchronization feature based on equivalent copy capability, this feature is used regardless of whether JRM for proxies is "on" or "off."

Writing into Shadow Volumes

Data may be written to a shadow volume even if the operation for writing is not especially intended. For example, executing mount(8) (excluding when using the -o ro option), fsck(8) or mkfs(8) results in the write operation.

When a proxy is rejoined, a master is restored, or a slice is reattached once a shadow volume is created, it is recommended to disable the just resynchronization mechanism mode (JRM) regardless of whether or not data is written into the shadow volume in order to ensure the data integrity.

A.2.27 Configuration Tables [Linux2.6]

Do not use editors such as vim(1) and sed(1) to edit configuration tables created with the sdxconfig Backup command or those saved in configuration files. To edit configuration tables, use the sdxconfig Convert command.

A.2.28 GDS Management View

Physical Disk Recognition

When any operation that changes the physical disk configuration, such as addition or deletion of disk units, is conducted during system operation, update physical disk information with new information. Execute [Update Physical Disk Information] on the [Operation] menu when:

- The power of the disk unit was turned on after the system was booted.
- The disk unit became unavailable for some kind of problem during system operation, but was recovered without system reboot.
- The configuration of devices was changed.

Object Status Monitoring

Objects indicated as failed in GDS Management View are only those in which GDS detected errors. Even if a hardware error occurs in a disk unit, the disk unit status is indicated as normal until the disk unit is accessed and the error is detected.

A.2.29 File System Auto Mount

File systems created on volumes in classes other than the root class (local classes and shared classes) cannot be mounted with the OS auto mount feature at OS startup.

This is because the OS auto mount process is executed before the GDS startup script `"/etc/*.*sfdsk*"` is executed.

When including file systems created on volumes in classes other than the root class in the `/etc/fstab` file, be sure to describe the `"noauto"` option in the fourth field. In the File System Configuration screen of GDS Management View, select `"no"` to `"mount."` If `"yes"` to `"mount"` is selected, the `noauto` option is not declared.

If the `noauto` option is not declared, the following message is displayed at OS startup and mounting fails.

```
mount: special device /dev/sfdsk/Class Name/dsk/Volume Name does not exist
```

Additionally, if the `noauto` option is not declared and if a specification to check file systems at OS startup is made in the sixth field, the following message is displayed and the OS does not start.

```
fsck.ext3: No such file or directory /dev/sfdsk/Class Name/dsk/Volume Name:
```



Information For Shared Classes

To mount or unmount file systems on volumes in shared classes when starting or exiting cluster applications, `/etc/fstab` file setting and Fsystem resource setting are necessary. When using GFS, see "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide." For details, see "PRIMECLUSTER Global File Services Configuration and Administration Guide."



Information For Local Classes [Linux2.6]

To automatically mount file systems created on volumes in local classes at OS startup, use a startup scrip such as `/etc/rc.local` that is executed after the GDS startup script `"/etc/*.*sfdsk*"` for mounting.

A.2.30 Raw Device Binding

For using a volume or a temporarily detached slice as a character (raw) device, use the raw device bound to the volume's or the temporarily detached slice's block device with the raw(8) command.

Example) Bind raw device "raw1" to the block device of Volume1 in Class1.

```
# raw /dev/raw/raw1 /dev/sfdsk/Class1/dsk/Volume1
```

By describing the setting for raw device bounding in the /etc/sysconfig/rawdevices file, the raw device can automatically be bound at OS startup. For details, see the raw(8) and rawdevices manuals.

If the setting for raw device bounding for a volume or a temporarily detached slice is described in the rawdevices file, the following message may be output at OS startup. Even if this message is output, the raw device will properly be bound after GDS initialization and there is no effect on the system.

```
Cannot locate block device '/dev/sfdsk/Class Name/dsk/Volue Name' (No such file or directory)
```



Information Volume and Temporarily Detached Slice raw Devices [Linux2.4]

If the Linux kernel version is 2.4, it is also possible to use a volume's raw device "/dev/sfdsk/Class Name/rdsk/Volume Name" or a temporarily detached slice's raw device "/dev/sfdsk/Class Name/rdsk/Disk Name.Volume Name".

A.2.31 Upgrading Instructions

Device Special File Incompatibility (GDS)

The following incompatibility exists between PRIMECLUSTER GDS 4.1A30 or before and 4.1A40 or later.

- PRIMECLUSTER GDS 4.1A30 or before and 4.2A00 [Linux2.4]
Device special files for volumes and temporarily detached slices are either the block device type or the character (raw) type.
 - Device special files for volumes

Block : /dev/sfdsk/Class Name/dsk/Volume Name
Character (raw): /dev/sfdsk/Class Name/rdsk/Volume Name
 - Device special files for temporarily detached slices

Block : /dev/sfdsk/Class Name/dsk/Disk Name.Volume Name
Character (raw): /dev/sfdsk/Class Name/rdsk/Disk Name.Volume Name
- PRIMECLUSTER GDS 4.1A40 or later [Linux2.6]
Device special files for volumes and temporarily detached slices are all block device special files.
 - Device special files for volumes

Block : /dev/sfdsk/Class Name/dsk/Volume Name
 - Device special files for temporarily detached slices

Block : /dev/sfdsk/Class Name/dsk/Disk Name.Volume Name

To use a volume or a temporarily detached slice as a character (raw) device, bind the raw device with the raw(8) command. For details, see "A.2.30 Raw Device Binding."

GDS Management View Information Message Incompatibility (GDS)

Based on the device special file incompatibility, the following incompatibility exists between information messages for slice detachment completion displayed from GDS Management View.

- PRIMECLUSTER GDS 4.1A30 or before

5202 The slice was detached. The following access path is available. Access path: <i>/dev/sfdsk/classname/rdsk/slicename</i>

- PRIMECLUSTER GDS 4.1A40 or later

5202 The slice was detached. The following access path is available. Access path: <i>/dev/sfdsk/classname/dsk/slicename</i>
--

Command Information Message Incompatibility (GDS)

Based on the device special file incompatibility, the incompatibilities between command information messages below exist. All message lines are displayed in GDS 4.1A30 or before and 4.2A00 [Linux2.4], but the italic text strings (the second lines of the following messages) are not displayed in GDS 4.1A40 or later [Linux2.6].

INFO: <i>volume</i> : started volume on node <i>node</i> <i>/dev/sfdsk/class/rdsk/volume</i> <i>/dev/sfdsk/class/dsk/volume</i>

INFO: <i>disk.volume</i> : enabled slice <i>/dev/sfdsk/class/rdsk/disk.volume</i> <i>/dev/sfdsk/class/dsk/disk.volume</i>

INFO: <i>disk.volume</i> : started slice on node <i>node</i> <i>/dev/sfdsk/class/rdsk/disk.volume</i> <i>/dev/sfdsk/class/dsk/disk.volume</i>

INFO: <i>disk.volume</i> : current node <i>node</i> took over and started slice <i>/dev/sfdsk/class/rdsk/disk.volume</i> <i>/dev/sfdsk/class/dsk/disk.volume</i>
--

INFO: <i>volume</i> : started volume on node <i>node</i> ,... <i>/dev/sfdsk/class/rdsk/volume</i> <i>/dev/sfdsk/class/dsk/volume</i>
--

Command Error Message Incompatibility (GDS)

The following incompatibility exists between PRIMECLUSTER GDS 4.1A30 and 4.1A40 or later.

4.1A30	4.1A40 or later
ERROR: <i>group</i> : group other than switch group exists in <i>class</i>	ERROR: <i>class</i> includes a group that cannot exist together with a switch group

sdxproxy Command's -e delay Option Incompatibility (GDS Snapshot)

The following incompatibility exists between PRIMECLUSTER GDS Snapshot 4.1A20 and 4.1A30 or later.

- PRIMECLUSTER GDS Snapshot 4.1A20
Even if delay time that is out of range of 0 to 1000 in milliseconds is specified with the -e delay option of the sdxproxy command, the sdxproxy command does not result in an error and the specified delay time will be in effect.
- PRIMECLUSTER GDS Snapshot 4.1A30 or later
If delay time that is out of range of 0 to 1000 in milliseconds is specified with the -e delay option of the sdxproxy command, the sdxproxy command results in an error and the following error message is output. *msec* indicates the delay time specified with -e delay=*msec*.

ERROR: <i>msec</i> : invalid delay value
--

A.3 General Points

A.3.1 Guidelines for Mirroring

Pay attention to the following guidelines when constructing mirroring configurations.

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

A.3.2 Guidelines for Striping

Pay attention to the following guidelines when constructing striping configurations.

- In order to improve I/O performance with striping, it is necessary to adjust the stripe width and the number of stripe columns depending on the way an application accesses the disk.
If the striping configuration is not appropriate, you cannot gain much performance improvement. And, depending on the way an application accesses the disk, the performance may not improve even after adjusting the stripe width or the number of stripe columns.
- Do not make the stripe widths too large.
The sizes of stripe groups and stripe volumes are rounded to the common multiple of stripe width times the number of stripe columns and cylinder size. Therefore, if the stripe width is too large, use of the disk area may be inefficient or a volume with the intended size may be created.
- Where possible, connect disks and lower level groups with the same available sizes to the same stripe group.
The available size of the stripe group (available capacity as volumes) equals the available size of the smallest disk or the lower level group connected to the stripe group multiplied by the number of stripe columns and rounded down to the common multiple of stripe width times the number of stripe columns and cylinder size.
When connecting disks or lower level groups with different available sizes to a stripe group, the larger disk or lower level group will only be able to use the capacity of the smaller disk or lower level group. For example, if a 4 GB disk and a 9 GB disk are connected to one stripe group, the 9 GB disk will only be able to use approximately 4 GB. Which means, the available size of stripe group will be approximately 8 GB (4 GB x 2).
- Where possible, striping across disks with similar performance specifications is recommended.
When striping disks with different performance specifications such as revolution speed, the performance becomes unbalanced and will depend on the slower disk performance.
- Using striping in combination with mirroring is recommended.
In a striping configuration, the risk of losing data from a disk failure increases as more disks are involved compared to a usual disk configuration.
By mirroring stripe groups, both the I/O load balancing and data redundancy can be achieved at the same time.

A.3.3 Guidelines for Concatenation

Pay attentions to the following guidelines when constructing concatenation configurations.

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

A.3.4 Guidelines for Combining Striping with Mirroring

Pay attention to the following guidelines when striping and mirroring simultaneously.

- Where possible, mirroring stripe groups with similar configuration is recommended. See "A.3.1 Guidelines for Mirroring" and "A.3.2 Guidelines for Striping" also.

Appendix B Log Viewing with Web-Based Admin View

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

Appendix C Web-Based Admin View Operating Environment Setting

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

Appendix D Command Reference

This appendix discusses the commands provided by GDS and GDS Snapshot. This appendix explains the format and facility of commands, specifiable options, and return values. GDS provides the following commands.

Command	Function
sdxcass	Class operations
sdxdisk	Disk operations
sdxgroup	Group operations
sdxvolume	Volume operations
sdxslice	Slice operations
sdxinfo	Display object configuration and status information
sdxattr	Change attribute values of an object
sdxswap	Swap disk
sdxfix	Restore a crashed object
sdxcopy	Synchronization copy operations
sdxroot	Root file system mirroring definition and cancellation [Itanium]
sdxparam	Configuration parameter operations
sdxconfig	Object configuration operations [Linux2.6]

GDS Snapshot provides the following commands.

Command	Function
sdxproxy	Proxy object operations
sdxshadowdisk	Shadow disk operations
sdxshadowgroup	Shadow group operations
sdxshadowvolume	Shadow volume operations



Information Commands That Operate Multiple Objects

When an error occurs in operation for part of the objects, the command may either continue operation for the other objects or terminate the process.

In either situation, referring to "Appendix E GDS Messages" check the meaning of the error message and take necessary action.

D.1 sdxclass - Class operations

SYNOPSIS

`sdxclass -R -c class`

DESCRIPTION

Use `sdxclass` to perform operations on class objects (excluding shadow class) specified by *class*. You must be superuser to use this command.

PRIMARY OPTIONS

You can use the following option.

-R

Remove

Removes the class definition specified by *class*.

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

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

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

For further details, see "D.2 sdxdisk - Disk operations."

SUB OPTIONS

Sub options are as follows:

-c class

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

RETURNED VALUE

Upon successful completion, a value of 0 is returned.

Otherwise, a non-zero value is returned.

D.2 sdxdisk - Disk operations

SYNOPSIS

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

DESCRIPTION

Use sdxdisk to perform operations on disk objects (excluding shadow disk) specified by *disk*. You must be superuser to use this command.

PRIMARY OPTIONS

You can use either of the following options.

-C

Connect

Connects one or more disks (keep disks, single disks, or undefined disks) specified by *disk*,... to a group specified by *group*. The *class* indicates the class name with which *disk* is registered. To connect *disk* to a group other than a switch group, specify a disk connected to all the nodes that belong to the scope of *class*.

If no group with the same name as *group* exists, a new group is created.

Spare disks cannot be connected to groups. Keep disks and single disks cannot be connected to existing groups. Additionally, a keep disk and a single disk, multiple keep disks, or multiple single disks cannot be connected to one group together.

The disk attribute will be changed to match the attribute of the group (mirror, stripe, concatenation, or switch) it is connected to. Disks and lower level groups connected to a group are mirrored, striped, concatenated or made switchable according to the type attribute.

Details about connecting disks to a mirror group, a stripe group, a concatenating group, and a switch group are explained below.

- When connecting to a mirror group

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

When one or more volumes already exist within the mirror group specified by *group*, the slice configuration of disks or lower level groups that are already connected to *group* will be automatically copied to the newly connected *disks*.

Also, when there is an activated volume within *group*, volume data in addition to the slice configuration will be automatically copied upon returning from the *sdxdisk* command, therefore increasing the mirroring multiplicity.

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

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

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

When a keep disk is connected, volumes are created for every physical slice within the keep disk. If the physical slices are not open, the created volumes are started and synchronization copying is performed automatically after returning from the *sdxdisk* command.

In classes that include switch groups, mirror groups cannot be created.

- When connecting to a stripe group

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

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

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

You cannot connect a single disk to a stripe groups.

In classes that include switch groups, stripe groups cannot be created.

- When connecting to a concatenation group

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

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

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

You cannot connect a single disk to a concatenation group.

In classes that include switch groups, concatenation groups cannot be created with this command. Additionally, to concatenation groups to which lower level switch groups are connected, disks cannot be added.

- When connecting to a switch group

One of two disks connected to a switch group becomes the active disk and the other one becomes the inactive disk. Use the `sdxattr -G` command to switch the disks from active to inactive and vice versa. A switch group can connect a maximum of two disks. If only one disk is connected, the disk acts as the active disk and an active disk switch cannot be performed.

To create a switch group specified by *group*, specify the active disk in the `-a actdisk` option. The other disk not specified by the `-a actdisk` option becomes the inactive disk. When connecting *disk* to an existing switch group specified by *group*, the *disk* becomes the inactive disk.

Unlike mirror groups, even if a switch group specified by *group* already includes volumes, synchronization copying to the newly connected *disk* is not performed. To perform a disk switch for continuous services in the event of an error in the active disk, previously create copies of data from the active disk to the inactive disk with the disk unit's copy function and so on.

The available size of a switch group (capacity available for volumes) conforms to the available size of the smallest disk connected to the switch group. If the available size of *group* decreases as a result of *disk* connection, a warning message is sent to standard error output.

In classes that include any one of the following objects, switch groups cannot be created.

- Disk other than an undefined disk
- Mirror group
- Stripe group
- Concatenation group to which no lower switch group is connected

class must be a shared class of which scope includes 2 nodes. The physical scope of the active disk and the inactive disk must also meet either the following conditions.

- The active disk and the inactive disk are connected to both the nodes included in the scope of *class* and are not connected to nodes not included in that scope.
- The active disk is connected to only one of the nodes included in the scope of *class* and the inactive disk is connected to the other node included in that scope.

-D

Disconnect

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

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

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

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

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

- When disconnecting from a mirror group

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

- When disconnecting from a stripe group

A disk cannot be disconnected from a stripe group with one or more existing volumes, or from a stripe group connected to a higher level group.

- When disconnecting from a concatenation group

The only disk that can be disconnected from a concatenation group is the disk that was concatenated last.

Disks containing volume areas cannot be disconnected from a concatenation group.

If the concatenation group specified by *group* is connected to a stripe group that is connected to a mirror group, disks cannot be disconnected.

- When disconnecting from a switch group

Inactive disks can be disconnected regardless whether or not volumes exist.

The active disk can be disconnected from a switch group if all the following conditions are satisfied.

- The switch group is not connected to a higher level concatenation group.
- The switch group includes no volume.
- The inactive disk is not connected to the switch group.

If the switch group includes volumes, before disconnecting the active disk, remove those volumes. If the switch group includes the inactive disk, switch the active disk to it with the `sdxattr -G` command and then disconnect the former active disk.

-M

Make

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

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

A root type *class* can include *device* of the keep type. However, if multiple keep type *devices* are to be registered together, as many or more undef type *devices* must be registered.



Note

Since the `sdxdisk` command initializes the registered physical disks (excluding *devices* with "keep" assigned as the type attributes), when registering a physical disk containing data, you must first create data backup.

-R

Remove

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

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

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

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

If removal of *disk* will result in *class* closure, the *disk* cannot be removed.

The *class* will be closed when it includes:

- less than 3 disks in ENABLE status and no disk normally accessible
- three to 5 disks in ENABLE status and less than 2 disks normally accessible
- six or more disks in ENABLE status and less than 3 disks normally accessible

SUB OPTIONS

Sub options are as follows:

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

When using the -C option and defining a new group name with the -g option, a new group is automatically created. Using this option sets the *attribute* of the created *group* to *value*.

The *attribute* indicates the attribute name, and *value* indicates the attribute value. Always separate *attribute* and *value* with an equal (=) sign. When indicating multiple attributes, specifiers should be combined using commas(,) as the delimiter.

If no group is created, indicating a different attribute value from the existing *group* will result in an error. To change the attribute value of an existing *group*, use the `sdxattr -G` command.

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

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

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

Sets the type attribute of *group*. If *class* is the root class, specifying "stripe" or "concat" will result in an error. If *class* is not a shared class of which scope includes 2 nodes, specifying "switch" will also result in an error.

mirror

Sets type attribute to "mirror."

stripe

Sets type attribute to "stripe."

concat

Sets type attribute to "concatenation."

switch

Sets type attribute to "switch."

width=*blks* (default is 32)

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

actdisk=*disk*

Sets the active disk of *group*. Specify a disk name of the active disk into *disk*. When *group* is an existing group, not specifying the -a type=switch option, or specifying a disk other than that specified by -d option into *disk*, will result in an error.

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

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

The *attribute* indicates the attribute name, and the *value* indicates the attribute value. Always separate *attribute* and *value* with an equal (=) sign. Specifiers should be combined using commas(,) as the delimiter.

If no class is created, this option is ignored.
To change the attributes of an existing class, use the `sdxattr -C` command.

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

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

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

Sets the attributes of *class* type.

root [Itanium]

Sets the type attribute to "root."

Objects within *class* of the root type can be used only on the current node.

Unlike with local type classes, system disks including the root file system can be registered with the root type *class*.

Only one root type class can be created within the system. If a class of the root type already exists, specifying another *class* of the root type will result in an error.

For the scope attribute, the node identifier of the current node is set automatically.

local [Linux2.6]

Sets the type attribute to "local."

Objects within *class* of the local type can be used only on the current node.

For the scope attribute, the node identifier of the current node is set automatically.

shared

Sets the type attribute to "shared."

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

A shared type class can include physical disks connected to all the nodes that belong to the scope. When the scope includes 2 nodes, disks connected to only one node in the scope can be registered as undefined disks.

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

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

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

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

It sets the operation of the hot spare.

on

Enables the hot spare.

off

Disables the hot spare. If the operation mode is set to off, spare disk automatic connection is restrained.

hsmode=exbox or **hsmode=bybox** (default is **exbox**)

Sets the spare disk selection mode for automatic connection by hot spare.

exbox

Sets the spare disk selection mode to the external mode.

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

bybox

Sets the spare disk selection mode to the internal mode.

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

-c class

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

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

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

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

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

The physical disk name can be specified in either the following formats:

sdX	(for normal hard disks)
emcpowerX	(for emcpower disks)

X indicates the device identifier.

The *type* can be indicated as one of the following. If no type is indicated, the default value of undef (undefined disk) is used. If *class* is not the root class, specifying "keep" will result in an error. If *device* is not connected to part of nodes included in the scope of *class*, or if a switch group exists in *class*, specifying a type other than "undef" will also result in an error.

keep [Itanium]

Keep disk. When it is registered with a class or connected to a group, the format and data of the disk will be retained.

single

Single disk. Single volumes can be created on it.

spare

Spare disk.

undef

Undefined disk.

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

When only one *device* is specified with "keep" assigned as its *type*, the specified *device* must have a reserved physical slice number and sufficient free disk space so that the private slice can be created on the *device*. When multiple *devices* are specified, *devices* with "keep" assigned to *type* must have reserved physical slice numbers and sufficient free disk space or swap devices with sufficient disk space.

With a system disk (disk with running */*, */usr*, */var*, */boot*, or */boot/efi* or a swap area) with "keep" assigned as its *type*, even if all or some of the physical slices are currently open, the *sdxdisk* command ends normally. However, if "keep" is specified for a disk other than a system disk, this command will result in an error where open physical slices exist. For example, if any of the physical slices are being used as a file system, unmount the file system to free up the physical slice, and then execute the *sdxdisk* command.

When "single" is specified for *type*, *device* will be registered as a single disk. For a "single" disk, you can create single volumes on it using the *sdxvolume* command without connecting the disk to any group.

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

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

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

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

-e *chkps* (when using **-M**)

Registers *device* with *class* even if the private slice exists in the *device*, as far as disk identification information (class and disk names) stored in the private slice matches identification information of a disk already registered with the *class*. For example, if *device* contains a copy of the private slice of a disk that is already registered with *class*, to register the *device* to the *class*, turn on this option.

If *class* is not a shared class, this command results in an error.

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

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

This option delays the issuing of the input/output request to the disk at the time of copying by milliseconds specified by *msec*, allowing adjustment for the influence on the application accessing the volume.

The value is set to 0 by default.

Values from 0 to 1000 may be specified for *msec*.

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

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

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

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

Specifies the created volume's attribute value when connecting *disk* of the keep type. This option setting is simply ignored if *disk* of the keep type is not specified.

Always use an equal sign (=) after *num*, and separate *volume* and *jrm* with a colon (:). When specifying the attributes values of multiple volumes, specifiers should be combined using commas (,) as the delimiter.

Specify the physical disk slice number (an integer 1 to 15) of a keep type *disk* storing volume data for *num*, the volume name for *volume*, and the created volume's just resynchronization mode ("on" or "off") for *jrm*.

When the keep type *disk* contains multiple physical slices of which size is nonzero, it is necessary to specify the corresponding volume attribute values for all the physical slices.

RETURNED VALUE

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

D.3 sdxgroup - Group operations

SYNOPSIS

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

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

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

DESCRIPTION

Use sdxgroup to perform operations on group objects (excluding shadow groups). You must be superuser to use this command.

PRIMARY OPTIONS

You can use one of the following options.

-C

Connect

Connects one or more groups (stripe groups, concatenation groups, or switch groups) specified by *lgroup*,... to a group (a mirror group, stripe group, or concatenation group) specified by *hgroup*. The *class* indicates the name of the class to which *lgroup* belongs. If *class* is the root class, this command results in an error.

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

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

Lower level groups and disks connected to the same higher level group are mirrored, striped, or concatenated according to the type attribute of the higher level group. Connecting a group to a higher level group does not change the type attribute of the lower level group.

You cannot connect groups when:

- *lgroup* is a mirror group
- *hgroup* is a switch group
- type attributes of *lgroup* and *hgroup* are the same

A group that already contains volumes cannot be connected to another group.

Details about connecting groups to a mirror group, a stripe group, and a concatenation group are explained below.

- When connecting to a mirror group

You can connect one or more groups (stripe group or concatenation group) specified by *lgroup*,... to *hgroup* which is a mirror group.

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

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

The available size of the mirror group (available capacity as volumes) will be the same as the available size of the smallest disk or the lower level group connected to the mirror group. If connecting *lgroup* decreases the available size of *hgroup*, a warning message will be sent to standard error output.

In classes that include switch groups, mirror groups cannot be created.

- When connecting to a stripe group

You can connect one or more groups (concatenation group) specified by *lgroup*,... to *hgroup* which is a stripe group. Groups specified by *lgroup*,..., will be connected to *hgroup* in the order they are listed.

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

When a stripe group specified by *hgroup* already exists, stripe columns will be added after the stripe columns that already exist in *hgroup*, in the order they are specified by *lgroup*,... However, you cannot increase the stripe columns by connecting groups to stripe groups with volumes, or to stripe groups connected to a higher level group.

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

In classes that include switch groups, stripe groups cannot be created.

- When connecting to a concatenation group [Linux2.6]

This command can connect one or more groups (switch groups) specified by *lgroup*,... to *hgroup* which is a concatenation group. Switch groups connected to the same concatenation group will be concatenated in the order they are specified in *lgroup*,.... Concatenation of a maximum of 64 groups is supported.

The available size (available capacity as volumes) of a concatenation group equals the total available size of lower level groups connected to the concatenation group.

By connecting lower level groups to an existing concatenation group, the available size of the concatenation group can increase. If the concatenation group specified by *hgroup* already exists, lower level groups are concatenated in the order they are specified in *lgroup*,... following the last concatenated lower level group in *hgroup*. However, to concatenation groups connected to higher level groups, lower level groups cannot be connected.

To concatenation groups to which disks are connected, switch groups cannot be connected.

-D

Disconnect

Disconnects group specified by *lgroup* from the higher level group *hgroup*. The *class* indicates class name to which *lgroup* belongs, and *hgroup* indicates the higher level group name to which *lgroup* is connected.

When *lgroup* is the only object connected to *hgroup*, *hgroup* will automatically be removed upon disconnecting *lgroup*. However, when *lgroup* is the only object connected to *hgroup*, and *hgroup* is connected to a higher level group, disconnection will result in an error. In such case, disconnect *hgroup* from its higher level group, and then disconnect *lgroup*.

You cannot disconnect *lgroup* if the disconnection may result in a change in the status of any existing volume within *hgroup*.

Restrictions that prevent group disconnection from a mirror group, a stripe group, and a concatenation group are explained below.

- When disconnecting from a mirror group

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

- When disconnecting from a stripe group

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

- When disconnecting from a concatenation group [Linux2.6]

Only the last concatenated lower level group can be disconnected from a concatenation group.

Lower level groups that have volume areas cannot be disconnected from concatenation groups.

-R

Remove

Remove the group definition specified by *group*. The *class* indicates the class name to which *group* belongs.

Disks and lower level groups connected to *group* will be disconnected. The disconnected disk's attribute will return to its original setting (keep disk, single disk, or undefined disk).

group cannot be removed when one or more volumes exist within *group*, or when *group* is connected to a higher level group.

SUB OPTIONS

Sub options are as follows:

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

When using the -C option and defining a new group name with the -h option, a new group *hgroup* is automatically created. Using this option sets the *attribute* of the created *hgroup* to *value*.

The *attribute* indicates the attribute name, and *value* indicates the attribute value. Always separate *attribute* and *value* with an equal (=) sign. When indicating multiple attributes, specifiers should be combined using commas(,) as the delimiter.

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

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

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

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

Sets the type attribute of *hgroup*.

mirror

Sets the type attribute to "mirror."

stripe

Sets the type attribute to "stripe."

concat [Linux2.6]

Sets the type attribute to "concatenation."

width=blks (default is 32)

Sets the stripe width of *hgroup*. The *blks* indicates the stripe width in block number (base 10).

One block is 512 bytes. For *blks*, you can indicate an integer that is two raised to the power (one and larger), which is equal to or smaller than the available size of the smallest group specified by *lgroup*,... If *hgroup* is not a stripe group, this option will result in an error.

-c class

The *class* indicates the class name to which the group belongs, where *group* is the target of the operation.

-e delay=msec (when using -C)

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

This option delays the issuing of the input/output request to the disk by milliseconds specified by *msec*, allowing adjustment for the effect on the application accessing the volume. Default is 0. Values from 0 to 1000 may be specified for *msec*.

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

-g group (when using -R)

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

-h hgroup (when using -C,-D)

The *hgroup* indicates the higher level group name to which the lower level group is connected or is to be connected, where the lower level group is the target of the operation.

-l lgroup (when using -D)

The *lgroup* indicates the lower level group name that is the target of the operation.

-l lgroup,... (when using -C)

The *lgroup* indicates the lower level group name that is the target of the operation. To connect multiple groups, separate each group name with a comma (,) as the delimiter.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.

Otherwise, a non-zero value is returned.

D.4 sdxvolume - Volume operations

SYNOPSIS

```
sdxvolume -F -c class [-v volume,...]
                [-e {allnodes|node=node[:node:...]}]

sdxvolume -M -c class {-g group|-d disk } -v volume -s size
                [-a attribute=value[,attribute=value...]][-e delay=msec]

sdxvolume -N -c class [-v volume,...]
                [-e [allnodes|node=node[:node:...]],delay=msec, mode=val, nosync, unlock]

sdxvolume -R -c class -v volume

sdxvolume -S -c class -v volume -s size
```

DESCRIPTION

Use sdxvolume to perform operations on volume objects (excluding shadow volumes) specified by *volume*.

You must be superuser to use this command.

PRIMARY OPTIONS

You can use either of the following options.

-F

offline

Stops one or more volumes specified by *volume*. If the *-v* option is omitted, all volumes within the *class* are taken offline. Offline volumes cannot be accessed.

If synchronization copying of *volume* is in process, it cannot be stopped. You can stop the synchronization copying using the `sdxcopy -C` command.

The stopped volume will be activated when the node is rebooted. (Except for when the volume is locked.) Attempting this operation while *volume* is in use results in an error.

-M

Make

Creates a volume specified by *volume*, within the highest level group specified by *group*, or within a single disk specified by *disk*. The *size* indicates the number of blocks on *volume*, *class* indicates the class name associated with the *group* or *disk*.

If *class* is the root class, a maximum of 14 volumes with their physical slice attribute set to "on" can be created within the same *group*.

If *class* is a local class or a shared class, the same *group* or *disk* can contain a maximum of 4 volumes with their physical slice attribute set to "on." A maximum of 224 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."

After volume creation is complete, the volumes are started on a node where the command was executed and become accessible through the following special files.

```
/dev/sfdsk/class/rdsk/volume [Linux2.4]
/dev/sfdsk/class/dsk/volume
```

If *group* is a mirror group, the system will automatically execute a synchronization copying upon returning from the `sdxvolume` command.

The features of volumes created when *group* is mirror group, stripe group and switch group are explained below.

- When group is a mirror group

To ensure data availability, GDS restricts the mirroring on a single piece of disk unit. In the case of mirror groups, a mirror volume that consists of mirror-multiplexing equal to the number of connected disks or lower level groups is created (maximum of eight). When only one disk or one lower level group is connected, the volume created within that mirror group will not be mirrored.

If the last block number of the volume that is created within a mirror group is larger than the available size of any of the spare disks registered with *class*, a warning message is sent to standard error output informing you that the hot spare feature is disabled.

- When group is a stripe group

In a stripe group, stripe volumes with columns equal to the number of connected disk or lower level groups are created. When only one disk or lower level group is connected, volume cannot be created.

- When group is a switch group

In a switch group, switch volumes with redundancy equivalent to the number of connected disks (a maximum of 2) are created. If only one disk is connected, an active disk switch cannot be performed.

If the active disk is not connected to a node where the command was executed, the volumes are not started. To use the created switch volumes, perform an active disk switch with the *sdxattr -G* command, or move to a node to which the active disk is connected, and then start the volumes with the *sdxvolume -N* command.

-N

oNline

Activates one or more volumes specified by *volume*,... If the *-v* option is omitted, all volumes within *class* are activated.

Activated volumes can be accessed.

If there is a slice in TEMP status on the *volume*, a warning message is sent to standard error output.

If *volume* is a mirror volume, the system will determine whether synchronization has been lost upon returning from the *sdxvolume* command and automatically execute a synchronization copying as needed (except for when *-e nosync* is specified).

If *volume* is a switch volume, it cannot be started on nodes to which the active disk is not connected. If *volume* belongs to the highest level concatenation group to which lower level switch groups are connected, it also cannot be started on nodes to which the active disk of *volume* is not connected.

-R

Remove

Removes the volume specified by *volume* and releases the disk area used in the group or the single disk.

If the specified volume is active, this command results in an error.



Note

Be aware that any data stored on *volume* will be lost.

-S

reSize

Expands the size of a volume specified by *volume* to *size* blocks.
class indicates the name of a class to which *volume* belongs.

volume must be a volume that belongs to any one of:

- A single disk
- A mirror group that consists of only one disk
- A mirror group that consists of only one lower level group

The size can be expanded even when the *volume* is active.

The first block of *volume* is not changed. If any area of a volume other than *volume* exists in the area of *size* blocks after the first block of *volume*, it results in an error.

SUB OPTIONS

Sub options are as follows:

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

Use this to set an *attribute* for the *volume*.

The *attribute* indicates the attribute name, and the *value* indicates the attribute value. Always separate *attribute* and *value* with an equal (=) sign.

When indicating multiple attributes, specifiers should be combined using commas(,) as the delimiter. You can indicate the following combination to *attribute* and *value*.

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

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

Sets the volume's JRM mode.

If a group other than a mirror group is specified in the -g option, this command results in an error.

on

JRM is "on."

off

JRM is "off."

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

Sets the physical slice attribute value of volume.

When using -g option to indicate a group in which a physical slice cannot be created (i.e. stripe group, concatenation group, or mirror group where the only object directly connected to mirror group is lower level group), this option cannot be set to "on." If *class* is the root type, this option cannot be set to "off."

on

Physical slice attribute value is set to "on."

Among slices composing the volume, the slices within the single disk, the disks that are connected to the switch group or the disks that are directly connected to the mirror group are registered to the partition table, and physical slices are created.

off

Physical slice attribute value is set to "off."

None of the slices consisting the volume is registered to the partition table, and physical slice will not be created. When physical slice attribute is set to "off", the slice cannot be detached.

-c class

The *class* indicates the class name to which the volume that is the target of the operation belongs, or the class name in which the volume is to be created.

-d *disk* (when using **-M**)

The *disk* indicates the single disk name in which the single volume will be created.

-e allnodes (when using **-F,-N**)

Stops or activates the volume on all nodes included in the scope of *class*.

Stopped nodes are ignored. *class* must be a shared class.

When neither this option nor **-e node=*node* [*:node:...*]** option is specified, volume is stopped or started only on the self-node.

-e delay=*msec* (when using **-M,-N**)

If synchronization is not maintained when creating or activating a mirror volume, synchronization copying will take place automatically (except for when **-e nosync** is specified).

This option delays the issuing of the input/output request to the disk at the time of copying by milliseconds specified by *msec*, allowing adjustment for the effect on the application accessing the volume.

The value is set to 0 by default.

Values from 0 to 1000 may be specified for *msec*.

If mirror group is not specified with *group*, this option is ignored.

-e mode=*val* (when using **-N**)

Specifies the access mode for one or more volumes that will be activated.

val indicates either of the following options.

rw

Sets access mode for read and write.

ro

Sets access mode for read only. Opening a read-only volume in write mode will result in an error.

Although volume will be activated in the access mode specified by *val*, the access mode attribute for volume will remain unchanged. Access mode specified by *val* ("Current Access Mode") is valid only while the volume is activated, and will become invalid once the volume is stopped. When the volume is restarted, it will start in the mode set by access mode attribute ("Default Access Mode"), except for when the access mode is specified at the time of restart.

In order to start a volume that is already activated on the current node in a different access mode, you must first stop the volume.

-e node=*node*[:*node*,...](when using **-F,-N**)

Stops or activates the volume on one or more specified nodes.

Stopped nodes are ignored. You must specify the node identifier of the node to stop or activate the volume to *node*. If a node not included in the scope of *class* is specified, the volume is not stopped or activated on any node. *class* must be a shared class.

If this option and the **-e allnodes** are both omitted, the volume is stopped or activated only on the current node.

-e nosync (when using **-N**)

Disables automatic synchronization copying after activating a mirror volume.
If mirror group is not specified with *group*, this option is ignored.



Note

Volumes that are activated using this option will not be mirrored.
In order to configure a mirroring environment, you must perform synchronization copying with the `sdxcopy -B` command.

-e unlock (when using **-N**)

The volume will be activated regardless to whether or not it is locked. Lock mode will not be changed unless you change it with the `stxattr -V` command.

-g group (when using **-M**)

The *group* indicates the group name in which the volume will be created.

-s size (when using **-M**)

Specifies the size of the *volume* being created, in blocks (base 10).
One block is 512 bytes.

When *group* indicates stripe group, the size of *volume* created will be *size* rounded up to a common multiple of stripe width multiplied by stripe columns and cylinder size. In other cases, the size of *volume* created will be *size* rounded up to the integer multiple of cylinder size.

-s size (when using **-S**)

Specifies the number of blocks (decimal number) to which the size of the specified *volume* is expanded. One block is 512 bytes.

The size of the expanded *volume* will be the size rounded up to the integer multiple of the cylinder *size*.

-v volume (when using **-M,-R**)

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

-v volume,... (when using **-F,-N**)

The *volume,...* indicates one or more volume names that is the target of the operation.
To indicate multiple volumes, separate each volume name with a comma (,) as the delimiter.

RETURNED VALUE

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

D.5 sdxslice - Slice operations

SYNOPSIS

```
sdxslice -F -c class -s slice,...
sdxslice -M -c class -d disk -v volume
    [-a attribute=value[,attribute=value]]
sdxslice -N -c class -s slice,...
sdxslice -R -c class {-d disk -v volume}-s slice }
    [-e delay=msec,waitsync]
sdxslice -T -c class -s slice,...
```

DESCRIPTION

Use sdxslice to perform operations on slice objects (excluding shadow slices). You must be superuser to use this command. Slice operations are impossible for the root class.

PRIMARY OPTIONS

You can use one of the following options.

-F

oFffline

Stops the slice or slices specified by *slice*,... Offline slices cannot be accessed. The *slice* indicates the mirror slice name detached from the mirror volume using -M option.

The *class* indicates the class name to which *slice* belongs.

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 specified by *volume* with mirroring multiplicity of two and higher, which is a part of the disk specified by *disk*. The *class* indicates the class name to which *volume* belongs.

Only when the physical slice attribute value of *volume* is "on", you can detach the slice. When the physical slice attribute is set to "off", you must turn it to "on" using the sdxattr -V command before executing this command.

Once detached, a special file is placed on the system. The path name is given below.

```
/dev/sfdsk/class/rdsk/disk.volume [Linux2.4]
/dev/sfdsk/class/dsk/disk.volume
```

Users can access the slice with this special file. You can use this slice to create a data backup of volume.

If *class* is a shared class, only the node that detached the slice can access the slice. Other nodes sharing the *class* cannot access it. If you need to access from other nodes, you can take over the access right with the -T option.

A slice can be detached even though the *volume* is active. You must ensure the integrity of backup data at the file-system layer or database layer.

If you are handling the volume as the file-system, for instance, there will be situations where you must regain integrity using the fsck command.



Note

Be aware that as long as a slice is not attached using the -R option, the degree of mirror multiplexing stays reduced.

-N

oNline

Activates the slice or slices specified by *slice*,... Activated slices can be accessed. The *slice* indicates the mirror slice name detached from the mirror volume using -M option. The *class* indicates the class name to which *slice* belongs.

-R

Remove

Reassembles the slice as part of the volume, where the slice is specified by *slice* or combination of *disk* and *volume*.

The *disk* and *volume* combination or *slice* indicates the mirror slice name disconnected from the mirror volume using -M option.

The *class* indicates the class name to which the slice belongs.

After (when using -e waitsync option, before) returning from the sdxslice command, the slice is automatically reassembled with the volume. If the volume is active at this time, a synchronization copy is executed.

Attempting this command while the slice is in use results in an error.

-T

Takeover

Takes over the slice or slices specified by *slice* from another node. When the takeover is complete, the slice will stop on the original node and be activated on the current node, allowing operation to the slice on the current node.

When a slice is attached to the volume after executing this command, the entire block will be copied regardless to the setting of JRM mode.

This option is effective only for a shared class.

The *slice* indicates the mirror slice name disconnected from the mirror volume using -M option.

The *class* indicates the class name to which *slice* belongs.

Attempting this command while the slice is in use results in an error.

SUB OPTIONS

Sub options are as follows:

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

Sets the *attribute* attribute of the slice to be *value*.

Both attribute values become invalid at the point when the slice is assembled with the volume.

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

Always separate *attribute* and *value* with an equal (=) sign.

Specifiers should be combined using a comma (,) as the delimiter.

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

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

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

Sets the slice's JRM mode.

on

JRM is "on."

off

JRM is "off."

mode=rw or **mode=ro** (**rw** when omitted)

Sets the slice access mode.

rw

Sets access mode for read and write.

ro

Sets access mode for read only. Opening a read-only slice in write mode will result in an error.

-c class

The *class* indicates the local or shared class name to which the slice belongs.

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

The *disk* indicates the disk name to which the slice belongs, where slice is the target of the operation.

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

This option delays the issuing of the input/output request to the disk at the time of synchronization copying of the slice detached from volume, by milliseconds specified by *msec*.

Always separate delay and *msec* with an equal (=) sign.

This option allows you to adjust the influence on the application accessing the volume.

The value is set to 0 by default. Values from 0 to 1000 may be specified for *msec*.

-e waitsync (when using **-R**)

When synchronization copying is executed, returns the command after the copying process is complete.

-s slice (when using **-R**)

The *slice* indicates slice that is the target of operation. Slice name should be specified in *disk.volume* format.

-s slice,... (when using **-F,-N, -T**)

The *slice* indicates one or more slice names that is the target of the operation. To indicate multiple slices, separate each slice name with a comma (,) as the delimiter.

Slice name should be specified in *disk.volume* format.

-v volume (when using **-M,-R**)

Specifies the name of volume comprising the slice that is the target of the operation.

RETURNED VALUE

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

D.6 sdxinfo - Display object configuration and status information

SYNOPSIS

sdxinfo [-ACDGSV] [-c *class*] [-o *object*] [-e *long*]

DESCRIPTION

Use sdxinfo to display configuration and status information of the objects which the current node shares. The sdxinfo command allows you to view configurations and status information of SDX objects, proxy objects and shadow objects.

PRIMARY OPTIONS

Primary options indicate the type of objects to display. If nothing is specified, only information on the pertinent object will be displayed. Object types can be any combination of the following.

-A

All

Displays all specified objects, and information on all related objects.

This is equivalent to -CDGSV.

If any other options are combined with this, they are ignored.

-C

Class

Displays all specified objects, and class information of the related objects.

-D

Disk

Displays all specified objects, and disk information of the related objects.

-G

Group

Displays all specified objects, and group information of the related objects.

-S

Slice

Displays all specified objects, and slice information of the related objects.

-V

Volume

Displays all specified objects, and volume information of the related objects.

SUB OPTIONS

Sub options are used to specify the names of objects to display.

If nothing is specified, the command is interpreted as if all objects at the current node had been specified.

-c *class*

The *class* indicates the class name whose information will be displayed.

If this option is omitted, this command is interpreted as if all classes had been specified.

In case of combining this option with the -o option, the objects related to the specified *object* within the *class* is displayed.

-e long

Displays detailed object information.

-o object

The *object* indicates the object name (class name, disk name, group name or volume name) whose information will be displayed.

If this option is omitted, this command is interpreted as if all object names had been specified.

In case of combining this option with the -c option, the objects related to the specified *object* within the *class* is displayed.

DISPLAYED INFORMATION

Information displayed in response to the `sdxinfo` command can be interpreted as follows.

Class information:

OBJ

Displays class as an object classification.

NAME

Displays the class name.

TYPE

Displays one of the following values.

root [Itanium]

root class

local [Linux2.6]

local class

shared

shared class

SCOPE

Displays the node names as scope attribute values. In a PRIMECLUSTER system, "(local)" is displayed for the root class, the node identifier (CF node name) is displayed for a local class, and node identifiers separated by colons (:) are displayed for a shared class.

HS

When the -e long option is used, this displays hot spare operation, which can be either of the following.

on

Enables the hot spare.

off

Disables the hot spare.

Be aware that spare disk automatic connection is prevented.

For a shadow class it always displays "on", but the hot spare feature is practically invalid since a shadow class cannot include a spare disk.

SPARE

Displays the number of spare disks that are not connected to the group.

SHADOW

When the -e long option is specified, one of the following is displayed as the class description.

0

Class created with the `sdxdisk -M` command.

1

Shadow class created with the `sdxshadowdisk -M` command.

HSMODE

Displays one of the following values to indicate the spare disk selection method for automatic connection by hot spare when the -e long option is specified. For a shadow class, an asterisk (*) is displayed.

exbox

External mode.

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

bybox

Internal mode.

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

Disk information:

OBJ

Displays disk as an object classification.

NAME

Displays the disk name.

TYPE

The type attribute value can be any of the following:

mirror

Mirror. It is connected to a mirror group.

stripe

Stripe. It is connected to a stripe group.

concat

Concatenation. It is connected to a concatenation group.

switch

Switch. It is connected to a switch group.

keep [Itanium]

Keep. When it is registered with a class or connected to a group, the format and data of the disk are retained.

single

Single. Single volumes can be created on it.

spare

Spare. "spare" is also displayed when it is connected to a group.

undef

Undefined. Its usage is not determined yet.

CLASS

Displays the class name to which the disk belongs.

GROUP

Displays the group name to which the disk is connected. If the disk is not connected to any group, an asterisk (*) is displayed.

DEVNAM

Displays the physical disk name in either the following formats. If the disk is not connected to the local node, an asterisk (*) is displayed.

sdX (for normal hard disks)
emcpowerX (for emcpower disks)

X indicates the device identifier

DEVBLKS

Displays the size of the physical disk. The size is given in blocks (sectors.) If the disk is not connected to the local node, 0 (zero) is displayed.

FREEBLKS

When the -e long option is used, displays the number of free allocatable blocks (sectors) for a new volume. If the disk is not a single disk, an asterisk (*) is displayed.

DEVCONNECT

Displays a list of node identifiers of the nodes to which the physical disk is connected, separated using colons (:) as delimiters. If PRIMECLUSTER has not been installed or the physical disk has not been registered in the PRIMECLUSTER resource database, an asterisk (*) is displayed.

STATUS

Disk status can be any of the following:

ENABLE

Available for work.

DISABLE

Not available for work.

SWAP

Not available for work, but available for disk exchanging.

E

When -e long option is used, displays error status of disk, which can be any of the following.

0

I/O error has not occurred on both the current node and the shared node.

1

I/O error has occurred on either the current node or the shared node.

Group information:**OBJ**

Displays group as an object classification.

NAME

Displays the group name.

CLASS

Displays the class name to which the group belongs.

DISKS

Displays the name of disks or lower level groups that are connected to the group, separated using colons (:) as delimiters. In case of a stripe group, names are listed in the order they are striped. Likewise, if it is a concatenation group, names are listed in the order they are concatenated.

BLKS

Displays the size of the group, which is the total of available size of group (available capacity as volumes) plus one private slice size. Size is given in blocks (sectors).

FREEBLKS

The number of free allocatable blocks for a new volume.
If the group is a lower level group, an asterisk (*) is displayed.

SPARE

Displays the number of spare disks that can be connected to the group. Unless it is a mirror group, an asterisk (*) is displayed.

MASTER

When the -e long option is used, it displays the group name of master group. When it is not a proxy group, an asterisk (*) is displayed.

TYPE

When -e long option is used, displays type attribute value, which can be any of the following.

mirror

Mirror group.

stripe

Stripe group.

concat

Concatenation group.

switch

Switch group

WIDTH

When the -e long option is used, displays stripe width in blocks (sectors). If the group is not a stripe group, an asterisk (*) is displayed.

ACTDISK

Displays the disk name of the active disk when the -e long option is specified. Unless it is a switch group, an asterisk (*) is displayed.

Volume information:**OBJ**

Displays volume as an object classification.

NAME

Displays the name of the volume.
If it is an area that cannot be allocated (area for private slice) or that can be allocated but have not (unallocated area), an asterisk (*) is displayed.

TYPE

When the -e long option is used, displays type attribute, which can be any of the following.

mirror

Mirror. It belongs to a mirror group.

stripe

Stripe. It belongs to a stripe group.

concat

Concatenation. It belongs to a concatenation group.

switch

Switch. It belongs to a switch group.

single

Single. It belongs to a single disk.

CLASS

Displays the class name to which the volume belongs.

GROUP

Displays the highest level group name to which the volume belongs. When it belongs to a single disk, an asterisk (*) is displayed.

DISK

When the -e long option is used, displays the name of the single disk to which the volume belongs. When it belongs to a group, an asterisk (*) is displayed.

MASTER

When the -e long option is used, it displays the volume name of master volume. When it is not a proxy volume, an asterisk (*) is displayed.

PROXY

When the -e long option is used, it displays the proxy volume status in one of the two ways as given below. When it is not a proxy volume, an asterisk (*) is displayed.

Join

The volume is being joined to a master volume.

Part

The volume is being parted from a master volume.

SKIP

Displays the skip-resynchronization mode setting, which can be either of the following. If the volume is neither a mirror volume nor a single volume, an asterisk (*) is displayed.

on

Skip resynchronization.

off

Execute resynchronization.



Note

Note that the interface for setting or changing this option is not available.

JRM

Displays the just resynchronization mode setting, which can be either on or off. If the volume is neither a mirror volume nor a single volume, an asterisk (*) is displayed.

on

JRM is "on."

off

JRM is "off."

MODE

When the -e long option is used, it displays the access mode attribute value (default access mode) for the current node, which can be either of the following. If it is either an area for private slice or an unallocated area, an asterisk (*) is displayed.

rw

Read and write mode.

ro

Read only mode.

CMODE

When `-e` long option is used, it displays the present access mode of the activated volume from the current node. If the volume is not activated, an asterisk (*) is displayed.

rw

Read and write mode.

ro

Read only mode.

LOCK

When the `-e` long option is used, displays the lock mode of current node, which can be either of the following. If it is either a private area or an unallocated area, an asterisk (*) is displayed.

on

The volume is locked from activating thereafter.

off

The volume is not locked from activating thereafter.

1STBLK

Displays the block (sector) number of the first block.

The block number is the logical block number, which is the offset in the group to which the volume belongs, and not the physical block number indicating the offset on the physical disk. However, when the volume belongs to a single disk, the block number will match the physical block number on the single disk. Also, when it belongs to a mirror group to which a disk is directly connected or a switch group, the block number will match the physical block number on the disk.

LASTBLK

Displays the block (sector) number of the last block.

The block number is the logical block number, which is the offset in the group to which the volume belongs, and not the physical block number indicating the offset on the physical disk. However, when the volume belongs to a single disk, the block number will match the physical block number on the single disk. Also, when it belongs to a mirror group to which a disk is directly connected or a switch group, the block number will match the physical block number on the disk.

BLOCKS

Displays the size in blocks (sectors).

STATUS

Displays the volume status of the current node, which can be any of the following.

ACTIVE

Ready for work.

STOP

Stopped.

INVALID

Stopped, and cannot be activated due to problem with data.

FREE

Not yet allocated as a volume.

PRIVATE

An area reserved for GDS control, so cannot be allocated as a volume.

PSLICE

When the -e long option is used, displays the physical slice attribute value, which can be either of the following. If it is either a private area or an unallocated area, an asterisk (*) is displayed.

on

Physical slice attribute of the volume is set to "on." Among slices comprising volumes, slices on single disks, on disks connected to switch groups and on disks directly connected to mirror groups are registered with partition table and have physical slices. If a lower level group is the only object directly connected to mirror group, the volume will not have a physical slice, regardless to this attribute being set to "on." Also, when the volume belongs to either a stripe group or a concatenation group, this attribute value will never be set to "on."

off

Physical slice attribute of the volume is set to "off." The volume has no physical slices, and none of the slices in the volume is registered to the partition table.

For a shadow volume it always displays off no matter whether the shadow slice is registered with partition table.

SNUM

When the -e long option is specified, the slice number of the slice configuring the volume is displayed. If the physical slice attribute is off or no physical slice configures the volume, an asterisk (*) is displayed.

PJRM

When the -e long option is specified, either of the following values is displayed to indicate the just resynchronization mechanism mode on the proxy volume.

If it is not a proxy volume parted from the master volume with the sdxproxy Part command, an asterisk (*) is displayed.

on

Proxy JRM is "on."

off

Proxy JRM is "off."

Slice information:

OBJ

Displays slice as an object classification.

NAME

When -e long option is used, this display the name of slice.

When the slice is not a mirror slice that is temporarily detached from the mirror volume using the sdxslice -M command, an asterisk (*) is displayed.

CLASS

Displays the class name to which the slice belongs.

GROUP

Displays the highest level group name to which the slice belongs.

If it is a single slice, an asterisk (*) is displayed.

DISK

Displays the name of the disk or the lower level group (i.e. the group to which this slice belongs, among the groups that are directly connected to the relevant highest level group) to which the slice belongs. If the highest level group is a switch group, the disk name of the active disk is displayed. If the highest level group is a stripe group or a concatenation group, an asterisk (*) is displayed.

VOLUME

Displays the volume name to which the slice belongs.

JRM

When the `-e` long option is used, displays the just resynchronization mode setting, which can be either on or off.

When the slice is not a mirror slice that is temporarily detached from the mirror volume using the `sdxslice -M` command, an asterisk (*) is displayed.

on

JRM is "on."

off

JRM is "off."

MODE

When `-e` long option is used, displays the access mode, which can be either of the following.

When the slice is not a mirror slice that is temporarily detached from the mirror volume using the `sdxslice -M` command, an asterisk (*) is displayed.

rw

Read and write mode.

ro

Read only mode.

STATUS

Displays the slice status on the current node, which can be any of the following.

ACTIVE

Ready for work.

STOP

Stopped.

INVALID

Due to a problem with data, temporarily detached from the volume.

COPY

Specifies a copy is underway, to maintain data uniformity.

TEMP

Temporarily detached from volume. Slice is operating in isolation.

TEMP-STOP

Temporarily detached from volume. Slice is stopping in isolation.

NOUSE

Stopped, with no operations possible.

COPY

When the `-e` long option is used, this displays one of the following copying process statuses. When the slice is not in COPY status, an asterisk (*) is displayed.

run

Copying is underway.

bg

Copying is in process in the background, but you can access valid data.

intr

Copying has been interrupted. Executing the `sdxcopy -I` command interrupts the copy.

wait

Since many copying processes are in progress, it is waiting to be scheduled.

CURBLKS

When the `-e` long option is used, this displays the number of blocks (sectors) that have already been copied.

When CURBLKS and the later described COPYBLKS match, all copying has been completed.

When the slice is not in COPY status or is being copied with TimeFinder or SRDF using GDS Snapshot, an asterisk (*) is displayed.

COPYBLKS

When the `-e` long option is used, this displays the number of blocks (sectors) that needs to be copied.

Usually this is the same size as the volume size it is registered with, but when just resynchronization copying is in process, the number of blocks that actually needs to be copied will be displayed.

When the slice is not in COPY status, an asterisk (*) is displayed.

DLY

When the -e long option is used, this displays the copy delay time in milliseconds. When not in COPY status, an asterisk (*) is displayed.

CPTYPE

When the -e long option is used, one of the following values is displayed as copy function which is used for copying process between the master and the proxy. When copying is not in process, or if the volume to which the slice belongs is not a target volume of copying process between a master volume and a proxy volume, an asterisk (*) is displayed.

soft

Copying is in process using the soft copy function provided by the GDS sfdsk driver.

EC

Copying is in process using the Equivalent Copy function.

OPC

Copying is in process using the One Point Copy function. When the master and the proxy have been joined, and the copy source volume is active, the soft copy function may be involved in part of the copying process.

REC

Copying is in process using the Remote Equivalent Copy function.

ROPC

Copying is in process using the Remote One Point Copy function. When the master and the proxy have been joined, and the copy source volume is active, the soft copy function may be involved in part of the copying process.

TF [Linux2.6]

Copying is in process with EMC TimeFinder.

SRDF [Linux2.6]

Copying is in process with EMC SRDF.

CPSOURCE

When the -e long option is used, the volume name of source proxy volume which is used to restore data is displayed. When the volume to which the slice belongs is not a target master volume of restoring process, an asterisk (*) is displayed.

FUNC

Displays one of the following values to indicate the session type of the disk unit's copy function when the -e long option is specified. If there is no session, an asterisk (*) is displayed.

EC

Source or target of the Equivalent Copy session.

OPC

Source or target of the One Point Copy session.

REC

Source or target of the Remote Equivalent Copy session.

ROPC

Source or target of the Remote One Point Copy session.

TF [Linux2.6]

Source or target of EMC TimeFinder's BCV pair.

SRDF [Linux2.6]

Source or target of EMC SRDF's SRDF pair.

CPSTAT

Displays one of the following values to indicate the session status of the disk unit's copy function when the -e long option is specified. If there is no session, an asterisk (*) is displayed.

equiv

Synchronized.

copy

In process of copying.

suspend

EC or REC session suspended.

split [Linux2.6]

BCV pair or SRDF pair split.

error

Suspended due to an error.

halt

Hardware suspended.

PARTNER

When the -e long option is specified, displays a destination slice name if the slice is the source in the session of the disk unit's copy function, or a source slice name if it is the destination. If there is no session, an asterisk (*) is displayed.

RETURNED VALUE

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

USAGE EXAMPLES

Displays all the objects within the current node.

```
sdxinfo
```

Displays information on all the disks registered with the class called "Class1."

```
sdxinfo -A -c Class1
```

Use this to check whether an object called "foo" is currently being used.

```
sdxinfo -o foo
```

CAUTION

Additional information may be displayed in accordance with new functions provided.

D.7 sdxattr - Set objects attributes

SYNOPSIS

```
sdxattr -C -c class -a attribute=value[,attribute=value,...]
sdxattr -D -c class -d disk -a attribute=value[,attribute=value]
sdxattr -G -c class -g group -a attribute=value[,attribute=value]
sdxattr -S -c class -s slice -a attribute=value[,attribute=value]
sdxattr -V -c class -v volume -a attribute=value[,attribute=value,...]
```

DESCRIPTION

Use sdxattr to change attribute values of objects (excluding shadow objects) on the current node. You must be superuser to use this command.

PRIMARY OPTIONS

Primary options are used to specify the category of the object whose attributes are to be set.

-C

Class

Set the attributes of the class specified by *class*.

-D

Disk

Set the attributes of the disk specified by *disk*.

The *class* indicates the class name with which *disk* is registered.

If *disk* is connected to a group, or if there is a volume within *disk*, it will result in an error, and the change in attribute will not be executed.

-G

Group

Set the attributes of the group specified by *group*.

The *class* indicates the class name to which *group* belongs.

When *group* is connected to another group, this option will result in an error and you cannot change the attribute. Also, this option will result in an error if there are one or more activated volumes within the group. Stop all volumes before executing this command.

-S

Slice

Set the attributes of the slice specified by *slice*.

The *class* indicates the class name to which the *slice* belongs.

-V

Volume

Set the attributes of the volume specified by *volume*.

The *class* indicates the class name to which the *volume* belongs.

SUB OPTIONS

Sub options are as follows:

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

Sets the *attribute* attribute of the *class* to be *value*.

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

Always separate *attribute* and *value* with an equal (=) sign.

Specifiers should be combined using a comma (,) as the delimiter.

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

When attempting to set multiple attributes, if any of them result in an error, no attributes are set.

type=local or **type=shared** [Linux2.6]

Changes the class type attribute.

When *class* includes active volumes, it can be changed from "local" to "shared" but cannot be changed from "shared" to "local." To change the class type from "shared" to "local", stop all the volumes and then execute this command.

Changes from "root" to "local" or "shared" and similarly from "local" or "shared" to "root" are impossible. Additionally, the class types cannot be changed with shared classes that include disks of which physical scope is one node or that include switch groups.

local

Change the type attribute to "local."

Objects within local type classes can be used only on the current node.

For the scope attribute, the node identifier of the current node is set automatically.

shared

Changes the type attribute to "shared."

By specifying this along with the scope attribute, objects within the class can be shared with multiple nodes including the current node within the class.

scope=node:node:...

For a "shared" type class, changes the nodes which share the class.

When there is an activated volume within *class*, you can add new nodes, but you cannot remove a node that has already been included in the scope. In order to remove a node, you must execute the command after first stopping all volumes.

If the *node* settings are not all complete, this will result in an error.

The *node* indicates a node identifier that is defined by PRIMECLUSTER.

The *scope* can indicate a maximum of 4 nodes.

Changing the scope of a class fails if the class is a shared class that includes a disk of which physical scope is 1 node, or that includes a switch group.

hs=on or **hs=off**

It sets the operation of the hot spare.

You can make changes regardless to whether there is an activated volume within *class*.

on

Enables the hot spare.

off

Disables the hot spare. Spare disk automatic connection is restricted.

hsmode=exbox or hsmode=bybox

Changes the spare disk selection mode for automatic connection by hot spare.

This operation is available regardless whether or not there are active volumes within *class*.

exbox

Changes the spare disk selection method to the external mode.

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

bybox

Changes the spare disk selection mode to the internal mode.

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

name=classname

Sets the class name to *classname*.

When there is an activated volume within *class*, it results in an error. First, you must stop all volumes.

When changing a class name through this operation, the paths of special files for volumes are also changed, so you must update the files in which the paths are described, such as */etc/fstab*.

```
/dev/sfdsk/classname/rdisk/ volume_name [Linux2.4]
/dev/sfdsk/classname/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 a comma (,) as the delimiter.

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

When attempting to set multiple attributes, if any of them result in an error, no attributes are set.

type=keep, type=single, type=spare or type=undef

Sets the SDX disk type attribute. If *disk* is not connected to part of nodes included in the scope of *class*, or if a switch group exists in *class*, changing the type attribute of *disk* fails.

keep [Itanium]

Sets the type attribute to "keep."

This disk will then be handled as a keep disk, and its format and data will be retained when it is connected to a group.

Single disks cannot be changed to keep disks.

single

Sets the type attribute to "single."

Single volume may be created within *disk* thereafter.

Keep disks cannot be changed to single disks.

spare

Sets the type attribute to "spare."

The disk will be used as a spare disk thereafter.

When the available size of *disk* is smaller than the available size of the smallest mirror group within *class*, a warning message notifying that the hot spare function may not be available will be sent to standard error output.

undef

Sets the type attribute to "undef."

Hereinafter, this disk will be regarded as an undefined disk, which use is not yet determined.

name=diskname

Sets the name of a disk to *diskname*.

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

Sets the *attribute* attribute of the *group* to be *value*.

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

Always separate *attribute* and *value* with an equal (=) sign. To set multiple attributes, specify sets of these specifiers in comma-delimited format.

Specify any of the following sets into *attribute* and *value*.

When multiple attributes are specified, the entire process is canceled in the event of an error in part of the processes.

name=groupname

Sets the name of the group to *groupname*.

actdisk=disk

Changes the active disk of the switch group specified by *group* to *disk*.

-a attribute=value (when using **-S**)

Sets the *attribute* attribute of the detached *slice* to be *value*.

Both attribute values become invalid at the point when the slice is assembled with the volume using the **-R** option.

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

Always separate *attribute* and *value* with an equal (=) sign.

Specifiers should be combined using a comma (,) as the delimiter.

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

When attempting to set multiple attributes, if any of them result in an error, no attributes are set.

jrm=off

Turns the slice's just resynchronization mechanism mode to "off."

It could be turned "off" regardless to the slice status.

To set the jrm to "on," attach the slice to the volume and then detach it again.

mode=rw or mode=ro

Changes the access mode of current node for slice.

It will result in an error when *slice* is activated. Execute after stopping it.

rw

Sets access mode for read and write.

ro

Sets access mode for read only.

Opening a read-only volume in write mode will result in an error.

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

Sets the *attribute* attribute of the *volume* to be *value*.

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

Always separate *attribute* and *value* with an equal (=) sign.

Specifiers should be combined using comma (,) as the delimiter.

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

When attempting to set multiple attributes, if any of them result in an error, no attributes are set.

jrm=on or jrm=off

Turn the JRM mode "on" or "off."

You can make changes regardless to whether *volume* is activated or not.

If *volume* belongs to a group other than a mirror group, this command results in an error.

on

JRM is "on."

off

JRM is "off."

lock=on or lock=off

Changes the lock mode of current node for *volume*.

You can make changes regardless to whether *volume* is activated or not.

If *class* is the root class, this command results in an error.

on

The volume is locked from activating thereafter.

off

The volume is not locked from activating thereafter.

mode=rw or mode=ro

Changes the access mode of current node for *volume*.

When *volume* is activated, it results in an error. First, you must stop the *volume*.

rw

Sets access mode for read and write.

ro

Sets access mode for read only.

Opening a read-only volume in write mode will result in an error.

name=*volumename*

Sets the name of the volume to *volumename*.
It will result in an error when *volume* is activated. Execute after stopping *volume*.

pjrm=off

The parted proxy volume's just resynchronization mechanism mode for rejoining or restoring is turned "off."

This can be changed regardless of whether volume is active or inactive.

This attribute value becomes invalid as volume is rejoined to the master volume with the Rejoin or RejoinRestore option of the *sdxproxy* command.

The value cannot be set to "on." To turn "on", the volume must be rejoined to the master volume once and then be parted again.

pslice=on or pslice=off

Turns the physical slice attribute value to be "on" or "off", respectively.
If *volume* is activated, to avoid an error, stop *volume* before indicating this option. This option will also result in an error if there is a detached slice within *volume*. In such case, attach the slice before indicating this option.

on

The physical slice attribute value of volume is set to be "on."
Among the slices consisting *volume*, any slice on a single disk, and any slices on disks that are directly connected to a mirror group will be registered to the partition table.
You cannot change this option to "on" when *volume* belongs to a group that cannot create a physical slice (stripe group, concatenation group, or a mirror group whose only directly-connecting group is a lower level group), or when there is a maximum number (four) of volumes with its physical slice attribute set to "on" within the same group or single disk.

off

The physical slice attribute value of volume is set to be "off."
If *class* is the root type, the value cannot be changed to "off."

-c *class*

The *class* indicates the class name to which the object belongs, or is to be changed, where object is the target of the change.

-d *disk*

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

-g *group*

The *group* indicates the group name that is the target of the change.

-s *slice*

The *slice* indicates the slice name that is the target of the change.
Slice name should be specified in *disk.volume* format.

-v *volume*

The *volume* indicates the volume name that is the target of the change.

RETURNED VALUE

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

D.8 sdxswap - Swap disk

SYNOPSIS

```
sdxswap -I -c class -d disk [-e delay=msec,nowaitsync ]  
sdxswap -O -c class -d disk
```

DESCRIPTION

Use sdxswap to make a disk (excluding a shadow disk) registered with GDS exchangeable and to restore the disk after swapping.

You must be superuser to use this command.

This command is 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.

The physical disk size of *disk* must be equal to or larger than the original physical disk size.

When the highest level group of *disk* is a mirror group, slice configuration or volume contents is copied as needed, and the sdxswap command returns the control once the copying is complete (when using -e nowaitsync option, before the copying process).

If a spare disk is substituted in place of *disk*, then once the redundancy of all related volumes has been restored, the spare disk is disconnected.

-O

swapOut

Detaches all slices within the disk (specified by *disk*), and sets the disk to exchangeable status. This command must be executed before swapping out a crashed disk.

The *class* indicates the class name with which the *disk* is registered.

If *disk* is already nonusable, the status is rechecked and a warning message is sent to standard error output.

The following details explain unexchangeable and exchangeable conditions when *disk* is not connected to a group and when the highest level group of *disk* is a mirror group, a stripe group, a concatenation group or a switch group.

- When *disk* is not connected to a group

When volumes exist in *disk*, the *disk* cannot be made exchangeable.

- When the highest level group of *disk* is a mirror group

When volumes exist in the highest level mirror group of *disk* and detaching slices within the *disk* can change the volume configurations and statuses, the *disk* cannot be made exchangeable.

For example, if there are volumes in the highest level mirror group of *disk*, and if only the disk specified by *disk* is connected to that group, detaching slices within the *disk* will change the configurations and statuses of the volumes. Therefore, the *disk* cannot be made exchangeable.

- When the highest level group of *disk* is a stripe group

When the highest level group of *disk* is a stripe group, the *disk* cannot be made exchangeable by detaching slices within the *disk*.

- When the highest level group of *disk* is a concatenation group

When *disk* is a disk connected to the highest level concatenation group or the active disk connected to a lower level switch group, it cannot be made exchangeable no matter whether or not volumes exist.

When *disk* is the inactive disk connected to a lower level switch group, it can be made exchangeable no matter whether or not volumes exist.

- When the highest level group of *disk* is a switch group

When *disk* is the inactive disk, the *disk* can be made exchangeable regardless whether or not there are volumes.

When *disk* is the active disk, the *disk* can be made exchangeable by detaching slices within the *disk* only if the switch group includes no switch volume or connected inactive disk.

When the switch group includes volumes, remove those volumes in order to make the active disk exchangeable. When it includes the connected inactive disk, perform an active disk switch with the `sdxattr -G` command and then make the former active disk exchangeable.

SUB OPTIONS

Sub options are as follows:

-c *class*

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

-d *disk*

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

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

When restoring the disk, data contained in the volume is copied as needed.

This option delays the issuing of the input/output request to the disk at the time of copying, by milliseconds specified by *msec*.

This option allows you to adjust the influence on the application accessing the volume.

The value is set to 0 by default.

Values from 0 to 1000 may be specified for *msec*.

-e *nowaitsync* (when using -l)

Returns the command before the copying is complete.

RETURNED VALUE

Upon successful completion, a value of 0 is returned.

Otherwise, a non-zero value is returned.

D.9 sdxfix - Restore a failed object

SYNOPSIS

```
sdxfix -C -c class
sdxfix -D -c class -d disk [-e online]
sdxfix -V -c class { -g group|-d disk } -v volume
```

DESCRIPTION

Use sdxfix to restore failed objects (excluding shadow objects). Data on restored disks or volume objects may no longer have consistency, and after the restoration, it is necessary to restore consistency using backup data or checking with the "fsck" command. The "sdxfix" command can be executed with superuser access privileges only.

PRIMARY OPTIONS

You can use one of the following options.

-C

Class

Restores closed *class* to a normal status on the local node.

This command can restore *class* when the *class* includes a configuration database normally accessible and:

- less than 3 disks in ENABLE status and 1 or more disks normally accessible
- three to 5 disks in ENABLE status and 2 or more disks normally accessible
- six or more disks in ENABLE status and 3 or more disks normally accessible

After restoration is complete, objects within the *class* are restored to the previous status. However, if *class* is a local class, volumes that were in STOP status before the class closure will be in ACTIVE status. Additionally, if *class* is a shared class, volumes that were in ACTIVE status before the class closure will be in STOP status.

-D

Disk

Restores the state of disk *disk* that detected an I/O error.

Reads all *disk* area, and if there's no problem, then clears the error status.

If there's a volume in the highest-level group to which *disk* belongs, or *disk* has a single volume, the volume should be stopped or inactive (STOP or INVALID) on all nodes (except when -e online is specified).

When *disk* is connected to a switch group, restoration fails. To clear an I/O error in a disk connected to a switch group, use the sdxswap -O command to make the disk exchangeable and then use the sdxswap -I command to make the disk useable.

-V

Volume

Restores a slice with invalid data (INVALID) or a not-ready slice (NOUSE) specified by a set of *disk* and *volume* or by a set of *group* and *volume* to the STOP status to restore the *volume* with invalid data (INVALID) to the STOP status.

The *volume* should be stopped or inactive (STOP or INVALID) on all nodes. The slice state specified in the combination of *disk* and *volume*, or *group* and *volume* should be INVALID or NOUSE.

Reads the entire slice specified in the combination of *disk* and *volume*, or *group* and *volume*, and if there's no problem, changes the state to STOP, and then changes the state of a stopped slice to INVALID.

SUB OPTIONS

You can use the following sub-options.

-c *class*

Specify a name of the class to which the object belongs.

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

Specify a name of the disk.

-d *disk* (when using **-V**)

When *volume* is a mirror volume, specify a name of the disk that is connected to the mirror group to which the *volume* belongs. This disk should have the INVALID mirror slice that needs to be restored to STOP.

Specify a single disk name when *volume* is a single volume.

When *volume* is a switch volume, specify a disk name of the active disk connected to a switch group that includes the *volume* into *disk*. Do not set an inactive disk name to *disk*.

-e *online* (when using **-D**)

Restores the object even when the highest level group to which *disk* belongs or the single disk specified by *disk* includes an active volume.

-g *group* (when using **-V**)

When *volume* is a mirror volume, specify a name of lower-level group that is connected to the mirror group to which the *volume* belongs. This group should have the INVALID mirror slice that needs to be restored to STOP.

Specify a name of the highest-level stripe group when *volume* is a stripe volume.

Specify a name of the highest-level concatenation group when *volume* belongs to the highest-level concatenation group.

-v *volume* (when using **-V**)

Specify a name of the volume.

RETURNED VALUE

When it is normally terminated, "0" is returned.

Otherwise, a non-zero value is returned.

D.10 sdxcopy - Synchronization copying operation

SYNOPSIS

```
sdxcopy -B -c class -v volume,...[-e delay=msec,nowaitsync]
sdxcopy -C -c class -v volume,...
sdxcopy -I -c class -v volume,...
sdxcopy -P -c class -v volume,... -e delay=msec
```

DESCRIPTION

Use sdxcopy to access to synchronization copying for volume objects (excluding shadow volumes) specified by *volume*.

You must be superuser to use this command.

PRIMARY OPTIONS

You can use one of the following options.

-B

Begin

Attaches slices that have been detached from mirror volumes specified by *volume*,... and executes synchronization copying. The command returns control after the synchronization copying is complete (before the copying process starts when using the -e nowaitsync option). *class* indicates the class name to which the *volume* belongs.

Slices on the *volume* that have a status of INVALID are attached, and then a synchronization copying is executed.

Slices with a status of TEMP* or NOUSE are not attached.

If there are any slices currently involved in a synchronization copy on the *volume*, this command will terminate with an error.

This command resumes copying from the point where it was interrupted with the -I option. Synchronization copying will be executed while the volume is activated or stopped.

-C

Cancel

Cancels synchronization copying in process or interrupted on the mirror volume or mirror volumes specified by *volume*,... After the cancel is completed, the system returns from the command. The *class* indicates the class name to which *volume* belongs.

-B option executes the copying process again.

-I

Interrupt

Interrupts the synchronization copying currently underway on a mirror volume or mirror volumes specified by *volume*,...

The command returns after the interrupt is completed.

The *class* indicates the class name to which *volume* belongs.

-B option executes the copying process from the point where it was interrupted.

-P

Parameter

Changes parameter related to the synchronization copying in process or interrupted on the mirror volume or mirror volumes specified by *volume*,...

The *class* indicates the class name to which *volume* belongs.

Copying in process will resume after the parameter is changed.

The current status of synchronization copying that is either in process or interrupted can be checked by executing the sdxinfo -S command.

SUB OPTIONS

Sub options are as follows:

-c *class*

The *class* indicates the class name to which *volume* belongs.

-e *delay=msec* (when using **-B,-P**)

Delays the issuing of the input/output request to the disk at the time of copying by milliseconds specified by *msec*. This option allows you to adjust the influence on the application accessing the volume.

The value is set to 0 by default.

If the copying process is either completed or canceled, the delay time will return to default (0). Values from 0 to 1000 may be specified for *msec*.

-e *nowaitsync* (when using **-B**)

Returns the command before the copying is complete.

-v *volume,...*

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

To indicate multiple *volumes*, separate each volume name with a comma (,) as the delimiter.

RETURNED VALUE

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

D.11 sdxroot – Root file system mirroring definition and cancellation [Itanium]

SYNOPSIS

sdxroot -M -c class -d disk[,disk,...]

sdxroot -R -c class -d disk[,disk,...]

DESCRIPTION

Use sdxroot to complete or cancel mirroring definition of system disks including root file systems. The "sdxroot" command can be executed with superuser access privileges only.

System disk means the physical disk on which the running Linux operating system is installed. Specifically, it means the entire disk that includes a slice currently running as any one of the following file systems (or a swap area).

/, /usr, /var, /boot, /boot/efi, or swap

PRIMARY OPTIONS

You can use one of the following options.

-M

Make

Checks that one or more system disks specified by *disk* are ready for mirroring (registered with a class and connected to a group) and creates remaining mirroring definitions (such as updates of files *fstab* and *elilo.conf*). After returning from this command, reboot the system immediately. After the system is rebooted, system disk mirroring will commence.

Specify a disk that includes a slice currently running as */* (root), */usr*, */var*, */boot*, */boot/efi*, or a swap area for *disk*. Among */* (root), */usr*, */var*, */boot*, */boot/efi*, and swap areas, although it is not required to specify disks with swap areas only, be sure to specify disks with */* (root), */usr*, */var*, */boot*, and */boot/efi*. Additionally, the disk specified by *disk* must be ready for mirroring (registered with a class and connected to a group).

When synchronization copying is being performed in groups to which the system disks are connected, the sdxroot command results in an error. In this situation, cancel the synchronization copying using the *sdxcopy -C* command, or execute this command after the synchronization copying is completed.

-R

Remove

Checks that one or more system disks specified by *disk* are ready for mirroring cancellation (disks are disconnected from the system disks' groups and only one system disk is connected to each group) and creates remaining mirroring cancellation definitions (such as update of the *fstab*, *elilo.conf* file). After returning from this command, reboot the system immediately. After the system is rebooted, system disk mirroring will be cancelled.

To totally cancel system disk management with GDS, after the system is rebooted, it is necessary to delete system disk related volumes, groups, disks, and classes.

Specify a disk that includes a slice currently running as */* (root), */usr*, */var*, */boot*, */boot/efi*, or a swap area for *disk*. The disk specified by *disk* must be ready for mirroring cancellation (disks are disconnected from the system disks' groups and only one system disk is connected to each group).

SUB OPTIONS

Sub options are as follows.

-c class

class indicates the class name to which *disk* belongs.

-d disk[,disk,...]

disk indicates a target disk. When connecting multiple disks, the disk names should be combined using commas (.).

RETURNED VALUE

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

USAGE EXAMPLES

Assuming the disk on which / (root), /usr, /var, /boot, and /boot/efi are installed and the disk allocated as a swap area are different, examples of the procedures for system disk mirroring and system disk mirroring cancellation are given below.

Procedure for system disk mirroring

1. Stop the running application programs.
In order to ensure mirroring definition, all the running application programs must be stopped. For the mirroring definition to be in effect, the system must be rebooted after going through this procedure.

When higher safety is required, create system disk backups.

2. Register the system disks with the root class.
In this example, the installation disk of / (root), /usr, /var, /boot, and /boot/efi is "sda", and the disk allocated as a swap area is "sdb."

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

3. Connect the system disks to a group.

```
# sdxdisk -C -c System -g Group1 -d Root1,Root2 -v  
1=root:on,2=usr:on,3=var:on,4=home:on,5=boot:on,6=efi:on  
# sdxdisk -C -c System -g Group2 -d Swap1,Swap2 -v 1=swap:on
```



Information When System Disks Have Unopen Physical Slices

After returning from the `sdxdisk -C` command, volumes created for unopen physical slices are started, and synchronization copying is performed. In this event, cancel the synchronization copying using the `sdxcopy -C` command, or after the synchronization copying is completed move to step 4. Physical slices on which file systems are mounted and those accessed as raw devices are considered to be open. Physical slices not displayed with the `mount(8)` command may not be open.

4. Check that mirroring definition is completed.

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

5. Reboot the system.

```
# shutdown -r now
```

6. Check that mirroring is in effect.
Using the mount command and the sdxinfo command, make sure that the system disks have been mirrored properly.



Note **Checking Physical Disk Information and Slice Numbers**

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

Procedure for system disk mirroring cancellation

1. Stop the running application programs.
In order to ensure mirroring cancellation, all the running application programs must be stopped. For the mirroring cancellation to be in effect, the system must be rebooted after going through this procedure.
When higher safety is required, create system disk backups.

2. Disconnect disks other than those used as system disks after this cancellation from the groups.

```
# sdxdisk -D -c System -g Group1 -d Root2
# sdxdisk -D -c System -g Group2 -d Swap2
```

3. Check that mirroring cancellation is completed.

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

4. Reboot the system.

```
# shutdown -r now
```

5. Check that the mirroring has been cancelled.
Using the mount command and the sdxinfo command, make sure that the system disk mirroring has been cancelled properly.

6. Cancel system disk management.

```
# sdxvolume -F -c System -v root
# sdxvolume -F -c System -v usr
# sdxvolume -F -c System -v var
# sdxvolume -F -c System -v home
# sdxvolume -F -c System -v boot
# sdxvolume -F -c System -v efi
# sdxvolume -F -c System -v swap
# sdxvolume -R -c System -v root
# sdxvolume -R -c System -v usr
# sdxvolume -R -c System -v var
# sdxvolume -R -c System -v home
# sdxvolume -R -c System -v boot
# sdxvolume -R -c System -v efi
# sdxvolume -R -c System -v swap
# sdxgroup -R -c System -g Group1
# sdxgroup -R -c System -g Group2
# sdxdisk -R -c System -d Root1
# sdxdisk -R -c System -d Root2
# sdxdisk -R -c System -d Swap1
# sdxdisk -R -c System -d Swap2
```

D.12 sdxparam - Configuration parameter operations

SYNOPSIS

```
sdxparam -G [-p param,...]  
sdxparam -S -p param=val [,param=val,...] [-e default]
```

DESCRIPTION

Use sdxparam to perform operations on GDS configuration parameter.
You must be superuser to use this command.

PRIMARY OPTIONS

You can use one of the following options.

-G

Get

Displays the current value of configuration parameter or parameters specified by *param*.
When using cluster system, parameter value of the current node will be displayed. If the -p option is omitted, all configuration parameters are displayed.

-S

Set

Sets the value specified by *val* to the configuration parameter or parameters specified by *param*.

When using cluster system, this option sets the parameter value of the current node. The new value becomes valid upon returning from the command, and rebooting the system will not change the value.

SUB OPTIONS

Sub options are as follows:

-e *default* (when using -S)

Resets all configuration parameter values to default.
When indicated at the same time as -p option, this option is ignored.

-p *param*,...(when using -G)

Displays the configuration parameter *param* value.

-p *param=val* [,*param=val*,...] (when using -S)

Sets *val* to configuration parameter *param*.

You can indicate the following combination to *param* and *val*.

copy_concurrency=num

Sets the maximum number of synchronization copying you can execute simultaneously to *num*.
The value is set to 8 by default.
Values from 1 to 1024 may be specified for *num*.

copy_delay=msec

This option delays the synchronization copying by milliseconds specified by *msec*, when the copying is initiated by an event other than hot spare.
The value is set to 0 by default.
Values from 0 to 1000 may be specified for *msec*.

spare_copy_delay=msec

This option delays the synchronization copying by milliseconds specified by *msec*, when the copying is initiated by hot spare.
The value is set to 50 by default. Values from 0 to 1000 may be specified for *msec*.

RETURNED VALUE

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

CAUTION

The default values and the range of values you can use for configuration parameters may change in the future.

D.13 sdxconfig – Object configuration operations [Linux2.6]

SYNOPSIS

```
sdxconfig Backup -c class[-o outfile] [-e update]  
sdxconfig Convert -e remove[,update] -c class -d disk,...  
    [-i infile] [-o outfile]  
sdxconfig Convert -e remove[,update] -c class -g group,...  
    [-i infile] [-o outfile]  
sdxconfig Convert -e rename[,update] -c class=classname  
    [-i infile] [-o outfile]  
sdxconfig Convert -e replace[,update] -c class -d disk=device[,disk=device,...]  
    [-i infile] [-o outfile]  
sdxconfig Convert -e replace[,update] -c class  
    -p device=newdevice[,device=newdevice,...]  
    [-i infile] [-o outfile]  
sdxconfig Remove -c class[-e keepid]  
sdxconfig Restore -c class -i infile [-e chkps,skipsync]
```

DESCRIPTION

Use sdxconfig to perform object configuration operations for classes specified by *class* (excluding shadow classes). The sdxconfig command can be executed with superuser access privileges only. This command must be executed in multi-user mode.

PRIMARY OPTIONS

You have a choice of the following options.

Backup

Outputs the object configuration of a local class or a shared class specified by *class* to a file specified by *outfile* (standard output by default) in configuration table format. Specify the class name targeted for configuration table creation for *class*.

If *class* includes switch groups, proxy objects, DISABLE disks, or TEMP slices, creation of the *class* configuration table fails.

Convert

Converts the configuration table of a class specified by *class* according to sub option specifications. Specify the class name contained in the configuration table for *class*.

Remove

Removes the object configuration of a local class specified by *class* from the system. All objects (volumes, groups and disks) within the *class* are removed. Specify the target local class name for *class*.

Even if the class object configuration is removed using this option, contents (data) of the removed volumes are not lost. By restoring the object configuration with the Restore option, the volume configuration and contents can be restored.

If *class* includes proxy objects, ACTIVE volumes, or TEMP or COPY slices, deletion of the *class* fails.

Restore

Restores the object configuration of a class specified by *class* according to the configuration table declared in a configuration file specified by *infile*. Specify the class name contained in the configuration table for *class*.

Even if the class object configuration is restored with this option, volume areas on the physical disks registered with the class are not initialized. After the object configuration is deleted with the Remove option, by restoring the object configuration using this option, the volume configuration and contents can be restored.

However, if the configuration table contains mirror volumes with the mirroring multiplicity of two or higher, after returning from the `sdxcfg` command, synchronization copying of the mirror volumes is performed automatically (excepting when using `-e skipsync`). In this event, destination slices are overwritten with data of source slices automatically selected, and data previously saved on the destination slices will be lost.

The class specified by *class* will be restored as a local class on the current node. To restore the *class* as a shared class, after this command execution it is necessary to change the type attribute and scope attribute of the *class* using the `sdxattr -C` command.

If the class specified by *class* already exists, this command results in an error. Additionally, if the physical disk size contained in the configuration table and the actual physical disk size do not match, restoration of the *class* object configuration fails.

SUB OPTIONS

Sub options are as follows.

-c class

class indicates the target class name.

-c class=classname (when using **Convert -e rename)**

Changes the class name in the configuration table from *class* to *classname*.

-d disk,... (when using **Convert -e remove)**

Removes *disk*,... from the configuration table. Specify the disk name of an undefined disk, a spare disk, a single disk, or a disk directly connected to a mirror group to be removed for *disk*.

If *disk* is a single disk, volumes and slices within the *disk* are also removed. If *disk* is the only disk connected to a mirror group, volumes and slices in the mirror group and the mirror group itself are also removed.

This option can be used along with the `-g` option.

If *disk* is connected to a concatenation group or a stripe group in the configuration table, removing the *disk* fails.

-d disk=device[,disk=device,...] (when using **Convert -e replace)**

Changes the physical disk of a disk specified by *disk* to *device* in the configuration table. *device* can also indicate a physical disk not connected to the domain.

Specify a disk name for *disk* and a physical disk name for *device*. It is necessary to separate *disk* and *device* with an equal sign (=). To change multiple physical disks, specify sets of these specifiers in comma-delimited format.

The physical disk names can be specified in one of the following formats.

sdX (for normal hard disks)

emcpowerX (for emcpower disks)

X indicates the device identifier.

This option cannot be used along with the `-p` option.

-e chkps (when using **Restore**)

Checks consistency of disk identification information (class and disk names) stored in the private slices of physical disks to be registered with *class* and the configuration table contained in the configuration file specified by *infile*.

Restoration of *class* does not take place if any of the following conditions is not satisfied.

- All physical disks contained in the configuration table have the private slices.
- The sizes of the private slices match between all physical disks contained in the configuration table.
- The class names stored in the private slices match between all physical disks contained in the configuration table.
- For all physical disks contained in the configuration table, the disk names stored in the private slices match the disk names assigned to the physical disks in the configuration table.

-e keepid (when using **Remove**)

Retains the private slices and disk identification information stored in the private slices of all disks registered with *class*.

By using this option, when *class* is restored using physical disks that were removed from the *class* or that were copied with the copy functions of disk units, configuration consistency can be checked with the `-e chkps` option of the `sdxconfig Restore` command.



Note

If the object configuration of *class* is deleted with this option, physical disks that were deleted from the *class* cannot be registered with a class with the `sdxdisk -M` command. Before registering those deleted physical disks with a class with the `sdxdisk -M` command, restore the object configuration using the `Restore` option once, and then execute the `Remove` option again without this option.

-e remove (when using **Convert**)

Removes disks or groups from the configuration table.

-e rename (when using **Convert**)

Renames the class of the configuration table.

-e replace (when using **Convert**)

Changes physical disks in the configuration table.

-e skipsync (when using **Restore**)

Leaves synchronization copying of mirror volumes created within *class* undone, assuming that equivalency of all mirror volumes contained in the configuration file specified by *infile* is ensured on user's hand. Even if slices are nonequivalent, their statuses will not be in `INVALID`.

-e update (when using **Backup, Convert**)

Overwrites the file specified by *outfile* with the configuration table when the *outfile* is an existing file.

-g group,... (when using **Convert -e remove**)

Removes *group,...* from the configuration table. All objects (volumes, slices, disks, lower level groups) within the *group* are deleted. Specify the deleted group name for *group*.

This option can be used along with the `-d` option.

When *group* is connected to a higher level group other than a mirror group in the configuration table, removing *group* fails.

-i *infile* (when using **Convert, Restore**)

Converts the configuration table or restores the object configuration of a class specified by *class* according to the configuration file specified by *infile*. Specify the path to a configuration file for *infile*, using the absolute path name or the relative path name from the current directory.

When using Convert, it is not required to specify this option. By default, a configuration table from standard input is converted.

-o *outfile* (when using **Backup, Convert**)

Sends the created or converted configuration table to a configuration file specified by *outfile*. Specify the path to a configuration file for *outfile*, using the absolute path or the relative path from the current directory.

If the file specified by *outfile* already exists, this command results in an error (excepting when using **-e update**).

By default, the configuration table is output to standard output.

-p *device=newdevice[,device=newdevice,...]* (when using **Convert -e replace**)

Changes the physical disk specified by *device* to another physical disk specified by *newdevice* in the configuration table. *newdevice* can also indicate a physical disk not connected to the domain.

Specify a physical disk name described in the configuration table for *device* and a new physical disk name for *newdevice*. It is necessary to separate *device* and *newdevice* with the equal sign (=). To change multiple physical disks to new disks, specify sets of these specifiers in comma-delimited format.

The physical disk names for *device* and *newdevice* can be specified in one of the following formats.

sdX (for normal hard disks)
emcpowerX (for emcpower disks)

X indicates the device identifier.

This option cannot be used along with the **-d** option.

RETURNED VALUE

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

D.14 sdxproxy - Proxy object operations

SYNOPSIS

sdxproxy Break -c *class* -p *proxy* [-e *force,restore*]

sdxproxy Cancel -c *class* -p *proxy*

sdxproxy Join -c *class* -m *master* -p *proxy*
[-a *mvol=pvol:jrm[:pslice][,mvol=pvol:jrm[:pslice],...]*]
[-e *delay=msec,softcopy,syncmode,waitsync*]

sdxproxy Part -c *class* -p *proxy*,... [-a *attribute=value*] [-e *instant,mode=val,unlock*]

sdxproxy Rejoin -c *class* -p *proxy*,... [-e *delay=msec,softcopy,waitsync*]

sdxproxy RejoinRestore -c *class* -p *proxy*,...
[-e *delay=msec,instant,nowaitsync,softcopy*]

sdxproxy Relate -c *class* -m *master* -p *proxy*

sdxproxy Restore -c *class* -p *proxy*,... [-e *instant,nowaitsync*]

sdxproxy Root -c *class* -m *master*,... [-e *boot*] [Itanium]

sdxproxy Root -c *class* -p *proxy*,... [-e *boot*] [Itanium]

sdxproxy Root -c *class* -m *master*,... -p *proxy*,... [-e *boot*] [Itanium]

sdxproxy Swap -c *class* -p *proxy*

sdxproxy Update -c *class* -p *proxy*,... [-e *instant,nowaitsync*]

DESCRIPTION

Use sdxproxy to perform operations on proxy objects.
The sdxproxy command can be executed with superuser access privileges only.



[See](#)

- A.1.8 Exception to Proxy Configuration
- A.1.9 Number of Proxy Volumes
- A.1.10 Proxy Volume Size
- A.1.11 Proxy Group Size
- A.2.16 Using the Advanced Copy Function in a Proxy Configuration
- A.2.17 Instant Snapshot by OPC
- A.2.19 Using EMC TimeFinder or EMC SRDF in a Proxy Configuration [Linux2.6]

PRIMARY OPTIONS

You have a choice of the following options.

Break

Cancels master-proxy relationship between a pair of volumes or groups. You can cancel the relationship when the objects are either joined or parted.

proxy indicates a proxy volume or a proxy group to be cancelled. A proxy volume within a proxy group cannot be specified for cancellation.

Even after the relationship is cancelled, the proxy volume or proxy volumes within the proxy group can be used as normal volumes that retain original volume data and attributes.

You can also cancel the relationship when master volume and proxy volume are in use. However, if the joined master volume is in use, the data integrity of the proxy volume should be ensured in the file system layer or database layer that is managing data, after the relationship is cancelled. For example, if a master volume is being used as a file system, you must use the `umount(8)` command to unmount the file system, before canceling the relationship.

The command will result in an error when:

- Copying is in process from the master volume to the proxy volume (except for when using `-e force`)
- Copying is in process from the proxy volume to the master volume (except for when using `-e restore`)

Cancel

Cancels (releases) sessions of the copy functions of disk units existing between parted proxies and masters.

Specify a parted proxy group or a parted proxy volume into *proxy*. When *proxy* is a proxy group, all parted proxies within the proxy group become targets. A parted proxy volume within a proxy group can also be specified.

However, when a BCV pair or an SRDF pair exists between the master group and the proxy group, sessions cannot be cancelled specifying a parted proxy volume within the proxy group.

Sessions can be canceled even if master volumes and proxy volumes are in use. Even after canceling the sessions, the masters and the proxies are left parted. The volume statuses of the masters and the proxies also remain unchanged. However, if sessions are canceled when copying from masters to proxies and vice versa are in process, data becomes INVALID. If that happens, perform copying again, and the volume statuses will be restored when the copying is complete.

Join

A pair of volumes or a pair of groups are related and joined as master and proxy.

When joining a pair of volumes, synchronization copying of the master volume to the proxy volume is performed after returning from the command (when using `-e waitsync`, before returning from the command).

When joining a pair of volumes, the following conditions must be satisfied.

- The master volume size and the proxy volume size match.
- The master volumes and the proxy volumes belong to different mirror groups or single disks.

When joining a pair of groups, proxy volumes with the same offset and size as master volumes within the master group will be created in the proxy group, and synchronization copying from the master volumes to the proxy volumes is performed after returning from the command (when using `-e waitsync`, before returning from the command). The access mode of the proxy volumes created in the proxy group is set to "ro" (read-only).

If a keep disk is connected to the master group or the proxy group and geometry such as the cylinder size does not match between the master group and the proxy group, geometry of the proxy group is changed conforming to that of the master group.

When joining a pair of groups, there are the following conditions and restrictions.

- For the root class, the smallest physical disk size directly connected to the proxy group must be larger than the last block number of a volume within the master group.
- For a local class or a shared class, the proxy group size must be larger than the last block number of a volume within the master group.
- If the master group has no volume and the proxy group already has volumes, joining them results in an error.
- The master group and the proxy group must be mirror groups.

When the copy functions of disk units are available, synchronization copying from masters to proxies is performed with those copy functions (except for when using `-e softcopy`).

A proxy volume that is joined cannot be accessed or activated. In order to access the proxy volume, part the volume from master using the `Part` option, or break the relationship with master using the `Break` option.

You can create multiple snapshots by joining another proxy to a master, which is already joined with a proxy. However, the total number of slices which belong to the master volume and slices which belong to the proxy volumes that are related to the master volume must not exceed 32.

Master volumes that are already related to a proxy cannot be joined to another master as a proxy. Also, a proxy that is already related to a master cannot be joined with another proxy.

The command will result in an error when:

- The proxy volume is activated.
- There is a slice that is temporarily detached or in copying process among any of the slices comprising the proxy volume or the master volume.
- Copying is in process between the master volume and the other proxy volume.
- The master volume is in `INVALID` status.

Part

Separates a proxy or proxies in joined status from the master. The master-proxy relationship will be maintained after parted. The parted proxy volume will be the snapshot containing the copy of the master volume data at the time of parting. By using the parted proxy volume, you can for instance, create a backup of the master volume at the time it was parted, or use it for other purposes.

proxy indicates a proxy volume or a proxy group in joined status. When proxy group is specified, all proxy volumes in the group will be parted. You can also indicate a proxy volume in joined status in a proxy group.

After the parting is complete, the proxy volume will be activated as an independent volume and become accessible using the following special file.

```
/dev/sfdisk/class/rdisk/volume_name [Linux2.4]
/dev/sfdisk/class/dsk/volume_name
```

When the proxy volume belongs to a "shared" type class, it will be activated on all nodes defined in the class scope.

You can part master and proxy volumes even if the master volume is active, but the data integrity of the parted proxy volume must be ensured in the file system layer or database layer that is managing data. For instance, if you are using the master volume as a file system, use the `umount(8)` command to unmount the file system, before parting.

The command will result in an error when:

- Copying is in process from the master volume to the proxy volume (except for when using `-e instant`)
- Copying is in process from the proxy volume to the master volume

Rejoin

Rejoins one or more parted proxy with the master.

proxy indicates a parted proxy volume, or proxy group. When a proxy group is indicated, all proxy volumes in the group will be rejoined. A parted proxy volume in a proxy group can also be specified.

Synchronization copying of master volumes to proxy volumes is performed after returning from the command (when using `-e waitsync`, before returning from the command). When the copy functions of disk units are available, synchronization copying is performed with those copy functions (except for when using `-e softcopy`).

When more than one proxy volume related to the same master volume is specified simultaneously, this command will result in an error.

The command will result in an error when:

- The proxy volume is active.
- There is a slice that is in copying process among any of the slices comprising the proxy volume or the master volume.
- Copying is in process between the master volume and the other proxy volume.
- The master volume is in INVALID status.



Note [Itanium]

When proxy volumes are running as system volumes, they cannot be stopped and thus rejoining fails. To rejoin such proxy volumes, firstly switch the boot environment using the `sdxproxy Root` command to free up the proxy volumes.

RejoinRestore

Rejoins a proxy or proxies in parted status with the master and restores the master volume data using the proxy volume data. Master volume data are restored by synchronization copying from the proxy volume to the master volume. When executing the command using this option, the master volume data are overwritten with the proxy volume data.

proxy indicates a proxy volume or a proxy group in parted status. When a proxy group is specified, all proxy volumes in the group will be rejoined and data of the related master volumes will be restored. You can also indicate a proxy volume in parted status in a proxy group.

Synchronization copying of master to proxy volume is performed before returning from the command (when using `-e waitsync`, after returning from the command). When the copy functions of disk units are available, synchronization copying from masters to proxies is performed with those copy functions (except for when using `-e softcopy`).

When more than one proxy volume related to the same master volume is specified simultaneously, this command will result in an error.

The command will result in an error when:

- The master volume or the proxy volume is activated.
- There is a slice that is in copying process among any of the slices comprising the proxy volume or the master volume.
- Copying is in process between the master volume and the other proxy volume.
- The master volume is in INVALID status.

Relate

Relates and parts a pair of volumes or a pair of groups as a master and a proxy. This operation does not change data, statuses and attributes of the master and the proxy. To the related master and proxy, sessions by the copy functions of disk units are not set.

To relate a pair of volumes, the master volume and the proxy volume must be equivalent in size. Those volumes also need to be volumes that belong to different mirror groups or different single disks.

To relate a pair of groups, the volume layout (offsets and sizes) of the master group and the proxy group must be same. Those groups also need to be mirror groups.

For masters to whom proxies are already related, other additional proxies can be related. However, the number of slices comprising a master volume and all proxy volumes related to the master volume is limited to 32 in total.

Masters to whom proxies are already related cannot be related as proxies to other masters, or for proxies already related to masters, other proxies cannot be related.

This command will result in an error when:

- A slice being copied or temporarily detached exists in the master volume or the proxy volume.
- Copying is in process between the master volume and another proxy volume.

Restore

Copies data from a parted proxy to a master and restores contents of the master. With the OPC function, data of the proxy at the moment is copied (overwritten) to the master. The command returns control after the copying is complete (right after the copying starts when using the `-e` instant option and the `-e` `nowaitsync` option). If the OPC function is unavailable, the command fails.

Specify one or more proxy groups or parted proxy volumes for *proxy*. When *proxy* is a proxy group, all parted volumes within the proxy group become targets. A parted proxy volume within a proxy group can also be specified. Do not specify multiple proxy volumes related to the same master volume simultaneously.

The Restore operations can be performed even if proxy volumes are active, but it is necessary to secure consistency of data copied to master volumes in the file system layer or database layer that is managing data. For example, if the proxy volume is used as a file system, unmount the file system with the `umount(8)` command and then perform restoration.

This command will result in an error when:

- The master volume is active.
- A slice being copied exists in the master volume or the proxy volume.
- Copying is in process between the master volume and another proxy volume.
- A proxy volume joined to the master volume exists.
- The proxy volume is in INVALID status.

Root [Itanium]

Configures master volumes and proxy volumes specified by *master,...* and *proxy,...* for using them as file systems or swap areas in an alternative boot environment. When a master group or a proxy group is specified, all volumes that belong to the specified group will be configured.

Volumes to be used in an alternative boot environment must conform to the following conditions.

- The volumes are related directly or indirectly as the master and the proxy (alternative volumes) to volumes declared as file systems or swap areas in the `/etc/fstab` file (current volumes).
- The volumes are parted.
- The access mode is `"rw"` (read and write).
- The volumes are in status other than INVALID (invalid data).
- The volumes are not copy destinations.
- The volumes are not running as file systems or swap areas.

It is not required to specify alternative volumes for all current volumes, but a volume to be used as the root file system in the alternative boot environment (alternative root volume) must always be specified.

Before returning from the command, the device names and the special file names contained in the `elilo.conf` file and the `fstab` file on the specified alternative root volume are changed to those for the specified alternative volume. Current volumes of which alternative boot volumes were not specified are included in `fstab` on the alternative root volume without change. After parting the current root volume and the alternative root volume using the Part option, to edit `fstab` on the current root volume, to edit `fstab` on the alternative root volume, or to perform configuration change such as volume creation or deletion, firstly configure the alternative boot environment using this command. When executing the `sdxproxy` command with this option after these configurations are changed, after returning from the command, check whether contents of `fstab` on the alternative root volume are correct. If an alternative volume that is a swap area used as a dump device is specified, configuration of the alternative volume for using it as a dump device takes place when starting the alternative boot environment.

For modifying the `elilo.conf` and `fstab` files on the alternative root volume, the alternative root volume is temporarily mounted on the `/.GDSPROXY` directory. This temporal mount point can be changed by specifying the mount point path in the environment variable `PROXY_ROOT`.

When alternative boot environment configuration is completed, the boot device names for the current boot environment and the alternative boot environment are output to standard output (excepting when using `-e boot`). Be sure to take a note of the output boot device names. By selecting the boot device name for the alternative boot environment in the EFI boot manager's boot option selector window, the environment can be switched to the alternative boot environment. Similarly, by selecting the boot device name for the current boot environment, the environment can be switch back to the current boot environment. With successful boot environment switchover, the boot environment will be

the default boot environment.

Swap

Swaps the master's slices with the proxy's slices.

proxy indicates a proxy volume or a proxy group in the joined status. A proxy volume within a proxy group cannot be specified for swapping.

You can swap the slices when master is in use.

The command will result in an error when:

- There is a slice that is in copying process among any of the slices comprising the proxy volume or the master volume.
- Copying is in process between the master volume and the other proxy volume.
- The proxy volume is in INVALID status.
- Between a master and a proxy, an EC session, a BCV pair, or an SRDF pair exists.

Update

Copies data from a master to a parted proxy and updates contents of the proxy. With the OPC function, data of the master at the moment is copied (overwritten) to the proxy. The command returns control after the copying is complete (right after the copying starts when using the `-e` instant option and the `-e nowait` option).

Updated proxy volumes become snapshots that have copies (replicas) of data of master volumes at the moment. By use of the updated proxy volumes, creating backups of master volumes at the moment and running other services become possible.

Specify one or more proxy groups or parted proxy volumes into *proxy*. When a proxy group is specified, all the parted proxy volumes within the proxy group become targets. A parted proxy volume within a proxy group can also be specified. Do not specify multiple proxy volumes related to the same master volume simultaneously.

The Update operations can be performed even if master volumes are active, but it is necessary to secure consistency of data copied to proxy volumes in the file system layer or database layer that is managing data. For example, if the master volume is used as a file system, unmount the file system with the `umount(8)` command and then perform update.

This command will result in an error when:

- The proxy volume is active.
- A slice being copied exists in the master volume or the proxy volume.
- Copying is in process between the master volume and another proxy volume.
- The master volume is in an INVALID status.

SUB OPTIONS

Sub options are as follows.

-a *attribute=value* (when using **Part**)

Sets *attribute* that is the attribute of the parted proxy volume to *value*. This attribute value becomes invalid when the proxy volume is rejoined to the master volume with the Rejoin or RejoinRestore option.

attribute indicates the attribute name, and *value* indicates the attribute value. It is necessary to separate *attribute* and *value* with an equal sign (=).

You can specify one of the following combinations to *attribute* and *value*.

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

Sets the just resynchronization mechanism mode for proxies.

on

Turns "on" the just resynchronization mechanism mode for proxies.

off

Turns "off" the just resynchronization mechanism mode for proxies.

-a *mvol=pvol:jrm[:pslice]* [*,mvol=pvol:jrm[:pslice],...*] (when using **Join**)

Indicates the proxy volume attributes.

mvol must always be followed by an equal (=) sign, and *pvol*, *jrm* and *pslice* are delimited by a colon (:). When indicating attributes for more than one proxy volume, combine specifiers with a comma (,) as the delimiter.

When joining a pair of groups, *mvol* indicates the volume name of the master volume within the master group. *pvol* indicates the volume name of the proxy volume that will be created in the proxy group corresponding to the master volume specified by *mvol*, *jrm* indicates the just resynchronization mechanism mode for the volume (on or off), and *pslice* indicates the physical slice attribute (on or off). You must specify attributes for all proxy volumes created for respective master volumes within the master group. If *:pslice* is omitted, the physical slice attribute of the proxy volume will be equivalent to that of the corresponding master volume.

When joining a pair of volumes, *mvol* indicates the volume name of the master volume, *pvol* indicates the volume name of the proxy volume, *jrm* indicates the just resynchronization mechanism mode for the volume (on or off), and *pslice* indicates the physical slice attribute (on or off). *mvol* and *pvol* should match with *master* and *proxy* respectively. When not using this option, the proxy volume attributes will be the same as before it was joined.

If *class* is the root type, "off" cannot be set to *pslice*.

-c *class*

class indicates the class name to which the master object or proxy object that will be the target of operation belongs

-e **boot** (when using **Root**) [**Itanium**]

Sets an alternative boot environment as the default boot environment. After returning from the `sdxproxy` command, reboot the system immediately, and the environment will be switched to the alternative boot environment.

When the `sdxproxy` command ends normally, the previous and new boot device names are output to standard output. Be sure to take a note of the output boot device names. By selecting a previous boot device name in the EFI boot manager's boot option selector window, it is possible to boot in the original boot environment. Additionally, by moving a previous boot device name to the top with the Change Boot Order menu on the EFI boot manager's maintenance menu, it is possible to set the original current boot environment as the default boot environment again.

If the root volume's slice status is changed under synchronization copy completion or due to an I/O error, or if a GDS daemon ends abnormally and it is restarted, the boot device for the current boot environment is set as the default boot device again. Therefore, when this option is used, after returning from the `sdxproxy` command it is necessary to reboot the system immediately.

-e delay=*msec* (when using **Join, Rejoin, RejoinRestore**)

Delays the issuing of the input/output request to the disk at the time of synchronization copying between master volume and proxy volume by milliseconds, specified by *msec*.

This option allows users to adjust the influence on the application accessing the master volume.

When copying is performed with the copy function of a disk unit, this option setting is ignored.

When this option is omitted, the delay will be set to 0.

Values from 0 to 1000 may be specified for *msec*.

-e force (when using **Break**)

Forcibly breaks master-proxy relationship even when copying is in process between master and proxy.

When using this option, copying process will be cancelled and the status of the proxy volume will become INVALID.

-e instant (when using **Part**)

Parts proxy volumes and creates virtual snapshots of master volumes with the OPC function even if copying from the master to the proxy is in process. After returning from the command, the parted proxy volume will become accessible before the copying is complete and will serve as a snapshot containing data of the master volume at the time of parted. When copying is in process from proxy to master, or if a hard copy function is not available, the command will result in an error.

-e instant (when using **RejoinRestore**)

Restoration will be completed instantly, and the command will be returned. After returning from the command, synchronization copying from proxy volume to master volume will automatically begin. Although the copying is still in process, restoration will appear to be complete. After returning from the command, the master volume can be activated and accessed, before the copying is complete. Master volume data will appear to have been overwritten by the proxy volume data at the time of executing the command.

-e instant (when using **Restore, Update**)

Instantly completes restore or update and returns from the command. After returning from the command, background copying with the OPC function is performed. Before the background copying is complete, you may start the copy destination volumes for access to valid data.

-e mode=*val* (when using **Part**)

Indicates the access mode of the proxy volume which will be activated. *val* indicates either of the following options.

rw

Sets access mode for read and write.

ro

Sets access mode for read only.

Opening a read-only volume in write mode will result in an error.

Although proxy volume will be activated in the access mode specified by *val*, the proxy volume's access mode attribute will remain unchanged. Access mode specified by *val* is valid only while the proxy volume is activated ("Current Access Mode") and will become invalid once the proxy volume is stopped. When the proxy volume is restarted, it will start in the mode set by access mode attribute ("Default Access Mode"), except for when the access mode is specified at the time of restart.

When this option is omitted, proxy volume will be activated in the access mode according to the access mode attribute set on each node.

-e nowaitsync (when using **RejoinRestore, Restore, Update**)

Returns control from the command right after copying starts. After returning from the command, wait until the copying is complete to start the copy destination volumes. To start copy destination volumes without waiting until copying is complete, use the **-e instant** option. When the **-e instant** option is specified simultaneously, this option is ignored.

-e restore (when using **Break**)

Cancels copying from a proxy to a master when such copying is in process and forces the command to break the relationship between the master and the proxy.

If copying is canceled and relationships are broken with this option, the master volumes after this operation becomes INVALID.

-e softcopy (when using **Join**, **Rejoin** or **RejoinRestore**)

The copy function of a disk unit will not be used for synchronization copying between master volume and proxy volume.

-e syncmode (when using **Join**)

When the REC function is used for synchronization copying from the master volume to the proxy volume, the transmission mode of REC is set to the synchronous mode. The default is the asynchronous Through mode. When the REC function is not used, this option setting is ignored.

-e unlock (when using **Part**)

The proxy volume will be activated regardless to whether it is locked. Lock mode will not be changed unless you change it with the `sdxattr -V` command.

-e waitsync (when using **Join** or **Rejoin**)

When performing synchronization copying, returns from the command after the copying is complete.

-m master (when using **Join**, **Relate**)

Specifies the master volume or the master group that is joined or related.

master indicates the volume name of the master volume, or the group name of the master group.

-m master,... (when using **Root**)

Specifies one or more master volumes or master groups to be handled. When specifying multiple volumes or groups, they must belong to the same class.

Specify the master volume name or the master group name for *master*.

When specifying multiple volumes or groups, these specifiers should be combined using commas (,) as the delimiter.

-p proxy (when using **Break**, **Cancel**, **Join**, **Swap** or **Relate**)

proxy indicates a proxy volume or a proxy group that is the target of the operation.

proxy indicates the volume name of the proxy volume, or the group name of the proxy group.

-p proxy,... (when using **Part**, **Rejoin**, **RejoinRestore**, **Restore**, **Root** or **Update**)

proxy indicates one or more proxy volumes or proxy groups that will be the target of the operation.

When indicating more than one volume or group, they must belong in the same class.

proxy indicates the volume name of the proxy volume or the group name of the proxy group that will be the target of the operation

When indicating more than one volume name or group name, combine them with a comma (,) as the delimiter.

RETURNED VALUE

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

D.15 sdxshadowdisk - Shadow disk operations

SYNOPSIS

```
sdxshadowdisk -C -c class -g group -d disk,...  
                [-a attribute=value[,attribute=value]]  
sdxshadowdisk -D -c class -g group -d disk  
sdxshadowdisk -M -c class  
-d device=disk[:type][,device=disk[:type],...]  
sdxshadowdisk -R -c class -d disk
```

DESCRIPTION

Use sdxshadowdisk to perform operations on shadow disks specified by *disk*. This command can be executed with superuser access privileges only.

PRIMARY OPTIONS

You have a choice of the following options.

-C

Connect

Connects one or more shadow disks (single type or undefined type) specified by *disk*,... to a group specified by *group*. *class* indicates the name of the shadow class with which *disk* is registered.

If there is no shadow group with the name specified by *group*, it is created automatically.

You cannot connect a shadow disk of the single type to an existing shadow group. Also, multiple shadow disks of the single type cannot be connected to the same shadow group simultaneously.

The type attribute of a shadow disk connected to a shadow group will be changed to match the type attribute of that group (mirror, stripe or concatenation). Shadow disks and lower level groups that are connected to the same shadow group will be mirrored, striped or concatenated, depending on their type attributes.

Details about connecting shadow disks to mirror type, stripe type, and concatenation type shadow groups are described below.

- When connecting to a shadow group of the mirror type

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

If a shadow disk is connected to a shadow group of the mirror type including a shadow volume, synchronization copying of the shadow volume is not performed. To ensure synchronization for a shadow volume of the mirror type, the mirror volume must be properly synchronized with GDS that manages the mirror volumes corresponding to shadow volumes.

By connecting a shadow disk of the single type including a shadow volume to a group of the mirror type, the shadow volume can also be changed from the single type to the mirror type.

The available size of a shadow group of the mirror type (available capacity as shadow volumes) will be the same size as that of the smallest shadow disk or lower level shadow group connected. If connecting *disk* results in a decrease in the available size of *group*, a warning message will be sent to standard error output.

- When connecting to a shadow group of the stripe type

Shadow disks specified by *disk,...* will be connected to *group* in the order they are listed. Disks connected to a stripe group in another domain should be connected in the same order. Alternatively, destination disks copied with the copy functions of disk units from disks connected to a stripe group should be connected in the same order. For the disk connecting order, check the DISKS field displayed with the `sdxinfo -G` command. Respective shadow disks and lower level shadow groups connected to the same shadow group of the stripe type will configure stripe columns, and will be striped in the order they were connected. When only one shadow disk or one lower level shadow group is connected to a shadow group of the stripe type, a shadow volume cannot be created within that shadow group. When striping "n" number of columns, "n" number of shadow disks or lower level shadow groups must be connected. Striping of two or more columns up to 64 columns is supported.

When a shadow group of the stripe type specified by *group* already exists, stripe columns will be added after the existing stripe columns in *group*, in the order they are specified by *disk,...* However, a shadow disk with the available size smaller than the stripe width cannot be connected to the existing shadow group of the stripe type. In addition, you cannot increase stripe columns by connecting shadow disks to a stripe group with a shadow volume, or to a stripe group connected to a higher level shadow group.

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

You cannot connect a shadow disk of the single type to a shadow group of the stripe type.

- When connecting to a shadow group of the concatenation type

Shadow disks connected to the same shadow group of the concatenation type will be concatenated in the order they are specified by *disk,...* Disks connected to a concatenation group in another domain should be connected in the same order. Alternatively, destination disks copied with the copy functions of disk units from disks connected to a concatenation group should be connected in the same order. For the disk connecting order, check the DISKS field displayed with the `sdxinfo -G` command. A maximum of 64 disks can be concatenated.

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

The available size of an existing shadow group of the concatenation type can be increased by connecting shadow disks. When a shadow group of the concatenation type specified by *group* already exists, shadow disks will be concatenated in the order they were specified by *disk,...* after the disk that was last concatenated in *group*. However, you cannot add a shadow disk to a lower level shadow group of the concatenation type if the highest level shadow group of the stripe type already has a shadow volume. Also if the order of connecting shadow groups from the higher level is the mirror type, the stripe type and the concatenation type, a shadow disk cannot be connected to the lowest level shadow group of the concatenation type. You cannot connect a shadow disk of the single type to a shadow group of the concatenation type.

-D

Disconnect

Disconnects a shadow disk specified by *disk* from a shadow group specified by *group*. *class* indicates the name of the shadow class with which the *disk* is registered, and *group* indicates the name of the shadow group to which *disk* is connected.

The disconnected shadow disk will have the original type attribute again (single or undefined).

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

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

Conditions that do not allow you to disconnect a shadow disk from a shadow group of the mirror type, stripe type or concatenation type are as below.

- When disconnecting from a shadow group of the mirror type

For example, you cannot disconnect *disk* from a shadow group of the mirror type specified by *group* if a shadow volume exists within the group, and *disk* is the only object connected to *group*.

- When disconnecting from a shadow group of the stripe type

You cannot disconnect a shadow disk from a shadow group of the stripe type including an existing shadow volume, or from a shadow group of the stripe type connected to a higher level shadow group..

- When disconnecting from a shadow group of the concatenation type

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

A shadow disk containing shadow volume data cannot be disconnected from a shadow group of the concatenation type.

You cannot disconnect a shadow disk from a lower level shadow group of the concatenation type if the highest level shadow group has an existing shadow volume. Also, if the order of connecting shadow groups from the higher level is the mirror type, stripe type and the concatenation type, a shadow disk cannot be disconnected from the lowest level shadow group of the concatenation type.

-M

Make

Registers one or more physical disks specified by *device* with a shadow class. *class* indicates the name of the destination shadow class. Once physical disks are registered, they can then be managed by GDS. Accordingly, the user can perform operations on the disk by use of the disk name specified by *disk*. However, *device* will be no longer managed by GDS if the current node is rebooted or if the GDS daemon on the current node is re-launched because the configuration information of a shadow class is only retained on the memory of the current node but not stored on the private slice.

If no shadow class with the same name as *class* exists, then it is automatically created. The type attribute of the shadow class is "local," and objects in the shadow class are only available on the current node.

A shadow class can include physical disks that are not registered with other classes in the current domain and on which the private slices of GDS exist. In other words, a shadow class can include physical disks that are registered with classes in other domains and physical disks to which the private slices of SDX disks are copied with the copy functions of disk units. Physical disks can be registered with the same shadow class if they are registered with classes that have the same names in other domains or if they are destinations to which the private slices of SDX disks registered with classes with the same names are copied with the copy functions of disk units. In addition, disks with the private slices of different sizes cannot be registered with the same shadow class.

While contents on physical disks (excluding keep disks) registered by the `sdxdisk` command are initialized, contents on physical disks registered by the `sdxshadowdisk` command are not changed.

-R

Remove

Removes a shadow disk specified by *disk* from a shadow class specified by *class*. *class* indicates the name of the shadow class with which *disk* is registered.

The removed shadow disk is no longer managed by GDS.

When the last shadow disk is removed from *class*, the shadow class definition is automatically removed as well.

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

SUB OPTIONS

Sub options are as follows:

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

When using the **-C** option and defining a new group name with the **-g** option, a new shadow group is automatically created. This option sets *attribute* to *value* for the created *group* attribute.

The *attribute* indicates the attribute name, and the *value* indicates the attribute value. The equal sign (=) is always necessary between attribute and value. When specifying multiple attributes, each specifier set must be separated by a comma (,).

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

You can specify one of the following combination to *attribute* and *value*.

If multiple attributes are specified and any error occurs, the entire process is canceled.

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

Sets the type attribute of *group*.

mirror

Sets type attribute to "mirror."

stripe

Sets type attribute to "stripe."

concat

Sets type attribute to "concatenation."

***width=blks* (default is 32)**

Sets the stripe width of *group*. *blks* indicates the stripe width in blocks (decimal numbers). One block is 512 bytes. For *blks*, you can specify an integer that is the power of two, is equal to or larger than one and is equal to or smaller than the available size of the smallest shadow disk among disks specified by *disk*,... If *group* is not the stripe type, this option will cause an error.

-c *class*

class indicates the name of the shadow class to which the target shadow disk is registered or is to be registered.

-d *device=disk[:type] [,device=disk[:type],...]* (when using **-M**)
device indicates the name of the physical disk, *disk*, the name of the disk, and *type*, the type attribute of the shadow disk. An equal sign (=) always follows *device*, and if *type* is specified it must be separated from *disk* by a colon (:). To register multiple *devices*, separate each specifier set as above with a comma (,).

The physical disk name can be specified in either the following formats:

sdX (for normal hard disks)
emcpowerX (for emcpower disks)

X indicates the device identifier.

If *device* is registered with a class in another domain, the same disk name as that in the domain must be specified to *disk*. If *device* is a destination to which the private slice of an SDX disk is copied with a disk unit's copy function, the same disk name as the SDX disk name must be specified to *disk*.

One of the following types can be specified to *type*. The default value for the registered shadow disk is the undefined type.

single

Single type.

undef

Undefined type.

If "single" is specified to *type*, *device* is registered as a shadow disk of the single type. For the shadow disk of the single type, a shadow volume of the single type can be created with the `sdxshadowvolume` command even if it is not connected to a shadow group.

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

disk indicates the name of the shadow disk that is the object of the operation..

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

disk indicates the name of the shadow disk that is the object of the operation. To connect multiple shadow disks, separate each disk name with a comma (,).

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

group indicates the name of the shadow group to which the shadow disk as the object of the operation is connected or is to be connected.

RETURNED VALUE

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

D.16 sdxshadowgroup - Shadow group operations

SYNOPSIS

```
sdxshadowgroup -C -c class -h hgroup -l lgroup,...  
                [-a attribute=value[,attribute=value]]  
sdxshadowgroup -D -c class -h hgroup -l lgroup  
sdxshadowgroup -R -c class -g group
```

DESCRIPTION

Use sdxshadowgroup to perform operations on shadow groups.
This command can be executed with superuser access privileges only.

PRIMARY OPTIONS

You have a choice of the following options.

-C

Connect

Connects one or more shadow groups (stripe type or concatenation type) specified by *lgroup*,... to a shadow group (mirror type or stripe type) specified by *hgroup*. *class* indicates the name of the shadow class to which *lgroup* belongs.

When no shadow group with the same name as *hgroup* exists, it is created automatically.

A shadow group specified by *hgroup* is referred to as a higher level shadow group, and a shadow group specified by *lgroup* is referred to as a lower level shadow group.

Lower level shadow groups and shadow disks connected to the same higher level shadow group are mirrored or striped according to the type attribute of the higher level shadow group. Connecting a shadow group to a higher level shadow group does not change the type attribute of the lower level shadow group.

You cannot connect shadow groups when:

- *lgroup* is the mirror group.
- *hgroup* is the concatenation group.
- Type attributes of *lgroup* and *hgroup* are the same.

In addition, a shadow group already including a shadow volume cannot be connected to another shadow group.

Details about connecting shadow groups to mirror type and stripe type shadow groups are described below.

- When connecting to a shadow group of the mirror type

One or more shadow groups (stripe type or concatenation type) specified by *lgroup*,... can be connected to a shadow group of the mirror type specified by *hgroup*.

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

If a lower level shadow group is connected to a shadow group of the mirror type with a shadow volume, synchronization copying for the shadow volume is not performed. To ensure synchronization for a shadow volume of the mirror type, the mirror volume must be properly synchronized with GDS that manages the mirror volume corresponding to the shadow volume.

The available size of a shadow group of the mirror type (available capacity as shadow volumes) will be the same as that of the smallest shadow disk or lower level shadow group connected. If connecting *lgroup* decreases the available size of *hgroup*, a warning message will be sent to

standard error output.

- When connecting to a shadow group of the stripe type

One or more shadow groups (concatenation type) specified by *lgroup*,... can be connected to a shadow group of the stripe type specified by *hgroup*. Shadow groups specified by *lgroup*,... will be connected to *hgroup* in the order they are listed. Lower level groups connected to a stripe group in another domain should be connected in the same order. Alternatively, destination disks copied with the copy functions of disk units from lower level groups connected to a stripe group should be connected in the same order. For the order of connecting lower level groups, check the DISKS field displayed with the *sdxinfo -G* command.

Respective shadow disks and lower level shadow groups connected to the same shadow group of the stripe type will configure stripe columns and will be striped in the order they were connected. When only one shadow disk or one lower level shadow group is connected, a shadow volume cannot be created within that shadow group. When striping "n" number of columns, "n" number of shadow disks or lower level shadow groups must be connected. Striping of two or more columns up to 64 columns is supported.

When a shadow group of the stripe type specified by *hgroup* already exists, stripe columns will be added after the existing stripe columns in *hgroup*, in the order they are specified by *lgroup*,... However, a shadow group with the available size smaller than the stripe width cannot be connected to an existing shadow group of the stripe type. In addition, you cannot increase stripe columns by connecting shadow groups to a stripe group with a shadow volume, or to a stripe group connected to a higher level shadow group.

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

-D

Disconnect

Disconnects a shadow group specified by *lgroup* from a shadow group specified by *hgroup*. *class* indicates the name of the shadow class to which *lgroup* belongs, and *hgroup* indicates the name of the higher level shadow group to which *lgroup* is connected.

When *lgroup* is the only object connected to *hgroup*, *hgroup* will automatically be removed upon disconnecting *lgroup*. However, when *lgroup* is the only object connected to *hgroup*, and *hgroup* is connected to a higher level shadow group, disconnection will result in an error. In such a case, disconnect *hgroup* from its higher level shadow group, and then disconnect *lgroup*.

You cannot disconnect *lgroup* if the disconnection may result in a change in the status of any existing shadow volume within *hgroup*.

Conditions that do not allow you to disconnect a shadow group from a shadow group of the mirror type or a shadow group of the stripe type are as below.

- When disconnecting from a higher level shadow group of the mirror type

For example, you cannot disconnect *lgroup* from a shadow group of the mirror type specified by *hgroup* if a shadow volume exists within that group, and *lgroup* is the only object connected to *hgroup*.

- When disconnecting from a shadow group of the stripe type

You cannot disconnect a lower level shadow group from a shadow group of the stripe type with an existing shadow volume, or from a shadow group of the stripe type connected to a higher level shadow group.

-R

Remove

Removes the shadow group definition specified by *group*. *class* indicates the name of the shadow class to which *group* belongs.

Shadow disks and lower level shadow groups connected to *group* will be disconnected. The disconnected shadow disk will have the original type attribute (single or undefined).

The definition cannot be removed when a shadow volume exists within *group*, or when *group* is connected to a higher level shadow group.

SUB OPTIONS

Sub options are as follows:

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

When using the -C option and defining a new group name with the -h option, a new shadow group, *hgroup*, is automatically created. This option sets *attribute* to *value* for the created *hgroup* attribute.

attribute indicates the attribute name, and *value* indicates the attribute value. An equal sign (=) is always necessary between *attribute* and *value*. When specifying multiple attributes, each specifier must be separated by a comma (,).

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

You can specify one of the following combinations to *attribute* and *value*.

If multiple attributes are specified and any error occurs, the entire process is canceled.

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

Sets the type attribute of *hgroup*.

mirror

Sets the type attribute to "mirror".

stripe

Sets the type attribute to "stripe".

width=*blks* (default is 32)

Sets the stripe width of *hgroup*. *blks* indicates the stripe width in blocks (decimal numbers). One block is 512 bytes. For *blks*, you can specify an integer that is the power of two, equal to or larger than one and equal to or smaller than the available size of the smallest shadow group among groups specified by *lgroup*,... If *hgroup* is not the stripe type, this option will cause in an error.

-c *class*

class indicates the name of the shadow class to which the shadow group as the object of the operation belongs.

-g *group* (when using -R)

group indicates the name of the shadow group that is the object of the operation.

-h *hgroup* (when using -C,-D)

hgroup indicates the name of the higher level shadow group to which the lower level shadow group as the object of the operation is connected or is to be connected.

-l *lgroup* (when using -D)

lgroup indicates the name of the lower level shadow group as the object of the operation.

-l *lgroup*,... (when using -C)

lgroup indicates the name of the lower level shadow group as the object of the operation. To connect multiple shadow groups, separate each group name with a comma (,).

RETURNED VALUE

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

D.17 sdxshadowvolume - Shadow volume operations

SYNOPSIS

```
sdxshadowvolume -F -c class [-v volume,...]  
sdxshadowvolume -M -c class {-g group | -d disk} -v volume -s size  
sdxshadowvolume -N -c class [-v volume,...] [-e mode=val]  
sdxshadowvolume -R -c class -v volume
```

DESCRIPTION

Use sdxshadowvolume to perform operations on shadow volumes specified by *volume*. This command can be executed with superuser access privileges only.

PRIMARY OPTIONS

You have a choice of the following options.

-F

offline

Stops one or more shadow volumes specified by *volume*,... By default all shadow volumes within *class* are turned off. Volumes turned off cannot be accessed.

This results in an error if the volume is in use.

-M

Make

Creates a shadow volume specified by *volume* to the highest level shadow group specified by *group* or to a shadow disk of the single type specified by *disk*. *size* indicates the number of blocks on *volume*, *class* indicates the name of the shadow class associated with *group* or *disk*.

The access mode of the created shadow volume is ro (read-only). Synchronization copying is not performed on shadow volumes of the mirror type. The attribute of just resynchronization mode is off. In addition, the physical slice attribute is always off even though the slice configuring the shadow volume is registered with partition table. A maximum of 224 shadow volumes can be created to the same *group* or *disk*.

When creation is completed, the shadow volume is activated and can be accessed using the following special files:

```
/dev/sfdsk/class/rdisk/volume [Linux2.4]  
dev/sfdsk/class/dsk/volume
```

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

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

The features of created shadow volumes when *group* is the mirror type and the stripe type are as below.

- When group is a mirror type

In a shadow group of the mirror type, shadow volumes of the mirror type with mirror-multiplexing equal to the number of shadow disks or lower level shadow groups connected (maximum of eight) are created. When only one shadow disk or one lower level shadow group is connected, the created shadow volume will not be mirrored.

Synchronization copying is not performed on shadow volumes even if the shadow volumes of the mirror type are created. To ensure synchronization for a shadow volume of the mirror type, the mirror volume must be properly synchronized with GDS that manages the mirror volume corresponding to the shadow volume.

- When group is a stripe group

In a shadow group of the stripe type, shadow volumes of the stripe type with columns equal to the number of shadow disks or lower level shadow groups connected (maximum of 64) are created. When only one shadow disk or one lower level shadow group is connected, a shadow volume cannot be created.

-N

oNline

Activates one or more shadow volumes specified by *volume*,... By default all shadow volumes within *class* are activated. The activated shadow volumes become accessible.

Synchronization copying is not performed on shadow volumes if the shadow volumes of the mirror type are activated. To ensure synchronization for a shadow volume of the mirror type, the mirror volume must be properly synchronized with GDS that manages the mirror volume corresponding to the shadow volume.

-R

Remove

Removes a shadow volume specified by *volume* and releases the disk area that the shadow group or the shadow disk of the single type was occupying.

This results in an error if the shadow volume is active.

No data stored on *volume* will be lost due to this removal of *volume*.

SUB OPTIONS

Sub options are as follows:

-c class

class indicates the name of the shadow class to which the shadow volume as the object of the operation belongs, or in which the shadow volume is to be created.

-d disk (when using -M)

disk indicates the name of a shadow disk of the single type in which the shadow volume of the single type will be created.

-e mode=val (when using **-N**)

Specifies the access mode for one or more shadow volumes that will be activated.

You can specify either of the following options to *val*.

rw

Sets the access mode to read and write.

ro

Sets the access mode to read only. Opening a read-only volume in the write mode will result in an error.

Although shadow volumes are activated in the access mode specified by *val*, the access mode attributes for the shadow volumes will remain unchanged. The access mode specified by *val* ("current access mode") is valid only while the shadow volume is active and will become invalid once the shadow volume is stopped. When the shadow volume is restarted, it will start in the mode according to the access mode attribute ("default access mode") unless the access mode is specified by this option.

In order to start a shadow volume that is already activated on the current node in a different access mode, you must first stop the shadow volume.

-g group (when using **-M**)

group indicates the name of a shadow group in which the shadow volume will be created.

-s size (when using **-M**)

Specifies the size of the *volume* being created in blocks (decimal numbers). One block is 512 bytes. When *group* is the stripe type, the size of *volume* created will be the *size* rounded up to a common multiple of the stripe width multiplied by stripe columns and the cylinder size. In other cases, the size of *volume* created will be the *size* rounded up to the integer multiple of the cylinder size.

-v volume (when using **-M,-R**)

volume indicates the name of the shadow volume as the object of the operation.

-v volume,... (when using **-F,-N**)

volume,... indicates the names of one or more shadow volumes as the objects of the operation. To specify multiple shadow volumes, separate each volume name with a comma (,).

RETURNED VALUE

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

D.18 Volume Creation Using Command

In this section, operation outline of volume creation is explained. Please use as a reference when configuring the environment. For details, see the Command Reference.



Information

For the order of mirroring system disks, see "USAGE EXAMPLES" in "D.11 sdxroot – Root file system mirroring definition and cancellation [Itanium]."



Note For PRIMECLUSTER Systems

In order to define the configuration of GDS objects such as classes and volumes, PRIMECLUSTER resources must be registered in advance. For details about registering resources, see "Appendix H Shared Disk Unit Resource Registration."

(1) Creating a mirror volume

The following example shows the procedures for creating a volume by mirroring physical disks named sda and sdb.

1) Registering disks to class

Register the physical disks with a class. When the specified class does not exist, it will be created automatically.

Example) Registering physical disks sda and sdb with local class "Class1", and name these disks Disk1 and Disk2.

```
# sdxdisk -M -a type=local -c Class1 -d sda=Disk1,sdb=Disk2
```



Information

When registering disks with a shared class, it is necessary to use the -a option to specify the scope attribute. See an example below.

```
# sdxdisk -M -c Class1 -a type=shared,scope=node1:node2 -d sda=Disk1,sdb=Disk2
```

2) Connecting the disks to a mirror group

Connect the disks to a mirror group. When the specified mirror group does not exist, it will be created automatically.

Example) Connecting "Disk1" and "Disk2" to mirror group "Group1."

```
# sdxdisk -C -c Class1 -g Group1 -d Disk1,Disk2
```

3) Creating a mirror volume

Create a volume within the mirror group.

Example) Creating a volume of 1,000 blocks within mirror group "Group1", and assigning a volume name "Volume1."

```
# sdxvolume -M -c Class1 -g Group1 -v Volume1 -s 1000
```

After returning from the command, synchronization copying will automatically begin.

4) Confirming the completion of procedure

Confirm that the synchronization copying is complete.

Example) Confirming synchronization copying of volume "Volume1" is complete.

If all the displayed slices' STATUS fields are "ACTIVE", synchronization copying is complete.

If the synchronization copying is still in progress, "COPY" will be displayed in the STATUS field. Using the -e long option, you can check the progress of the synchronization copying. For details, see "D.6 sdxinfo - Display object configuration and status information."

(2) Creating a single volume

The following example shows the procedures for creating a single volume using a physical disk named sda.

1) Registering a disk to class

Register the physical disk with a class. When the specified class does not exist, it will be created automatically.

Example) Registering physical disk sda to shared class "Class1", which is shared on node1 and node2, and assigning the name "Disk1."

```
# sdxdisk -M -c Class1 -a type=shared,scope=node1:node2 -d sda=Disk1:single
```

2) Creating a single volume

Create a volume within the single disk.

Example) Creating a volume of 1,000 blocks within single disk "Disk 1", and assigning a volume name "Volume1."

```
# sdxvolume -M -c Class1 -d Disk1 -v Volume1 -s 1000
```

(3) Creating a large-capacity volume (using concatenation)

The following example shows the procedures for creating a volume by concatenating physical disks named sda and sdb.

1) Registering disks to class

Register the physical disks with a class. When the specified class does not exist, it will be created automatically.

Example) Registering physical disks sda and sdb with shared class "Class1", which is shared on node1 and node2, and naming those disks Disk1 and Disk2.

```
# sdxdisk -M -c Class1 -a type=shared,scope=node1:node2 -d sda=Disk1,sdb=Disk2
```

2) Connecting the disks to a concatenation group

Connect the disks to a concatenation group. When the specified concatenation group does not exist, it will be created automatically.

Example) Connecting "Disk1" and "Disk2" to concatenation group "Group1."

```
# sdxdisk -C -c Class1 -g Group1 -d Disk1,Disk2 -a type=concat
```

3) Creating a large-capacity volume

Create a volume within the concatenation group.

Example) Creating a volume of 1,000,000,000 blocks within concatenation group "Group1", and assigning a volume name "Volume1."

```
# sdxvolume -M -c Class1 -g Group1 -v Volume1 -s 1000000000 -a pslice=off
```

(4) Creating a stripe volume

The following example shows the procedures for creating a volume by striping physical disks named sda and sdb.

1) Registering disks to class

Register the physical disks with a class. When the specified class does not exist, it will be created automatically.

Example) Registering physical disks sda and sdb with shared class "Class1", which is shared on node1 and node2, and naming those disks Disk1 and Disk2.

```
# sdxdisk -M -c Class1 -a type=shared,scope=node1:node2 -d sda=Disk1,sdb=Disk2
```

2) Connecting the disks to a stripe group

Connect the disks to a stripe group. When the stripe group does not exist, it will be created automatically.

Example) Connecting "Disk1" and "Disk2" to stripe group "Group1."

```
# sdxdisk -C -c Class1 -g Group1 -d Disk1,Disk2 -a type=stripe,width=32
```

3) Creating a stripe volume

Create a volume within the stripe disk.

Example) Creating a volume of 1,000 blocks within stripe group "Group1", and assigning a volume name "Volume1."

```
# sdxvolume -M -c Class1 -g Group1 -v Volume1 -s 1000 -a pslice=off
```

(5) Creating a mirror volume (Combining striping and mirroring)

The following example shows the procedures for creating a volume by constructing stripe groups with physical disks sda and sdb and other physical disks sdc and sdd respectively and then mirroring the two stripe groups.

1) Registering disks to class

Register the physical disks with a class. When the specified class does not exist, it will be created automatically.

Example) Registering physical disks sda, sdb, sdc and sdd to shared class "Class1", which is shared on node1 and node2, and assigning the names "Disk1", "Disk2", "Disk3" and "Disk4" respectively.

```
# sdxdisk -M -c Class1 -a type=shared,scope=node1:node2 \
-d sda=Disk1,sdb=Disk2,sdc=Disk3,sdd=Disk4
```

2) Connecting the disks to a stripe group

Connect the disks to a stripe group. When the stripe group does not exist, it will be created automatically.

Example) Connecting "Disk1" and "Disk2" to stripe group "Group1."

```
# sdxdisk -C -c Class1 -g Group1 -d Disk1,Disk2 -a type=stripe,width=32
```

Connecting Disk3 and Disk4 to stripe group Group2.

```
# sdxdisk -C -c Class1 -g Group2 -d Disk3,Disk4 -a type=stripe,width=32
```

3) Connecting the stripe groups to a mirror group

Connect the stripe groups to a mirror group. When the specified mirror group does not exist, it will be created automatically.

Example) Connecting the stripe group "Group1" and "Group2" to mirror group "Group3."

```
# sdxgroup -C -c Class1 -h Group3 -l Group1,Group2 -a type=mirror
```

The "-a type=mirror" option is omissible.

4) Creating a mirror volume

Create a volume within the highest level mirror group.

Example) Creating a volume of 1,000 blocks within mirror group "Group3", and assigning a volume name "Volume1."

```
# sdxvolume -M -c Class1 -g Group3 -v Volume1 -s 1000 -a pslice=off
```

After returning from the command, synchronization copying will automatically begin.

5) Confirming the completion of procedure

Confirm that the synchronization copying is complete.

Example) Confirming synchronization copying of volume "Volume1" is complete.

```
# sdxinfo -S -o Volume1
OBJ   CLASS  GROUP  DISK   VOLUME  STATUS
-----
slice Class1  Group3 Group1 Volume1 ACTIVE
slice Class1  Group3 Group2 Volume1 COPY
```

If the all the displayed slices' STATUS fields are "ACTIVE", synchronization copying is complete. If the synchronization copying is still in progress, "COPY" will be displayed in the STATUS field. Using the -e long option, you can check the progress of the synchronization copying.

For details, see "D.6 sdxinfo - Display object configuration and status information."

D.19 Snapshot Creation Using Command

In this section, operation outline of snapshot creation using command is explained. Please use it as a reference when configuring the environment. For details, see the Command Reference.

1) Joining the proxy volume with master volume

Join the proxy volume with the master volume, and copy the master volume data to the proxy volume. Before joining them, stop the proxy volume.

Example) Master volume Volume1 and proxy volume Volume2 will be joined.

```
# sdxvolume -F -c Class1 -v Volume2
# sdxproxy Join -c Class1 -m Volume1 -p Volume2
```

After returning from the command, synchronization copying will automatically be performed.



Information

- If Class1 is a shared class, use the -e allnodes option with the sdxvolume -F command to stop Volume2 on all the nodes within the class scope.
- -m option and -p option can indicate groups as well as volumes. When indicating a group, all volumes within the group will be copied. When specifying groups, it is necessary to use the -a option.

2) Confirming the completion of copying

Confirm that the synchronization copying is complete.

Example) Confirming synchronization copying of proxy volume (Volume2) is complete.

```
# sdxinfo -S
OBJ   CLASS   GROUP   DISK   VOLUME  STATUS
-----
slice Class1  Group1  Disk1  Volume1  ACTIVE
slice Class1  Group1  Disk2  Volume1  ACTIVE
slice Class1  Group2  Disk3  Volume2  STOP
slice Class1  Group2  Disk4  Volume2  STOP
```

If all the displayed slices' STATUS fields are "STOP", synchronization copying is complete. If the synchronization copying is still in progress, "COPY" will be displayed in the STATUS field.

3) Creating a snapshot

In order to create a snapshot, part the proxy volume from the master volume after confirming that synchronization copying is complete.

```
# sdxproxy Part -c Class1 -p Volume2
```



Note

In order to create a snapshot properly, you should either stop the application, or secure integrity at the file system layer or database layer that is managing the data.

For instance, if you are using the file system, use the umount(8) command to unmount the file system before separating the proxy volume from the master volume.

4) Creating a backup

In order to make a data backup from the snapshot you created, follow the procedures of whatever backup tool you are using.

If you are not using the snapshot function after creating the backup, cancel the master-proxy relationship.

5) Rejoining the proxy volume with master volume

In order to synchronize the proxy volume with the master volume again, remove the created snapshot, and rejoin the proxy volume with the master volume. Before rejoining them, stop the proxy volume.



Information

If Class1 is a shared class, use the -e allnodes option with the sdxvolume -F command to stop Volume2 on all the nodes within the class scope.

```
# sdxvolume -F -c Class1 -v Volume2  
# sdxproxy Rejoin -c Class1 -p Volume2
```

Synchronization copying will be performed as in step 1).

When creating more than one backup using the same proxy volume, repeat steps 2) to 4).

6) Canceling the master-proxy relationship

Cancel the master-proxy relationship, and finish using the snapshot function.

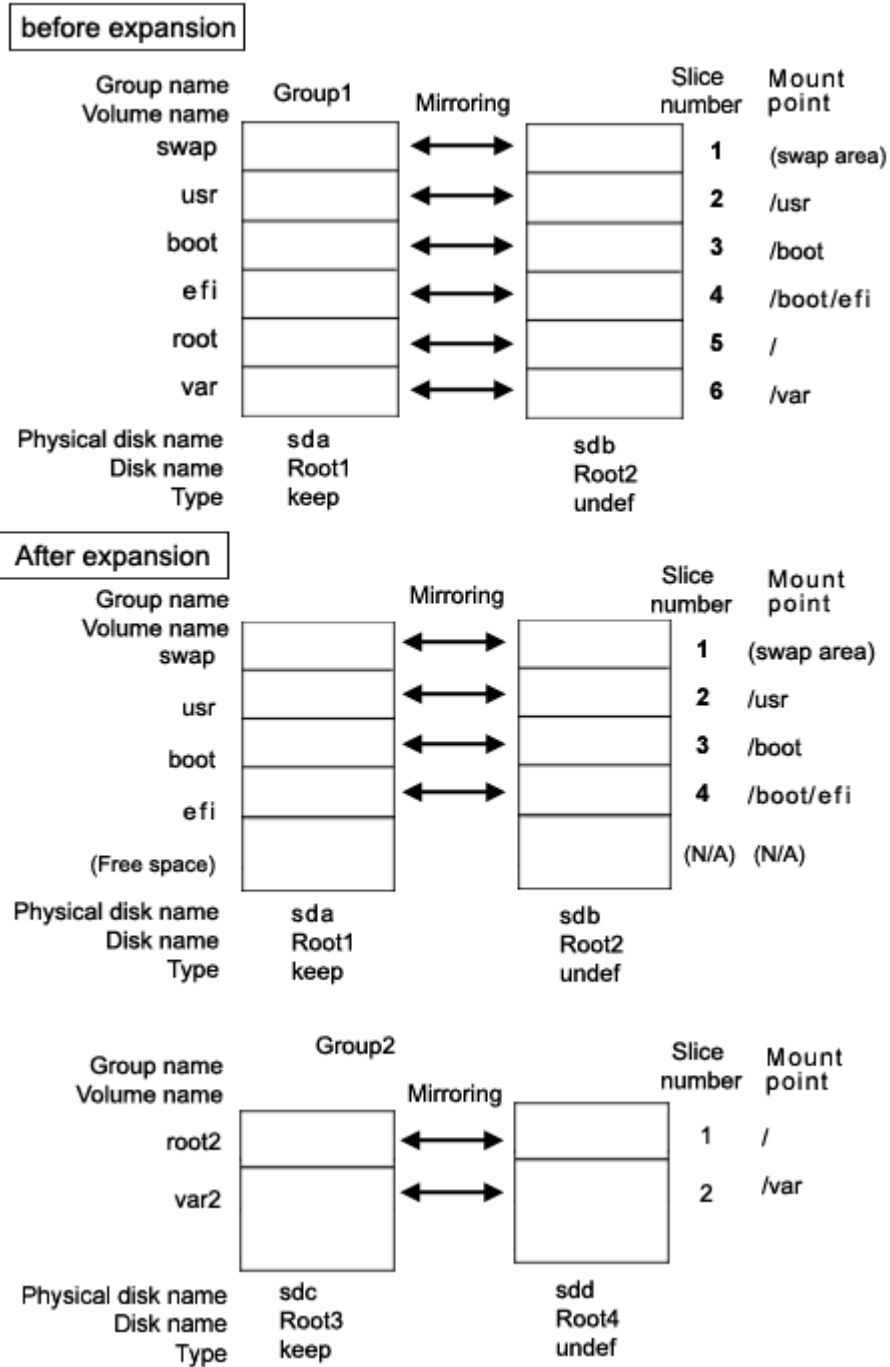
```
# sdxproxy Break -c Class1 -p Volume2
```

D.20 Volume Expansion Using Commands [Itanium]

This section explains the procedural flow for expansion of /, /usr, and /var file systems with the snapshot functions of GDS Snapshot that can be performed while operations are running.

[Procedure]

Assuming the configuration as below, the procedures for expanding /var file system space are explained. As shown in the following figure, the / file system must be transferred to another volume no matter whether or not it is expanded.





Note

The cylinder sizes of keep disks Root1 and Root3 must match.



Information

System disk mirroring is not a requirement. The configuration without sdb and sdd that are shown in the figure above is also supported. However, in systems that require high availability, it is recommended to use the mirroring configuration as above.

1) Mirroring the system disk

1-1) Stop the running application programs.

In order to ensure mirroring definition, all the running application programs must be stopped. When higher safety is required, create system disk backups.

1-2) Register the system disk to the root class.

```
# sdxdisk -M -c System -a type=root -d sda=Root1:sdb=Root2:undef
```

1-3) Connect the system disk to a group.

```
# sdxdisk -C -c System -g Group1 -d Root1,Root2 -v
1=swap:on,2=usr:on,3=boot:on,4=efi:on, 5=root:on,6=var:on
```

1-4) Check that mirroring definition is completed.

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

1-5) Reboot the system.

```
# shutdown -r now
```

1-6) Check that mirroring is in effect.

Using the mount(8) command and the sdxinfo command, make sure that the system disk has been mirrored properly.

2) Creating proxy volumes

Create proxy volumes for the / file system and the expanded /var file system. At this point, one keep disk only should be connected to the group that will include the proxy volumes. This example shows the procedure for creating proxy volumes in the following configuration.

Group name	Group2	Slice	Mount
Volume name		number	point
root2		1	(N/A)
var2		2	(N/A)
(Free space)		(N/A)	(N/A)
Physical disk name	sdc		
Disk name	Root3		
Type	keep		



Point

The disk must have sufficient free disk space following the last block of the proxy volume for the /var file system to be expanded.

2-1) Check the current volume sizes.

```
# sdxinfo -V -c System
```

OBJ	NAME	CLASS	GROUP	SKIP	JRM	1STBLK	LASTBLK	BLOCKS	STATUS
volume swap		System	Group1	off	on	0	1049759	1049760	ACTIVE
volume *		System	Group1	*	*	1049760	1071359	21600	PRIVATE
volume *		System	Group1	*	*	1071360	2946239	1874880	FREE
volume usr		System	Group1	off	on	2946240	2965679	19440	ACTIVE
volume boot		System	Group1	off	on	2965680	3096899	131220	ACTIVE
volume efi		System	Group1	off	on	3096900	3228119	131220	ACTIVE
volume root		System	Group1	off	on	3228120	3403079	174960	ACTIVE
volume var		System	Group1	off	on	3403080	3638519	235440	ACTIVE

For the volume sizes, check the BLOCKS field displayed with the sdxinfo -V command. In this example, the root size is 174960 blocks and the var size is 235440 blocks.

2-2) Register the disk with the root class.

```
# sdxdisk -M -c System -d sdc=Root3:keep,sdd=Root4:undef
```



Note

When registering multiple keep disks with a class together, as many or more undefined disks must also be registered.

2-3) Connect only one keep disk with a group.

```
# sdxdisk -C -c System -g Group2 -d Root3
```

2-4) Create volumes.

The volume sizes should be those shown in step 2-1).

```
# sdxvolume -M -c System -g Group2 -v root2 -s 174960
# sdxvolume -M -c System -g Group2 -v var2 -s 235440
```

2-5) Check the created volume sizes.

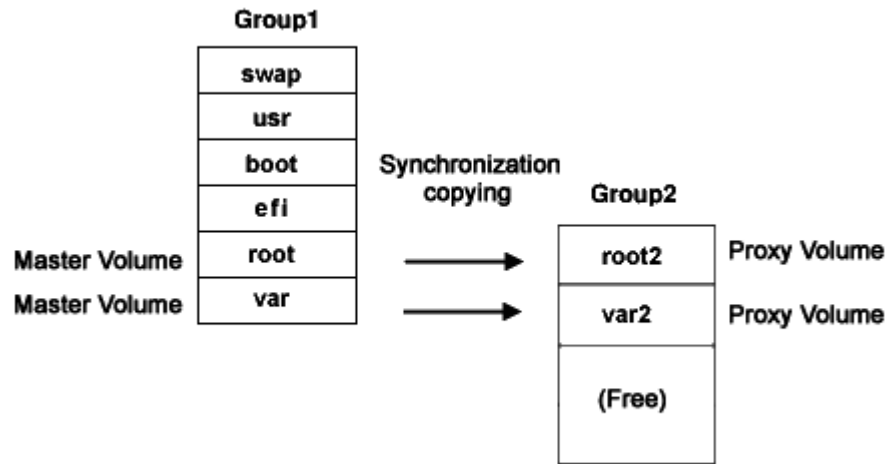
Make sure that the sizes of the volumes created in step 2-4) match the sizes shown in step 2-1).

```
# sdxinfo -V -c System
```

OBJ	NAME	CLASS	GROUP	SKIP	JRM	1STBLK	LASTBLK	BLOCKS	STATUS
volume swap		System	Group1	off	on	0	1049759	1049760	ACTIVE
volume *		System	Group1	*	*	1049760	1071359	21600	PRIVATE
volume *		System	Group1	*	*	1071360	2946239	1874880	FREE
volume usr		System	Group1	off	on	2946240	2965679	19440	ACTIVE
volume boot		System	Group1	off	on	2965680	3096899	131220	ACTIVE
volume efi		System	Group1	off	on	3096900	3228119	131220	ACTIVE
volume root		System	Group1	off	on	3228120	3403079	174960	ACTIVE
volume var		System	Group1	off	on	3403080	3638519	235440	ACTIVE
volume *		System	Group2	*	*	0	21599	21600	PRIVATE
volume root2		System	Group2	off	on	21600	196559	174960	ACTIVE
volume var2		System	Group2	off	on	196560	431999	235440	ACTIVE
volume *		System	Group2	*	*	432000	3376079	2944080	FREE

3) Joining the proxy volumes

Join the created proxy volumes with the volumes for / and /var (master volumes) to copy data in / and /var to the proxy volumes. This example shows the procedure for joining the proxy volumes in the following configuration.



3-1) Join the proxy volumes.

```
# sdxvolume -F -c System -v root2,var2
# sdxproxy Join -c System -m root -p root2
# sdxproxy Join -c System -m var -p var2
```

3-2) Check that synchronization copying is completed.

```
# sdxinfo -S -c System
OBJ   CLASS  GROUP  DISK  VOLUME  STATUS
-----
slice System Group2 Root1  swap    ACTIVE
slice System Group2 Root2  swap    ACTIVE
slice System Group1 Root1  usr     ACTIVE
slice System Group1 Root2  usr     ACTIVE
slice System Group1 Root1  boot    ACTIVE
slice System Group1 Root2  boot    ACTIVE
slice System Group1 Root1  efi     ACTIVE
slice System Group1 Root2  efi     ACTIVE
slice System Group1 Root1  root    ACTIVE
slice System Group1 Root2  root    ACTIVE
slice System Group1 Root1  var     ACTIVE
slice System Group1 Root2  var     ACTIVE
slice System Group2 Root3  root2   STOP
slice System Group2 Root3  var2    COPY
```

If the synchronization copying is in progress, "COPY" will be displayed in the STATUS field for the proxy volume slice. If "STOP" is displayed, synchronization copying is completed.

4) Parting the proxy volumes

With successful synchronization copying, the master volumes and the proxy volumes become equivalent. By parting those master volumes and proxy volumes, snapshots of the master volumes can be created on the proxy volumes.

4-1) Secure consistency of the file systems.

In order to secure consistency of snapshot file systems, it is necessary to prevent file system update. However, some file systems such as /, /usr, and /var are required for system operation and cannot be cancel the mounting statuses during system operation. There, use the following methods to minimize writing to the system disk and writing not updated on the system disk.

- a) Boot the system in single user mode (optional)
- b) Stop application programs having write access to the system disks (optional)
- c) Execute the sync(1) command to write file system data updated on the memory but not written yet.

Even if a), b), and c) are all put in action, file system update will not be completely prevented. Therefore, snapshot file systems may have inconsistency as in the case with after system panic occurrence.

When a), b), c) are all put in action, the snapshot file systems will be like those after system panic occurrence in single user mode.

When only c) is put in action skipping a) and b), the snapshot file systems will be like those after system panic occurrence during system operation.

In both situations, inconsistency may occur in the file systems, and it is necessary to check and repair consistency as instructed in step 5-1).

4-2) Part the proxy volumes.

```
# sdxproxy Part -c System -p root2,var2
```

4-3) When the system was rebooted in single user mode as instructed in step 4-1) a), reboot it in multi-user mode.

4-4) When the application programs were stopped as instructed in step 4-1) b), start them.

5) Configuring the alternative boot environment

Configure the environment to boot from the proxy volumes.

5-1) Check and repair the file systems on the proxy volumes.

The file systems on the proxy volumes may have inconsistency, and it is necessary to check and repair them with the fsck(8) command.

```
# fsck /dev/sfdsk/System/dsk/root2
# fsck /dev/sfdsk/System/dsk/var2
```

5-2) Configure the alternative boot environment.

```
# sdxproxy Root -c System -p root2,var2
```

With successful alternative boot environment configuration, the following message will be output.

```
SDX:sdxproxy: INFO: completed definitions of alternative boot environment:
  current-boot-device=Root1 Root2
  alternative-boot-device=Root3
```

Be sure to take a note of the output boot device names for the current boot environment (current-boot-device values) and those for the alternative boot environment (alternative-boot-device values).

6) Expanding snapshots

6-1) Cancel the master-proxy relationship.

```
# sdxproxy Break -c System -p root2
# sdxproxy Break -c System -p var2
```

6-2) Expand the snapshot volume size.

This example shows the procedure for expanding the size of the snapshot volume for /var to 706320 blocks.

```
# sdxvolume -S -c System -v var2 -s 706320
```

6-3) Expand the snapshot ext3 file system size.

This example shows the procedure for expanding the size of the snapshot ext3 file system for /var to 706320 blocks.

```
# resize2fs /dev/sfdsk/System/dsk/var2 706320
```



Note

For details about the resize2fs command, see the description of resize2fs(8) in the manual.

7) Mirroring snapshots

Add a disk to the group that includes the snapshot volumes to mirror the snapshot volumes.

```
# sdxdisk -C -c System -g Group2 -d Root4
```

8) Switching to the alternative boot environment

By switching the environment to the alternative boot environment, replace the /var file system with the expanded volume.

8-1) Boot the system from the alternative boot environment.

Among boot devices shown in the EFI boot manager's boot option selector window, select one from the boot devices for the alternative boot environment indicated by the message in step 5-2).

```
EFI Boot Manager ver 1.10

Please select a boot option

Root1
Root2
Root3 ← Select
...

Use ↑ and ↓ to change option(s). Use Enter to select an option
```

8-2) Check that the boot was successful.

Using the mount(8) command the sdxinfo command, check that the system was booted normally in the alternative boot environment and that no errors exist in the GDS object statuses. As needed, check also that the file systems for the alternative boot environment are correct and that applications can run properly.



Information

If the system was not booted normally, restore the original boot environment. To restore the original boot environment, among boot devices shown in the EFI boot manager's option selector window, select one from the boot devices for the current boot environment indicated by the message in step 5-2).

```
EFI Boot Manager ver 1.10

Please select a boot option

Root1 ← Select
Root2
Root3
...

Use ↑ and ↓ to change option(s). Use Enter to select an option
```

9) Deleting unnecessary volumes

After checking that the system was booted normally in the alternative boot environment, delete the volumes for the / file system and the original /var file system that build the previous boot environment.

```
# sdxvolume -F -c System -v root,var
# sdxvolume -R -c System -v root
# sdxvolume -R -c System -v var
```

Appendix E GDS Messages

This section explains the messages for GDS drivers, daemons and commands.

E.1 Web-Based Admin View Messages (0001-0099)

For message numbers 0001 to 0099, see "Web-Based Admin View Operation Guide" for details.

E.2 Driver Messages

Driver messages are output to a log file or console via the syslog interface.

Message Numbers

The message numbers used to identify messages described in E.2.1 to E.2.3 do not appear in messages actually output by GDS drivers.

Variable Names

Italicized words in the messages are variable names, and the actual output will vary depending on the situation. The meaning and the format of the variable names used in the message explanations are described below.

Variable names	Descriptions
<i>v_devno</i>	Volume device number (hexadecimal number)
<i>v_maj</i>	Volume major number (decimal number)
<i>v_min</i>	Volume minor number (decimal number)
<i>p_devno</i>	Physical slice device number (hexadecimal number)
<i>p_maj</i>	Physical slice major number (decimal number)
<i>p_min</i>	Physical slice minor number (decimal number)
<i>pdevtree</i>	Physical disk device tree path name. e.g.) /dev/sda
<i>pdevinst</i>	Physical disk instance name. e.g.) sd77
<i>blknodk</i>	I/O requested disk block number (decimal number)
<i>blknosl</i>	I/O requested slice block number (decimal number)
<i>length</i>	I/O transfer requested block size (in bytes, decimal number)
<i>resid</i>	Residing block size (in bytes, decimal number)
<i>errno</i>	System call error number (decimal number)
<i>oflag</i>	Open flag (hexadecimal number)
<i>second</i>	Elapsed time (in seconds, decimal number)
<i>lbolt</i>	Elapsed time after system boot (in ticks, hexadecimal number)
<i>details</i>	Details

Explanation

Messages output by the driver are shown below in the order of severity. There are three levels of severity.

Level of importance	Descriptions	facility.level
PANIC	This message is displayed when an event that will stop the system is detected.	kern.notice
WARNING	This message is output when an abnormal event is detected. It will not always result in an error.	kern.warning
NOTICE	This message is output to record driver operation and the like.	kern.notice

facility.level is the priority of a message passed from the GDS driver to syslogd(8). The output destinations of messages are defined in the /etc/syslog.conf configuration file of syslogd(8) and can be changed by modifying the definitions in /etc/syslog.conf. For details, see the description about syslog.conf(5) in this manual.

E.2.1 Warning Messages (22000-22099)

22000 WARNING: sfdsk: read error on slice:

volume info : devno(maj,min)=v_devno(v_maj,v_min)
device info : devno(maj,min)=p_devno(p_maj,p_min)
 devname=pdevtree (pdevinst)
error info : blk in disk=blknodk, blk in slice=blknos/
 length=length, resid=resid, errno=errno

Explanation

Read request sent to slice terminated abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

22001 WARNING: sfdsk: write error on slice:

volume info : devno(maj,min)=v_devno(v_maj,v_min)
device info : devno(maj,min)=p_devno(p_maj,p_min)
 devname=pdevtree (pdevinst)
error info : blk in disk=blknodk, blk in slice=blknos/
 length=length, resid=resid, errno=errno

Explanation

Write request sent to slice terminated abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk driver.

22002 WARNING: sfdsk: read error on disk:

device info : devno(maj,min)=*p_devno(p_maj,p_min)*
devname=*pdevtree (pdevinst)*
error info : blk in disk=*blknodk*, blk in slice=*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.

22003 WARNING: sfdsk: write error on disk:

device info : devno(maj,min)=*p_devno(p_maj,p_min)*
devname=*pdevtree (pdevinst)*
error info : blk in disk=*blknodk*, blk in slice=*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.

22004 WARNING: sfdsk: read and writeback error on slice:

volume info : devno(maj,min)=*v_devno(v_maj,v_min)*
device info : devno(maj,min)=*p_devno(p_maj,p_min)*
devname=*pdevtree (pdevinst)*
error info : blk in disk=*blknodk*, blk in slice=*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.

22005 WARNING: sfdsk: open error on slice:

volume info : devno(maj,min)=*v_devno(v_maj,v_min)*
device info : devno(maj,min)=*p_devno(p_maj,p_min)*
devname=*pdevtree (pdevinst)*
error info : blk in disk=*blknodk*, blk in slice=*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.

22006 WARNING: sfdsk: open error on disk:
device info : devno(maj,min)=*p_devno(p_maj,p_min)*
 devname=*pdevtree (pdevinst)*
error info : oflag=*oflag*, errno=*errno*

Explanation

Open request sent to disk returned abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

22007 WARNING: sfdsk: close error on disk:
device info : devno(maj,min)=*p_devno(p_maj,p_min)*
 devname=*pdevtree (pdevinst)*
error info : oflag=*oflag*, errno=*errno*

Explanation

Close request sent to disk returned abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

22008 WARNING: sfdsk: NVURM read error on disk:
volume info : devno(maj,min)=*v_devno(v_maj,v_min)*
device info : devno(maj,min)=*p_devno(p_maj,p_min)*
 devname=*pdevtree (pdevinst)*
error info : blk in disk=*blknodk*, blk in slice=*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 on the disk for just resynchronization.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

22009 WARNING: sfdsk: NVURM write error on disk:
volume info : devno(maj,min)=*v_devno(v_maj,v_min)*
device info : devno(maj,min)=*p_devno(p_maj,p_min)*
 devname=*pdevtree (pdevinst)*
error info : blk in disk=*blknodk*, blk in slice=*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 on the disk for just resynchronization.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

22010 WARNING: sfdsk: volume status log write error on disk:
device info : devno(maj,min)=*p_devno(p_maj,p_min)*
 devname=*pdevtree (pdevinst)*
error info : blk in disk=*blknodk*, blk in slice=*blknosl*
 length=*length*, resid=*resid*, errno=*errno*

Explanation

Write request of volume status log sent to disk terminated abnormally.
Volume status log records if the volume closed normally in the event of a system crash.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

22011 WARNING: sfdsk: failed to abort I/O requests on disk:
device info : devno(maj,min)=*p_devno(p_maj,p_min)*
 devname=*pdevtree (pdevinst)*
error info : errno=*errno*

Explanation

The request to cancel I/O requests on the physical device indicated in device info (an ioctl request for the mphd or mplb driver) ended abnormally.

Resolution

Collect investigation material and contact your local customer support.

22012 WARNING: sfdsk: hook for device is not GDS entry.

Explanation

Inconsistency was found in management information of disks registered with GDS, but it does not affect the operation.

Resolution

No resolution is required.

E.2.2 Information Messages (24000-24099)

24000 sfdsk: driver started up

Explanation

The driver has been installed into the system.

24001 sfdsk: received shutdown request

Explanation

A shutdown request from the sdxservd daemon has been received.

24002 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.

24003 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.

24004 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.

24005 NOTICE: sfdsk: trying to open slice:

volume info: devno(maj,min)=*v_devno(v_maj,v_min)*
device info: devno(maj,min)=*p_devno(p_maj,p_min)*
devname=*pdevtree(pdevinst)*

Explanation

A request for opening the slice is issued.

24006 NOTICE: sfdsk: copy timeout. no response from sdxservd daemon:

volume info: devno(maj,min)=*v_devno(v_maj,v_min)*

Explanation

Synchronization copying was not performed since there was no response from the sdxservd daemon.

24007 NOTICE: sfdsk: processing has taken long time on disk:

device info: devno(maj,min)=*p_devno(p_maj, p_min)*
devname=*pdevtree (pdevinst)*
request info: lapsed seconds=*second*, start lbolt=*lbolt*
details

Explanation

While *second* seconds have passed since the I/O request indicated in *details* was issued for the physical device indicated in device info, the I/O request is not complete yet.

24008 NOTICE: sfdsk: processing has taken long time on volume:
volume info: devno(maj,min)=v_devno(v_maj,v_min)
request info: lapsed seconds=second, start lbolt=lbolt
details

Explanation

While *second* seconds have passed since the I/O request indicated in *details* was issued for the volume indicated in volume info, the I/O request is not complete yet.

24009 NOTICE: sfdsk: abort I/O requests on disk:
device info: devno(maj,min)=p_devno(p_maj,p_min)
devname=pdevtree (pdevinst)

Explanation

The process to cancel I/O requests on the physical disk indicated in device info has started.

24010 NOTICE: sfdsk: succeeded to abort I/O requests on disk:
device info: devno(maj,min)=p_devno(p_maj,p_min)
devname=pdevtree (pdevinst)

Explanation

The request to cancel I/O requests on the physical device indicated in device info (an ioctl request for the mphd or mplb driver) ended normally.

24011 NOTICE: sfdsk: HBA='adapter' varyio=enabled

Explanation

Variable length IO is available for *adapter*, a SCSI host bus adapter, and SCSI disks connected to *adapter*.

24012 NOTICE: sfdsk: HBA='adapter' varyio=disabled

Explanation

Variable length IO is unavailable for *adapter*, a SCSI host bus adapter, and SCSI disks connected to *adapter*.

24013 NOTICE: sfdsk: variable length IO is enabled

Explanation

The sfdsk driver can make use of variable length IO.

24014 NOTICE: sfdsk: variable length IO is disabled

Explanation

The sfdsk driver cannot make use of variable length IO.

E.2.3 Internal Error Messages (26000)

26000 severity: sfdsk: internal error, details

Explanation

Internal error occurred. *details* indicates the cause of the error.
severity indicates the severity of the message.

Resolution

Collect investigation material and contact your local customer support.

E.3 Daemon Messages

Daemon messages are output to the GDS log file.

/var/opt/FJSV/sdx/msglog/sdxservd.log

Output format to log file is as follows.

Mon Day HH:MM:SS SDX:daemon: severity: message

Mon gives the month the message was output, *Day* the date, *HH* the hour, *MM* the minute, *SS* the second, *daemon* the daemon program name, *severity* the severity of the message, and *message* the message text.

Depending on the settings of the syslog, daemon message will also be output to syslog log file and console.

Message Numbers

The message numbers used to identify messages described in E.3.1 to E.3.5 do not appear in messages actually output by GDS daemons.

Variable Names

Italicized words in the messages are variable names, and the actual output will vary depending on the situation. The meaning and the format of the variable names used in the message explanations are described below.

Variable names	Descriptions
<i>class</i>	Class name
<i>disk</i>	Disk name
<i>group</i>	Group name
<i>lgroup</i>	Lower level group name
<i>hgroup</i>	Higher level group name
<i>volume</i>	Volume name
<i>disk.volume</i>	Slice name
<i>object.volume</i>	Slice name
<i>proxy</i>	Proxy object name
<i>object</i>	Object name
<i>type</i>	Type attribute value
<i>status</i>	Object status
<i>device</i>	Physical disk name (sdX, emcpowerX) X indicates the device identifier.
<i>pslice</i>	Physical slice name (sdXn, emcpowerXn) X indicates the device identifier and n indicates the slice number. The device tree file name may also be used. e.g.) "sd@2,0:a,raw"
<i>psdevtree</i>	Device tree path name to a physical slice e.g.) /dev/sd@2,0:a
<i>v_devno</i>	Volume device number (hexadecimal number)
<i>v_maj</i>	Volume major number (decimal number)
<i>v_min</i>	Volume minor number (decimal number)
<i>p_devno</i>	Physical slice device number (hexadecimal number)
<i>p_maj</i>	Physical slice major number (decimal number)
<i>p_min</i>	Physical slice minor number (decimal number)

Variable names	Descriptions
<i>blknodk</i>	I/O requested disk block number (decimal number)
<i>blknosl</i>	I/O requested slice block number (decimal number)
<i>length</i>	I/O transfer requested block size (in bytes, decimal number)
<i>resid</i>	Residing block size (in bytes, decimal number)
<i>errno</i>	System call error number (decimal number)
<i>sdxerrno</i>	GDS defined internal error number (decimal number)
<i>node</i>	Node identifier, or node name
<i>oflag</i>	Open flag (hexadecimal number)
<i>ioctlcmd</i>	The ioctl command name
<i>timeofday</i>	Character string indicating current time and date
<i>details</i>	Details
<i>sdxfunc</i>	GDS function name
<i>exitstat</i>	Value indicating command end status (decimal number)
<i>cmdline</i>	Character indicating command line

Explanation

Messages output by the daemon are shown below in the order of its severity. There are four levels of severity.

Level of importance	Descriptions	facility.level
HALT	This message is output when an event that will halt all services provided by GDS is detected.	user.crit
ERROR	This message is output when an event that will halt some services is detected.	user.err
WARNING	This message is output when an abnormal event is detected. It will not always result in a service halt.	user.warning
INFO	This message is output to record the daemon operation and the like. Usually, it can be ignored.	user.info

facility.level is the priority of a message passed from the GDS driver to syslogd(8). The output destinations of messages are defined in the /etc/syslog.conf configuration file of syslogd(8) and can be changed by modifying the definitions in /etc/syslog.conf. For details, see the description about syslog.conf(5) in this manual.

E.3.1 Halt Messages (40000-40099)

40000 HALT: failed to create a new thread, errno=*errno*

Explanation

Function `pthread_create()` terminated abnormally. Process cannot be continued. The daemon process will be exited.

Resolution

When error number information *errno* is insufficient for identifying the cause and recovery, collect investigation material and contact your local customer support.

40001 HALT: cannot open driver administrative file, errno=*errno*

Explanation

GDS driver(`sfdisk`) administrative file cannot be opened. Process cannot be continued. The daemon process will be exited.

This message is output when files under `/dev/sfdsk` directory cannot be accessed.

Resolution

Collect investigation material, and contact your local customer support.

40002 HALT: startup failed, no enough address space

Explanation

Startup failure due to unsuccessful memory allocation. Process cannot be continued. The daemon process will be exited.

Resolution

Confirm you have sufficient memory or swap area.

40003 HALT: failed to respawn *daemon* daemon, *osfunc*=*osfunc*, *errno*=*errno*

Explanation

The *daemon* terminated abnormally and failed to restart. The failure was caused by abnormal termination of OS *osfunc* function.

The error number is *errno*. This message is output via `syslog`.

Resolution

When error number information is insufficient for identifying the cause, collect investigation material, and contact your local customer support.

40004 HALT: cannot start node-down recovery for remote node *node*, no enough space, *osfunc*=*osfunc*, *errno*=*errno*

Explanation

Could not recover the crashed remote node *node* due to unsuccessful memory allocation. Process cannot be continued. The daemon process will be exited.

Resolution

The OS *osfunc* function terminated abnormally, and error number is *errno*. Confirm you have sufficient memory or swap area, and recover.

E.3.2 Error Messages (42000-42099)

42000 ERROR: read error on *status slice object.volume*, class=*class*:
volume info:devno(maj,min)=*v_devno(v_maj,v_min)*
device info:devno(maj,min)=*p_devno(p_maj,p_min)*
 devname=*device*
error info: blk in disk=*blknodk*, blk in slice=*blknosl*,
 length=*length*, resid=*resid*, errno=*errno*

Explanation

Read request sent to slice *object.volume* in *status* status terminated abnormally. Read request sent to volume configured by this slice, or to slice accessible in isolation was returned as error. You must attempt recovery as soon as possible since the application may cease normal operation.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and promptly recover the disk.

42001 ERROR: write error on *status slice object.volume*, class=*class*:
volume info:devno(maj,min)=*v_devno(v_maj,v_min)*
device info:devno(maj,min)=*p_devno(p_maj,p_min)*
 devname=*device*
error info: blk in disk=*blknodk*, blk in slice=*blknosl*,
 length=*length*, resid=*resid*, errno=*errno*

Explanation

Write request sent to slice *object.volume* in *status* status terminated abnormally. Write request sent to volume configured by this slice, or to slice accessible in isolation was returned as error. You must attempt recovery as soon as possible since the application may cease normal operation.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and promptly recover the disk.

42002 ERROR: open error on *status slice object.volume*, class=*class*:
volume info:devno(maj,min)=*v_devno(v_maj,v_min)*
device info:devno(maj,min)=*p_devno(p_maj,p_min)*
 devname=*device*
error info: oflag=*oflag*, errno=*errno*

Explanation

Open request sent to slice *object.volume* in *status* status terminated abnormally. Open request sent to volume configured by this slice, or to slice accessible in isolation was returned as error. You must attempt recovery as soon as possible since the application may cease to operate normally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and promptly recover the disk.

42003 ERROR: read error and writeback error on *status slice object.volume*, class=*class*:
volume info:devno(maj,min)=*v_devno(v_maj,v_min)*
device info:devno(maj,min)=*p_devno(p_maj,p_min)*
devname=*device*
error info: blk in disk=*blknodk*, blk in slice=*blknosl*,
length=*length*, resid=*resid*, errno=*errno*

Explanation

Read request and writeback request sent to slice *object.volume* in *status* status terminated abnormally. Read request sent to volume configured by this slice, or to slice accessible in isolation was returned as error.

You must attempt recovery as soon as possible since the application may not operate normally. Writeback is a process to read data from other slices in the event of a read error.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and promptly recover the disk.

42006 ERROR: *volume*: closed down volume, class=*class*

Explanation

volume has closed down.

Resolution

Promptly attempt recovery by identifying the cause of failure by referring to GDS log message that was output immediately before the error.

For information on recovery, see "F.1.3 Volume Status Abnormality."

42007 ERROR: *class*: cannot startup class, no valid configuration database, *sdxfcrmo=errno*

Explanation

Could not startup *class*, since valid *class* configuration database could not be found. This message is output when all disks (or the majority of disks) registered with *class* are unavailable.

Resolution

See "F.1.4 Class Status Abnormality."

42008 ERROR: *class*: cannot startup class, too few valid configuration database replicas, *sdxfcrmo=errno*

Explanation

Could not startup *class*, due to insufficient number of valid databases. This message is output when many disks registered with *class* are unavailable.

Resolution

See "F.1.4 Class Status Abnormality."

42009 ERROR: *class*: closing down class, no valid configuration database

Explanation

class was closed since no valid *class* configuration database was found. This message is output when all disks (or the majority of disks) registered with *class* are unavailable.

Resolution

See "F.1.4 Class Status Abnormality."

42010 ERROR: *class*: closing down class, too few valid configuration database replicas

Explanation

class was closed due to insufficient number of valid database.
This message is output when many disks registered with *class* are unavailable.

Resolution

See "F.1.4 Class Status Abnormality."

42011 ERROR: failed to send request message on node *node*, *details*

Explanation

Sending request message from *node* was unsuccessful.

Resolution

Collect investigation material and contact your local customer support.

42012 ERROR: timeout on receiving reply message from node *node*, *details*

Explanation

Timeout occurred while receiving a reply message from a remote node *node*.

Resolution

Collect investigation material and contact your local customer support.

42013 ERROR: rejected request message on remote node *node*, *details*

Explanation

Processing a request message on a remote node *node* was unsuccessful.

Resolution

Investigate the *node* message log and take necessary actions. If recovery is impossible, collect investigation material and contact your local customer support.

42014 ERROR: *class*: failed to start *type* volumes, *status* volume *volume* exists, *node*=*node*

Explanation

Starting volumes within the class *class* failed on the node *node* since the *status* volume *volume* exists. *type* is the *class* type attribute.

Resolution

volume is in abnormal status. First, you must recover normal status. For information on recovery, see "F.1.3 Volume Status Abnormality."

42015 ERROR: *class*: failed to start and standby *type* volumes, *status* volume *volume* exists, *node*=*node*

Explanation

Starting and putting on standby volumes within the class *class* failed on the node *node* since the *status* volume *volume* exists. *type* is the *class* type attribute.

Resolution

volume is in abnormal status. First, you must recover normal status. For information on recovery, see "F.1.3 Volume Status Abnormality."

42016 ERROR: *class*: failed to stop and standby *type* volumes, *status* volume *volume* exists, *node*=*node*

Explanation

Stopping and putting on standby volumes within the class *class* failed on the node *node* since the *status* volume *volume* exists. *type* is the *class* type attribute.

Resolution

volume is in abnormal status. First, you must recover normal status. For information on recovery, see "F.1.3 Volume Status Abnormality."

42017 ERROR: *class*: failed to stop *type* volumes, *status* volume *volume* exists, *node*=*node*

Explanation

Stopping volumes within the class *class* failed on the node *node* since the *status* volume *volume* exists. *type* is the *class* type attribute.

Resolution

volume is in abnormal status. First, you must recover normal status. For information on recovery, see "F.1.3 Volume Status Abnormality."

42018 ERROR: *class*: failed to start *type* volumes, class closed down, *node*=*node*

Explanation

Starting volumes failed since the class *class* has been closed down. *type* is the *class* type attribute.

Resolution

Recover closed *class*. There may be multiple disk failures. Identify the cause based on object status, GDS log message, and syslog message.
For information on recovery, see "F.1.4 Class Status Abnormality."

42019 ERROR: *class*: failed to start and standby *type* volumes, class closed down, *node*=*node*

Explanation

Starting and putting on standby volumes failed. since the class *class* has been closed down. *type* is the *class* type attribute.

Resolution

Recover closed *class*. There may be multiple disk failures. Identify the cause based on object status, GDS log message, and syslog message.
For information on recovery, see "F.1.4 Class Status Abnormality."

42020 ERROR: *class*: failed to stop and standby *type* volumes, class closed down, *node*=*node*

Explanation

Stopping and putting on standby volumes failed. since the class *class* has been closed down. *type* is the *class* type attribute.

Resolution

Recover closed *class*. There may be multiple disk failures. Identify the cause based on object status, GDS log message, and syslog message.
For information on recovery, see "F.1.4 Class Status Abnormality."

42021 ERROR: *class*: failed to stop *type* volumes, class closed down, node=*node*

Explanation

Stopping volumes failed since the class *class* has been closed down. *type* is the *class* type attribute.

Resolution

Recover closed *class*. There may be multiple disk failures. Identify the cause based on object status, GDS log message, and syslog message.

For information on recovery, see "F.1.4 Class Status Abnormality."

42022 ERROR: *class*: closing down class, cluster-wide lock failure, sdxerno=*sdxerno*

Explanation

Abnormal exclusive control between cluster system nodes occurred. Since the process cannot be continued, *class* will be closed.

Resolution

Collect investigation material and contact your local customer support.

42023 ERROR: *class*: cannot startup class, cluster-wide lock failure, sdxerno=*erno*

Explanation

Abnormal exclusive control between cluster system nodes occurred. Since the process cannot be continued, *class* could not be started.

Resolution

Collect investigation material and contact your local customer support.

42024 ERROR: *class*: closing down class, cluster communication failure, sdxerno=*sdxerno*

Explanation

Transmission failure between cluster system nodes occurred. Since the process cannot be continued, *class* will be closed.

Resolution

Collect investigation material and contact your local customer support.

42025 ERROR: *class*: cannot operate in cluster environment, created when cluster control facility not ready

Explanation

The class *class* cannot be used in a cluster environment since it was created when the cluster control facility was inactive. This message is output when one of the following operations is performed.

- a. After *class* was created resource registration was performed on a node where resource registration was not complete.
- b. In a cluster environment where resource registration is complete, *class* was created in single user mode.
- c. A single node with *class* was changed over to a cluster system.

Resolution

See (1) in "F.1.10 Cluster System Related Error."

42026 ERROR: *proxy*: failed to copy with OPC, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

While performing copying between proxy volume *proxy* and master volume with the OPC function, an I/O error occurred, and the copying process failed. "source" specifies the original slice name, "target" the mirror slice name, and "class" to which the original slice and mirror slice belong.

Resolution

Identify the cause by referring to the log messages for GDS, disk driver, and Fujitsu Storage Systems ETERNUS that were output right before the occurrence of the error, and restore.

42027 ERROR: *proxy*: failed to copy with EC, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

While performing copying between proxy volume *proxy* and master volume with the EC function, an I/O error occurred, and the copying process failed. "source" specifies the original slice name, "target" the mirror slice name, and "class" to which the original slice and mirror slice belong.

Resolution

Identify the cause by referring to the log messages for GDS, disk driver, and Fujitsu Storage Systems ETERNUS that were output right before the occurrence of the error, and restore.

42028 ERROR: *proxy*: failed to copy with TimeFinder, source=*disk*, target=*disk*, class=*class*

Explanation

Copying failed due to an I/O error caused while conducting copying between the proxy group *proxy* and the master group with TimeFinder. "source" is the name of the copy source disk, "target" is the name of the copy destination disk, and "class" is the name of the class to which the copy source and destination disks belong.

Resolution

Identify the cause by referring to the log messages for GDS, disk driver and Symmetrix device that were output right before the occurrence of the error, and restore.

42029 ERROR: *proxy*: failed to copy with SRDF, source=*disk*, target=*disk*, class=*class*

Explanation

Copying failed due to an I/O error caused while conducting copying between the proxy group *proxy* and the master group with SRDF. "source" is the name of the copy source disk, "target" is the name of the copy destination disk, and "class" is the name of the class to which the copy source and destination disks belong.

Resolution

Identify the cause by referring to the log messages for GDS, disk driver and Symmetrix device that were output right before the occurrence of the error, and restore.

42030 ERROR: *proxy*: failed to start OPC, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

An error occurred between proxy volume *proxy* and the master volume when OPC was initiated. "source" indicates the copy source slice name, "target" indicates the copy destination slice name, and "class" indicates a class to which the copy source and destination slices belong.

Resolution

Identify the cause based on log messages of GDS, disk drivers, Fujitsu Storage Systems ETERNUS and such, and promptly recover the disk.

42031 ERROR: *proxy*: failed to stop OPC, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

When stopping OPC between proxy volume *proxy* and master volume, an error occurred. "source" specifies the original slice name, "target" the mirror slice name, and "class" to which the original slice and mirror slice belong.

Resolution

Identify the cause by referring to the log messages for GDS, disk driver, Fujitsu Storage Systems ETERNUS that were output right before the occurrence of the error, and restore.

42032 ERROR: *proxy*: failed to start EC session, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

An error occurred between proxy volume *proxy* and the master volume when EC was initiated. "source" indicates the copy source slice name, "target" indicates the copy destination slice name, and "class" indicates a class to which the copy source and destination slices belong.

Resolution

Identify the cause based on log messages of GDS, disk drivers, Fujitsu Storage Systems ETERNUS and such, and promptly recover the disk.

42033 ERROR: *proxy*: failed to stop EC session, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

When stopping EC session between proxy volume *proxy* and master volume, an error occurred. "source" specifies the original slice name, "target" the mirror slice name, and "class" to which the original slice and mirror slice belong.

Resolution

Identify the cause by referring to the log messages for GDS, disk driver, and Fujitsu Storage Systems ETERNUS that was output right before the occurrence of the error, and restore.

42034 ERROR: *proxy*: failed to suspend EC session, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

When temporarily suspending EC session between proxy volume *proxy* and master volume, an error occurred. "source" specifies the original slice name, "target" the mirror slice name, and "class" to which the original slice and mirror slice belong.

Resolution

Identify the cause by referring to the log messages for GDS, disk driver, and Fujitsu Storage Systems ETERNUS that were output right before the occurrence of the error, and restore.

42035 ERROR: *proxy*: failed to resume EC session, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

When resuming EC session between proxy volume *proxy* and master volume, an error occurred. "source" specifies the original slice name, "target" the mirror slice name, and "class" to which the original slice and mirror slice belong.

Resolution

Identify the cause by referring to the log messages for GDS, disk driver, and Fujitsu Storage Systems ETERNUS that were output right before the occurrence of the error, and restore.

42036 ERROR: *proxy*: failed to establish BCV pair, STD=*disk*, BCV=*disk*, class=*class*

Explanation

Error occurred while establishing a BCV pair between the proxy group *proxy* and the master group. "STD" is the disk name of the standard device, "BCV" is the disk name of the BCV device, and "class" is the name of the class to which the standard and BCV devices belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or Symmetrix device, and promptly recover the disk.

42037 ERROR: *proxy*: failed to cancel BCV pair, STD=*disk*, BCV=*disk*, class=*class*

Explanation

Error occurred while canceling a BCV pair between the proxy group *proxy* and the master group. "STD" is the disk name of the standard device, "BCV" is the disk name of the BCV device, and "class" is the name of the class to which the standard and BCV devices belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or Symmetrix device, and promptly recover the disk.

42038 ERROR: *proxy*: failed to split BCV pair, STD=*disk*, BCV=*disk*, class=*class*

Explanation

Error occurred while splitting a BCV pair between the proxy group *proxy* and the master group. "STD" is the disk name of the standard device, "BCV" is the disk name of the BCV device, and "class" is the name of the class to which the standard and BCV devices belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or Symmetrix device, and promptly recover the disk.

42039 ERROR: *proxy*: failed to re-establish BCV pair, STD=*disk*, BCV=*disk*, class=*class*

Explanation

Error occurred while re-establishing a BCV pair between the proxy group *proxy* and the master group. "STD" is the disk name of the standard device, "BCV" is the disk name of the BCV device, and "class" is the name of the class to which the standard and BCV devices belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or Symmetrix device, and promptly recover the disk.

42040 ERROR: *proxy*: failed to establish SRDF pair, source=*disk*, target=*disk*, class=*class*

Explanation

Error occurred while establishing an SRDF pair between the proxy group *proxy* and the master group. "source" is the name of the source disk, "target" is the name of the target disk, and "class" is the name of the class to which the source and target disks belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or Symmetrix device, and promptly recover the disk.

42041 ERROR: *proxy*: failed to cancel SRDF pair, source=*disk*, target=*disk*, class=*class*

Explanation

Error occurred while canceling an SRDF pair between the proxy group *proxy* and the master group. "source" is the name of the source disk, "target" is the name of the target disk, and "class" is the name of the class to which the source and target disks belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or Symmetrix device, and promptly recover the disk.

42042 ERROR: *proxy*: failed to split SRDF pair, source=*disk*, target=*disk*, class=*class*

Explanation

Error occurred while splitting an SRDF pair between the proxy group *proxy* and the master group. "source" is the name of the source disk, "target" is the name of the target disk, and "class" is the name of the class to which the source and target disks belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or Symmetrix device, and promptly recover the disk.

42043 ERROR: *proxy*: failed to re-establish SRDF pair, source=*disk*, target=*disk*, class=*class*

Explanation

Error occurred while re-establishing an SRDF pair between the proxy group *proxy* and the master group. "source" is the name of the source disk, "target" is the name of the target disk, and "class" is the name of the class to which the source and target disks belong.

Resolution

Identify the cause by referring to log messages previously output on the GDS, disk driver or Symmetrix device, and promptly recover the disk.

42044 ERROR: *disk* is is bound to RAW device. disabled disk, class=*class*

Explanation

The disk *disk* registered with GDS was found that it was bound to a RAW device, and the disk was disabled.

Resolution

Cancel the bind to a RAW device using the raw(8) command for the disk registered with GDS. You may not bind disks registered with GDS to RAW devices.

E.3.3 Warning Messages (44000-44099)

44000 WARNING: read error on *status slice object.volume*, class=*class*:
volume info:devno(maj,min)=*v_devno(v_maj,v_min)*
device info:devno(maj,min)=*p_devno(p_maj,p_min)*
 devname=*device*
error info: blk in disk=*blknodk*, blk in slice=*blknosl*,
 length=*length*, resid=*resid*, errno=*errno*

Explanation

Read request sent to slice *object.volume* in *status* status terminated abnormally. Slice with abnormality will be detached.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44001 WARNING: write error on *status slice object.volume*, class=*class*:
volume info:devno(maj,min)=*v_devno(v_maj,v_min)*
device info:devno(maj,min)=*p_devno(p_maj,p_min)*
 devname=*device*
error info: blk in disk=*blknodk*, blk in slice=*blknosl*,
 length=*length*, resid=*resid*, errno=*errno*

Explanation

Write request sent to slice *object.volume* in *status* status terminated abnormally. Slice with abnormality will be detached.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44002 WARNING: open error on *status slice object.volume*, class=*class*:
volume info:devno(maj,min)=*v_devno(v_maj,v_min)*
device info:devno(maj,min)=*p_devno(p_maj,p_min)*
 devname=*device*
error info: oflag=*oflag*, errno=*errno*

Explanation

Open request sent to slice *object.volume* in *status* status terminated abnormally. Slice with abnormality will be detached.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44003 WARNING: read error and writeback error on *status slice object.volume*, class=*class*:
volume info:devno(maj,min)=*v_devno(v_maj,v_min)*
device info:devno(maj,min)=*p_devno(p_maj,p_min)*
devname=*device*
error info: blk in disk=*blknodk*, blk in slice=*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.

44004 WARNING: NVURM write error on disk *disk*, class=*class*:
volume info:devno(maj,min)=*v_devno(v_maj,v_min)*
volume=*volume*, class=*class*
device info:devno(maj,min)=*p_devno(p_maj,p_min)*
devname=*device*
error info: blk in disk=*blknodk*, blk in slice=*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.

44005 WARNING: *object.volume*: detached *status slice* by an I/O error, class=*class*

Explanation

Since an I/O error occurred on slice *object.volume* in *status* status, the slice was detached from the volume.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44006 WARNING: open error on private slice *pslice*, oflag=*oflag*, errno=*errno*

Explanation

Open request sent to disk private slice *pslice* terminated abnormally. It will automatically search for a normal alternate disk and attempt recovery.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44007 WARNING: read error on private slice *p_devno(p_maj,p_min)*, offset=*blkno*sl, length=*length*, resid=*resid*, errno=*errno*

Explanation

Read request sent to disk private slice *p_devno(p_maj,p_min)* terminated abnormally. It will automatically search for a normal alternate disk and attempt recovery.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44008 WARNING: write error on private slice *p_devno(p_maj,p_min)*, offset=*blkno*sl, length=*length*, resid=*resid*, errno=*errno*

Explanation

Write request sent to disk private slice *p_devno(p_maj,p_min)* terminated abnormally. It will automatically search for a normal alternate disk and attempt recovery.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44009 WARNING: close error on private slice *p_devno(p_maj,p_min)*, errno=*errno*

Explanation

Close request sent to disk private slice *p_devno(p_maj,p_min)* terminated abnormally. It will automatically search for a normal alternate disk and attempt recovery.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44010 WARNING: *sdxfunc: pslice: open error, errno=errno*

Explanation

The open request for the physical slice *pslice* terminated abnormally.

The following messages may be output when a node is booted, investigation material is collected (*pcsnap* or *sdxsnap.sh* is executed), or the physical disk information update menu in GDS Management View is executed.

- a. WARNING: *pd_get_info: pslice: open error, errno=6*
- b. WARNING: *pd_set_orig_all: pslice: open error, errno=6*

Resolution

A disk failure may have occurred. Identify the cause by referring to disk driver log messages, and recover the disk.

However, in the following situations, GDS is behaving normally and messages a) and b) may be ignored.

- In the messages a) and b), *pslice* is a physical slice of a disk unit previously removed. In this situation, delete the device special file for *pslice*, and these messages will no longer be output.
- In the messages a) and b), *pslice* is a physical disk slice of a physical disk not registered with GDS.

44011 WARNING: *sdxfunc: pslice: read error, errno=errno*

Explanation

Read request sent to physical slice *pslice* terminated abnormally.

Resolution

Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44012 WARNING: *sdxfunc: pslice: write error, erro=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.

44013 WARNING: *sdxfunc: pslice: ioctl error, request=ioctlcmd, erro=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.

44014 WARNING: *sdxfunc: pslice: close error, erro=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.

44015 WARNING: *volume: failed to enable JRM, no available NVURM space, class=class*

Explanation

Resuming just resynchronization process on *volume* was unsuccessful due to insufficient number of normal disks.

Resolution

Check the disk status within *class*. Disk failure may have occurred. Identify the cause by referring to disk driver log message, and recover the disk.

44016 WARNING: *volume: failed to retrieve NVURM from disk disk, class=class*

Explanation

NVURM read request of *volume* from *disk* was unsuccessful. Just resynchronization will switch to copying of the entire disk.

NVURM is 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.

44017 WARNING: *disk: device: disabled disk, class=class*

Explanation

disk is disabled since its data is invalid. *device* is the physical disk name corresponding to *disk*.

Resolution

Disk configuration may have been wrongly changed or disk data could be damaged. Check for improper connection change of I/O cables and disk swap.

Also check for disk driver log message regarding the disk in question, and attempt recovery.

44018 WARNING: *volume*: volume synchronization failed, class=*class*

Explanation

Synchronization copying of *volume* was unsuccessful.

Resolution

Attempt recovery by identifying the cause of failure by referring to GDS log message and disk driver message that were output immediately before the error.

44019 WARNING: *volume*: volume just resynchronization failed, class=*class*

Explanation

Just resynchronization copying of *volume* was unsuccessful.

Resolution

Attempt recovery by identifying the cause of failure by referring to GDS log message and disk driver message that were output immediately before the error.

44020 WARNING: *class*: unknown class file found

Explanation

Class file *class* which does not exist in class database file was found while booting the system.

**44021 WARNING: invalid configuration database ID information, sdxerno=*sdxerno*, class=*class*:
*psdevtree***

Explanation

Since the configuration database ID information was invalid, private slice *psdevtree* was not used as the configuration database for class *class*.

Resolution

Disk configuration may have been wrongly changed or disk data could be damaged. Check for improper connection change of I/O cables and disk swap. Also check for disk driver log message regarding the disk in question, and attempt recovery.

44022 WARNING: *class*: too few valid configuration database replicas

Explanation

There are not sufficient valid configuration databases for *class*.
This message is displayed when the majority of disks registered with *class* are unavailable.
Leaving it as is may cause serious problems.

Resolution

For details, see "F.1.4 Class Status Abnormality."

**44023 WARNING: cannot open message logging file, 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.

**44024 WARNING: cannot write message logging file, erro=erro
/var/opt/FJSVsdx/msglog/sdxservd.log**

Explanation

Writing to GDS log file was unsuccessful. This message will be output via syslog. Although the message will not be output on GDS logfile, it does not affect other processes.

Resolution

Collect investigation material and contact your local customer support.

44025 WARNING: failed to reply message to node *node*, details

Explanation

Replying to remote node *node* was unsuccessful.

Resolution

Collect investigation material and contact your local customer support.

**44026 WARNING: class: failed to change class resource status on remote node *node*,
status=*new-status*, sdxermo=*sdxermo***

Explanation

Changing class resource status to *new-status* on a remote node *node* was unsuccessful.

Resolution

Collect investigation material and contact your local customer support.

44027 WARNING: sdxfunc: no enough address space, osfunc=*osfunc*, erro=*erro*

Explanation

The OS *osfunc* function returned an error. The error number is *erro*.

Resolution

Confirm you have sufficient memory or swap area.

44028 WARNING: respawned *daemon* daemon successfully

Explanation

Although *daemon* terminated abnormally, it completed normally by restarting. This message will be output via syslog.

Resolution

To identify the cause of *daemon* terminating abnormally, collect investigation material and contact your local customer support.

44029 WARNING: device: failed to restore VTOC on disk, sdxermo=*sdxermo*

Explanation

Recovering physical disk *device* format information was unsuccessful. Disk failure may have occurred.

Resolution

Use parted(8) command to recover format information.

44031 WARNING: *volume*: cannot copy to one or more slices in volume

Explanation

Copying process could not be performed on some slices within *volume*.

Resolution

Execute synchronization copying as needed.

44032 WARNING: *device*: write error, *ermo*=*ermo*

Explanation

A write error occurred in physical disk *device* and recovering the partition table of *device* failed. A disk failure may have occurred.

Resolution

To recover the partition table of *device*, use the parted(8) command. If *device* is write-locked by the disk unit's copy function, no action is required.

44033 WARNING: *device*: read error, *ermo*=*ermo*

Explanation

A read error occurred in physical disk *device* and recovering the partition table of *device* failed. A disk failure may have occurred.

Resolution

To recover the partition table of *device*, use the parted(8) command.

44036 WARNING: *proxy*: too many EC/OPC sessions

Explanation

The number of EC or OPC sessions within the physical disk (LU) or the disk array body has reached the upper limit of supported concurrent sessions. For this reason, a new EC or OPC session cannot be started. Copying is started using the soft copy function.

Resolution

To make copying by EC or OPC available, cancel the relationship between the proxy volume *proxy* and the master volume, wait until the running session ends and try this command again. Alternatively, according to need, cancel the running session using the sdxproxy Cancel command, the sdxproxy Break command, or [Operation]:[Proxy Operation]:[Break] and try this command again.

44037 WARNING: *proxy*: failed to start OPC, *source*=*disk.volume*, *target*=*disk.volume*, *class*=*class*

Explanation

An error occurred when OPC was started between the proxy volume *proxy* and the master volume. "source" is the copy source slice name, "target" is the copy destination slice name, and "class" is the name of the class to which the copy source slice and the copy destination slice belong. Copying is started between *proxy* and the master using the soft copy function.

Resolution

The cause of this error may derive from a faulty disk array unit. Identify the cause by referring to the log messages for GDS, disk driver, and Fujitsu Storage Systems ETERNUS that were output right before the occurrence of the error, and restore.

44038 WARNING: proxy: failed to start EC session, source=disk.volume, target=disk.volume, class=class

Explanation

An error occurred when an EC session was started between the proxy volume *proxy* and the master volume. "source" is the copy source slice name, "target" is the copy destination slice name, and "class" is the name of the class to which the copy source slice and the copy destination slice belong. Copying is started between *proxy* and the master using the soft copy function.

Resolution

The cause of this error may derive from a faulty disk array unit. Identify the cause by referring to the log messages for GDS, disk driver, and Fujitsu Storage Systems ETERNUS that were output right before the occurrence of the error, and restore.

E.3.4 Information Messages (46000-46199)

**46000 INFO: read error and writeback success on status slice object volume, class=class:
volume info:devno(maj,min)=v_devno(v_maj,v_min)
device info:devno(maj,min)=p_devno(p_maj,p_min)
devname=device,
error info: blk in disk=blknodk, blk in slice=blknosl,
length=length**

Explanation

Read request sent to slice *object.volume* in status *status* was unsuccessful, but was restored by writeback process.
Writeback is a process to read data from other slices in the event of a read error.

**46001 INFO: NVRAM configuration parameter has been updated:
parameter='param'
old = 'old_value'
new = 'new_value'**

Explanation

Parameter *param* value stored in NVRAM (non-volatile memory) on the body unit was updated from *old_value* to *new_value*.

46002 INFO: volume: temporarily disabled JRM, class=class

Explanation

Although just resynchronization process on *volume* will be temporarily disabled, it will automatically attempt recovery.

46003 INFO: disk: failed to connect spare disk for disk disk, group=group, class=class, sdxfcrmo=sdxfcrmo

Explanation

Connecting spare disk *disk* to *group* instead of *disk* was unsuccessful.

46004 INFO: *disk*: failed to connect spare disk for group *lgroup*, the highest level group=*hgroup*, class=*class*, sdxerno=*sdxerno*

Explanation

The attempt to connect spare disk *disk* to the highest level group *hgroup* in place of group *lgroup* failed.

46005 INFO: *disk*: connected spare disk for disk *disk*, group=*group*, class=*class*

Explanation

Connected spare disk *disk* to group instead of *disk*.

46006 INFO: *disk*: connected spare disk for group *lgroup*, the highest level group=*hgroup*, class=*class*

Explanation

Spare disk *disk* was connected to the highest level group *hgroup* in place of group *lgroup*.

46007 INFO: *group*: free blocks are reduced in group, class=*class*

Explanation

Free blocks of *group* were reduced.

46008 INFO: *volume*: reallocated NVURM space and enabled JRM successfully, class=*class*

Explanation

Just resynchronization process of *volume* was resumed.

46009 INFO: *volume*: retrieved NVURM from disk *disk* successfully, class=*class*

Explanation

NVURM read request of *volume* sent from *disk* was successful. Just resynchronization process will be resumed.

NVURM is volume update area map information stored in the disk for just resynchronization.

46010 INFO: *volume*: no need to retrieve NVURM, sdxerno=*sdxerno*, class=*class*

Explanation

NVURM was not retrieved. Entire copying of volume will start.

46011 INFO: *disk*: *pslice*: failed to open physical special files exclusively, erro=*ermo*

Explanation

Opening *disk* physical slice *pslice* exclusively was unsuccessful.

46012 INFO: *disk: device:* disk ID information is invalid, *sdxermo=sdxermo*

Explanation

Disk ID information of *disk* is invalid. *disk* will be automatically disabled.

46013 INFO: *disk: enabled disk, class=class*

Explanation

Disabled *disk* was enabled.

46014 INFO: *volume: volume synchronization started, class=class*

Explanation

Started synchronization copying process of *volume*.

46015 INFO: *volume: volume just resynchronization started, class=class*

Explanation

Started just resynchronization process of *volume*.

46016 INFO: *volume: volume synchronization canceled, class=class*

Explanation

Canceled synchronization copying of *volume*.

46017 INFO: *volume: volume just resynchronization canceled, class=class*

Explanation

Canceled just resynchronization process of *volume*.

46018 INFO: *volume: volume synchronization completed successfully, class=class*

Explanation

Synchronization copying process of *volume* completed successfully.

46019 INFO: *volume: volume just resynchronization completed successfully, class=class*

Explanation

Just resynchronization process of *volume* completed successfully.

46020 INFO: *object: failed to update configuration database, class=class*

Explanation

Updating configuration database was unsuccessful since all (or majority of) disks registered with *class* was unavailable. Usually, *class* will be closed following this message.

46021 INFO: sdxservd daemon started up

Explanation

The sdxservd daemon started up. GDS will be started.

46022 INFO: started local volumes, *timeofday*

Explanation

All volumes under local class have been started.

46023 INFO: started root volumes

Explanation

All volumes under root class have been started.

46024 INFO: stopped all services by shutdown request, *timeofday*

Explanation

All GDS services were stopped in response to a shutdown request.

46025 INFO: cannot open class database file, *ermo=ermo*

Explanation

Could not open class database file while booting the system. Will automatically attempt recovery.

46026 INFO: class database file corrupted

Explanation

Corrupted class database file was detected while booting the system. Will automatically attempt recovery.

46027 INFO: *class*: cannot open class file, *ermo=ermo*

Explanation

Could not open class file *class* while booting the system. Will automatically attempt recovery.

46028 INFO: *class*: class file corrupted

Explanation

Corrupted class file *class* was detected while booting the system. Will automatically attempt recovery.

46029 INFO: class database file updated successfully

Explanation

Class database file was updated.

46030 INFO: *class*: class file updated successfully

Explanation

Class file *class* was updated.

46031 INFO: cannot write class database file, *ermo=ermo*

Explanation

Could not write to class database file. Will automatically attempt recovery.

46032 INFO: *class*: cannot write class file, *ermo=ermo*

Explanation

Could not write to class file *class*. Will automatically attempt recovery.

**46033 INFO: cannot check configuration database ID information, *sdxfcrmo=sdxfcrmo*,
class=class:
*psdextree***

Explanation

Opening or reading private slice *psdextree* was unsuccessful, and could not check configuration database ID information of *class*.

**46034 INFO: cannot check configuration database, *sdxfcrmo=sdxfcrmo*, *class=class*:
*psdextree***

Explanation

Opening or reading private slice *psdextree* was unsuccessful, and could not check configuration database of *class*.

**46035 INFO: configuration database corrupted, *sdxfcrmo=sdxfcrmo*, *class=class*:
*psdextree ...***

Explanation

Since there was a check-sum error in the configuration database, private slice *psdevtree...* was not used as the configuration database for class *class*.

**46036 INFO: configuration database defeated, *sdxfcrmo=sdxfcrmo*, *class=class*:
*psdextree ...***

Explanation

The configuration database for class *class* stored on private slice *psdevtree...* was determined to be invalid in a validity check.

**46037 INFO: *class*: valid configuration database replicas exist on:
psdextree ...**

Explanation

The valid configuration database for class *class* was determined. *psdevtree...* is the private slice storing the valid configuration database.

46038 INFO: *class*: starting up class

Explanation

class will be started.

**46039 INFO: cannot update configuration database replica, *sdxerno*=*sdxerno*, *class*=*class*:
*psdextree***

Explanation

Updating a replica of the configuration database for class *class* stored on private slice *psdevtree* failed.

**46040 INFO: *class*: relocated configuration database replicas on:
psdextree ...**

Explanation

A replica of the configuration database for class *class* was relocated onto private slice *psdevtree...*

46041 INFO: *disk*: disconnected spare disk from group *group*, *class*=*class*

Explanation

Disconnected spare disk *disk* from *group*.

46042 INFO: *group*: free blocks are increased in group, *class*=*class*

Explanation

Number of free blocks in *group* increased.

46043 INFO: failed to create a new thread, *erno*=*erno*

Explanation

Function `pthread_create()` terminated abnormally.

46044 INFO: cannot open configuration parameter file *filename*, *erno*=*erno*

Explanation

Opening configuration parameter *filename* was unsuccessful.

46045 INFO: cannot read configuration parameter file, ermo=*ermo*

Explanation

Reading configuration parameter files was unsuccessful.

46046 INFO: received unexpected data from sfdsk driver and ignored

Explanation

Unexpected data was received from sfdsk driver and was ignored.

46047 INFO: received unexpected event from sfdsk driver and ignored, *details*

Explanation

Unexpected event was received from sfdsk driver and was ignored.
details displays the details about the event.

46048 INFO: *class*: class closed down, *node*=*node*

Explanation

class on *node* was closed.

**46049 INFO: command executed:
*cmdline***

Explanation

Command *cmdline* was executed.

**46050 INFO: command exited, *exit-status*=*exitstat*:
*cmdline***

Explanation

Processing *cmdline* is complete.

**46051 INFO: trying to execute command:
*cmdline***

Explanation

The *cmdline* command is about to be executed.

**46052 INFO: failed to execute command:
*cmdline***

Explanation

The *cmdline* command failed.

46053 INFO: *class*: changed class resource status on remote node *node*, *old-status* => *new-status*

Explanation

Class resource status on remote node *node* was changed from *old-status* to *new-status*.

46054 INFO: *class*: changed class resource status on current node *node*, *old-status* => *new-status*

Explanation

Class resource status on current node *node* was changed from *old-status* to *new-status*.

46055 INFO: *class*: started *type* volumes, *node=node*

Explanation

Starting volumes that belong to the class *class* was completed on the node *node*. *type* is the *class* type attribute.

46056 INFO: *class*: started and stood by *type* volumes, *node=node*

Explanation

Starting and putting on standby volumes that belong to the class *class* was completed on the node *node*. *type* is the *class* type attribute.

46057 INFO: *class*: stopped and stood by *type* volumes, *node=node*

Explanation

Stopping and putting on standby volumes that belong to the class *class* was completed on the node *node*. *type* is the *class* type attribute.

46058 INFO: *class*: stopped *type* volumes, *node=node*

Explanation

Stopping volumes that belong to the class *class* was completed on the node *node*. *type* is the *class* type attribute.

46059 INFO: cannot connect spare disk, cluster-wide lock failure, *class=class*, *sdxermo=sdxermo*

Explanation

Due to the occurrence of an abnormal exclusive control between cluster system nodes, the spare disk could not be connected.

46060 INFO: cannot connect spare disk, too few valid configuration database replicas, class=*class*, disk=*disk*

Explanation

Could not connect spare disk due to insufficient number of valid configuration databases.

46061 INFO: cannot connect spare disk, hot spare disabled, class=*class*, disk=*disk*

Explanation

Spare disk *disk* could not be connected since the hot spare is disabled.

46062 INFO: *class*: started class-down recovery for remote node *node*

Explanation

Closed class *class* on remote node *node* will be recovered.

46063 INFO: *class*: class-down recovery failed, already class-down on current node *node*

Explanation

Attempted recovery of closed class *class* on remote node *node*. Recovery was unsuccessful since the class was also closed on current node *node*.

46064 INFO: *class*: class-down recovery failed, sdxerno=*sdxerno*

Explanation

Recovering closed class *class* was unsuccessful.

46065 INFO: *class*: class-down recovery completed successfully

Explanation

Recovering closed class *class* was successful.

46066 INFO: *class*: started node-down recovery for remote node *node*

Explanation

Started node-down recovery on remote node *node*.

46067 INFO: *class*: started shutdown recovery for remote node *node*

Explanation

Started shutdown recovery on remote node *node*.

46068 INFO: *class*: node-down recovery failed, already class-down on current node *node*

Explanation

Recovering node-down was unsuccessful since *class* was in closed status on current node.

46069 INFO: *class*: shutdown recovery failed, already class-down on current node *node*

Explanation

Recovering shutdown was unsuccessful since *class* was in closed status on current node.

46070 INFO: *class*: node-down recovery failed, *sdxermo*=*sdxermo*

Explanation

Recovering *class* from node-down was unsuccessful.

46071 INFO: *class*: shutdown recovery failed, *sdxermo*=*sdxermo*

Explanation

Recovering *class* from shutdown was unsuccessful.

46072 INFO: *class*: node-down recovery completed successfully

Explanation

Recovering *class* from node-down completed successfully.

46073 INFO: *class*: shutdown recovery completed successfully

Explanation

Recovering *class* from shutdown completed successfully.

46074 INFO: *object.volume*: failed to update slice error information, class closed down, *class*=*class*

Explanation

Updating error information of slice *object.volume* was unsuccessful due to *class* in closed status.

46075 INFO: *volume*: failed to disable JRM, class closed down, *class*=*class*

Explanation

Disabling just resynchronization process on *volume* was unsuccessful due to *class* in closed status.

46076 INFO: *object.volume*: failed to detach slice, class closed down, *class*=*class*

Explanation

Detaching slice *object.volume* was unsuccessful due to *class* in closed status.

46077 INFO: *volume*: failed to restart volume, class closed down, class=*class*

Explanation

Restarting *volume* was unsuccessful due to *class* in closed status.

46078 INFO: open error on *status slice object.volume*, class=*class*

Explanation

Open request sent to slice *object.volume* in *status* status terminated abnormally.

46079 INFO: *class*: trying to identify class master, *details*

Explanation

Trying to identify the class master for shared class *class*.

46080 INFO: *class*: identified class master, node=*node*

Explanation

The master for shared class *class* has been identified as node *node*.

46081 INFO: *class*: searching class master

Explanation

Searching the master for shared class *class*.

46082 INFO: *class*: class master found, node=*node*

Explanation

The master for shared class *class* was found to be node *node*.

46083 INFO: *class*: class master not found

Explanation

The master for shared class *class* cannot be found.

46084 INFO: *class*: got class master privilege

Explanation

Trying to obtain master privilege for shared class *class*.

46085 INFO: *class*: broadcasted class master information to remote nodes

Explanation

The master information for shared class *class* has been broadcast to the remote nodes.

46086 INFO: *class*: received confirmations of class master information from remote node *node*

Explanation

Received master confirmation for shared class *class* from remote node *node*.

46087 INFO: waiting for outstanding event operations, *details*

Explanation

Waiting for event operations in process.

46088 INFO: completed outstanding event operations

Explanation

Event operations in process have been completed.

46089 INFO: *class*: trying to release class master privilege, *details*

Explanation

Trying to release master privilege for shared class *class*.

46090 INFO: *class*: released class master privilege

Explanation

Master privilege for shared class *class* has been released.

46091 INFO: *proxy*: started to copy with OPC, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

Copying between proxy volume *proxy* and master volume with the OPC function started. "source" specifies the original slice name, "target" the mirror slice name, and "class" the name of the class to which the original slice and mirror slice belong.

46092 INFO: *proxy*: completed copying with OPC, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

Copying between proxy volume *proxy* and master volume with the OPC function is completed. "source" specifies the original slice name, "target" the mirror slice name, and "class" the name of the class to which the original slice and mirror slice belong.

46093 INFO: *proxy*: canceled copying with OPC, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

Copying between proxy volume *proxy* and master volume with the OPC function was cancelled. "source" specifies the original slice name, "target" the mirror slice name, and "class" the name of the class to which the original slice and mirror slice belong.

46094 INFO: *proxy*: EC session started, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

EC session between proxy volume *proxy* and master volume started. "source" specifies the original slice name, "target" the mirror slice name, and "class" the name of the class to which the original slice and mirror slice belong.

46095 INFO: *proxy*: completed copying with EC, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

Copying between proxy volume *proxy* and master volume with the EC function is completed. "source" specifies the original slice name, "target" the mirror slice name, and "class" the name of the class to which the original slice and mirror slice belong.

46096 INFO: *proxy*: canceled copying with EC, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

Copying between proxy volume *proxy* and master volume with the EC function was cancelled. "source" specifies the original slice name, "target" the mirror slice name, and "class" the name of the class to which the original slice and mirror slice belong.

46097 INFO: *proxy*: EC session stopped, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

EC session between proxy volume *proxy* and master volume with the EC function was stopped. "source" specifies the original slice name, "target" the mirror slice name, and "class" the name of the class to which the original slice and mirror slice belong.

46098 INFO: *proxy*: EC session suspended, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

EC session between proxy volume *proxy* and master volume with the EC function has been temporarily suspended. "source" specifies the original slice name, "target" the mirror slice name, and "class" the name of the class to which the original slice and mirror slice belong.

46099 INFO: *proxy*: EC session resumed, source=*disk.volume*, target=*disk.volume*, class=*class*

Explanation

EC session between proxy volume *proxy* and master volume with the EC function was resumed. "source" specifies the original slice name, "target" the mirror slice name, and "class" the name of the class to which the original slice and mirror slice belong.

46100 INFO: *proxy*: established BCV pair, STD=*disk*, BCV=*disk*, class=*class*

Explanation

BCV pair between the proxy group *proxy* and the master group was established. "STD" is the disk name of the standard device, "BCV" is the disk name of the BCV device, and "class" is the name of the class to which the standard and BCV devices belong.

46101 INFO: *proxy*: completed copying with TimeFinder, source=*disk*, target=*disk*, class=*class*

Explanation

Copying between the proxy group *proxy* and the master group with TimeFinder is complete. "source" is the name of the copy source disk, "target" is the name of the copy destination disk, and "class" is the name of the class to which the copy source and destination disks belong.

46102 INFO: *proxy*: canceled copying with TimeFinder, source=*disk*, target=*disk*, class=*class*

Explanation

Copying between the proxy group *proxy* and the master group with TimeFinder was canceled. "source" is the name of the copy source disk, "target" is the name of the copy destination disk, and "class" is the name of the class to which the copy source and destination disks belong.

46103 INFO: *proxy*: canceled BCV pair, STD=*disk*, BCV=*disk*, class=*class*

Explanation

BCV pair between the proxy group *proxy* and the master group was canceled. "STD" is the disk name of the standard device, "BCV" is the disk name of the BCV device, and "class" is the name of the class to which the standard and BCV devices belong.

46104 INFO: *proxy*: split BCV pair, STD=*disk*, BCV=*disk*, class=*class*

Explanation

BCV pair between the proxy group *proxy* and the master group was split. "STD" is the disk name of the standard device, "BCV" is the disk name of the BCV device, and "class" is the name of the class to which the standard and BCV devices belong.

46105 vINFO: *proxy*: re-established BCV pair, STD=*disk*, BCV=*disk*, class=*class*

Explanation

BCV pair between the proxy group *proxy* and the master group was re-established. "STD" is the disk name of the standard device, "BCV" is the disk name of the BCV device, and "class" is the name of the class to which the standard and BCV devices belong.

46106 INFO: *proxy*: established SRDF pair, source=*disk*, target=*disk*, class=*class*

Explanation

SRDF pair between the proxy group *proxy* and the master group was established. "source" is the name of the source disk, "target" is the name of the target disk, and "class" is the name of the class to which the source and target disks belong.

46107 INFO: *proxy*: completed copying with SRDF, source=*disk*, target=*disk*, class=*class*

Explanation

Copying between the proxy group *proxy* and the master group with SRDF is complete. "source" is the name of the copy source disk, "target" is the name of the copy destination disk, and "class" is the name of the class to which the copy source and destination disks belong.

46108 INFO: *proxy*: canceled copying with SRDF, source=*disk*, target=*disk*, class=*class*

Explanation

Copying between the proxy group *proxy* and the master group with SRDF was canceled. "source" is the name of the copy source disk, "target" is the name of the copy destination disk, and "class" is the name of the class to which the copy source and destination disks belong.

46109 INFO: *proxy*: canceled SRDF pair, source=*disk*, target=*disk*, class=*class*

Explanation

SRDF pair between the proxy group *proxy* and the master group was canceled. "source" is the name of the source disk, "target" is the name of the target disk, and "class" is the name of the class to which the source and target disks belong.

46110 INFO: *proxy*: split SRDF pair, source=*disk*, target=*disk*, class=*class*

Explanation

SRDF pair between the proxy group *proxy* and the master group was split. "source" is the name of the source disk, "target" is the name of the target disk, and "class" is the name of the class to which the source and target disks belong.

46111 INFO: *proxy*: re-established SRDF pair, source=*disk*, target=*disk*, class=*class*

Explanation

SRDF pair between the proxy group *proxy* and the master group was re-established. "source" is the name of the source disk, "target" is the name of the target disk, and "class" is the name of the class to which the source and target disks belong.

46112 INFO: *file* not found

Explanation

The *file* file was not found.

46113 INFO: physical disk not found, devno(maj,min)=p_devno(p_maj,p_min)

Explanation

The physical disk registered with GDS was not found. *p_devno(p_maj,p_min)* is the device number (major number, minus number) of the physical disk used when it was registered with a class.

E.3.5 Internal Error Messages (48000)

48000 daemon: severity: module: internal error, details

Explanation

An internal error occurred. *details* gives the cause of the error, *daemon* gives the name of the daemon program, *severity* gives the severity of the message, and *module* the module name (usually the internal function name) where the event was detected.

The following messages may be output when an I/O error occurs.

a) sdxservd: ERROR: *module*: internal error, sdxfunc=dbrw_read_dc(), rval=-1, sdxerrno=4120

b) sdxservd: ERROR: *module*: internal error, sdxfunc=dbrw_read_dc(), rval=-1, sdxerrno=4121

The message c) below may be output when a node is booted. This means the return value of the OS read(2) system call was 0 and the error number was 2. If the return value is 0, the event is not an error and the error number has no meaning, but it is recorded as a WARNING because the return value is not normal (a positive value).

c) sdxservd: WARNING: *module*: internal error, osfunc=read, rval=0, errno=2

Resolution

In the event of a) or b), some other I/O error message will be output. Refer to the explanation of resolution for that message, and take necessary actions.

In the event of c), if no error message was output around the same period of time, GDS was behaving normally and the message c) may be ignored. If some other error message was output around the same period of time, refer to the explanation and resolution for that message, and take necessary actions. In other events, collect investigation material and contact your local customer support.

E.4 Command Messages

Command messages will be sent to standard output or standard error output. Output format is as follows.

SDX:command: severity: message

The *command* gives the command name, *severity* the severity of the message, and *message* the message text body.

Command message will also be output to the following log file.

/var/opt/FJSVsdx/msglog/sdxservd.log

Output format is as follows.

Mon Day HH:MM:SS SDX:daemon: severity: message

Mon gives the month the message was output, *Day* the date, *HH* the hour, *MM* the minute, *SS* the second, *daemon* the daemon program name, *severity* the severity of the message, and *message* the message text body.

Message Numbers

The message numbers used to identify messages described in E.4.1 to E.4.5 do not appear in messages actually output by GDS commands.

Variable Names

Italicized words in the messages are variable names, and the actual output will vary depending on the situation. The meaning and the format of the variable names used in the message explanation are described below.

Variable names	Descriptions
<i>class</i>	Class name
<i>disk</i>	Disk name
<i>group</i>	Group name
<i>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>master</i>	Master object name
<i>proxy</i>	Proxy object name
<i>object</i>	Object name (or physical disk name)
<i>status</i>	Object status
<i>device</i>	Physical disk name (sdX, emcpowerX) X indicates the device identifier.
<i>pslice</i>	Physical slice name (sdXn, emcpowerXn) X indicates the device identifier. n indicates the slice number.
<i>node</i>	Node identifier or node name.
<i>attribute</i>	Attribute name
<i>value</i>	Attribute value
<i>param</i>	Parameter name
<i>val</i>	Parameter value
<i>size</i>	Number of blocks (in sectors), (decimal number)

Variable names	Descriptions
<i>option</i>	Command options.
<i>usage</i>	Syntax when using command.
<i>letter</i>	Characters
<i>details</i>	Details
<i>errno</i>	System call error number (decimal)
<i>sdxerrno</i>	Internal error number defined by GDS (decimal)
<i>string</i>	Other character strings

Explanation

Messages output by command are shown below in the order of severity. There are four levels of severity.

Level of importance	Descriptions
ERROR	This message is output to indicate the cause when a command terminates abnormally.
WARNING	This message is output as a warning when an abnormal event is detected. It will not always result in an abnormal termination of the command.
INFO	This message is output to record the command operation.
TO FIX	This message is output to show you how to fix the command. Usually it outputs command syntax.

E.4.1 Error Messages (60000-60399)

60000 ERROR: connection timeout

Explanation

Connection failed due to no response from sdxservd daemon.

Resolution

Confirm that sdxservd daemon process was started normally.

60001 ERROR: not privileged

Explanation

Executed user is not a superuser.

Resolution

You need to have superuser authority for execution.

60002 ERROR: *option*: illegal option

Explanation

option is illegal.

Resolution

See the Fix Messages described later, or the "Appendix D Command Reference."

60003 ERROR: syntax error

Explanation

Executed command has a syntax error.

Resolution

See the Fix Messages described later, or the "Appendix D Command Reference."

60004 ERROR: *string*: name too long

Explanation

Object name, file name or node identifier specified with *string* is too long.

Resolution

Specify the proper name.

60005 ERROR: *object*: name contains invalid character "*letter*"

Explanation

Object name *object* contains an invalid character *letter*.

Resolution

You can only use alphanumeric characters, "-" (hyphen), and "_" (underscore) when specifying an object name.

60006 ERROR: *object*: name starting with "_" or "-" is invalid

Explanation

Object name *object* starting with "-" (hyphen) or "_" (underscore) is invalid.

Resolution

Specify an object name that starts with an alphanumeric character.

60007 ERROR: device: illegal physical disk name

Explanation

Physical disk name *device* is illegal.

Resolution

See "Appendix D Command Reference."

60008 ERROR: *object*: object names must be unique within a class

Explanation

Object names within a class must be unique.

Resolution

Specify a unique object name.

60010 ERROR: *module*: environment error, *details*

Explanation

The command cannot be executed due to an error in the environment. The error causes only a failure in command execution but has no effect on the system. Volumes can be accessed continuously. *module* is the name of the module on which this event was detected (usually an internal variable name). *details* indicates the error in detail.

Resolution

Collect investigation material and contact your local customer support.

60011 ERROR: *class* is a shadow class

Explanation

The class *class* is a shadow class. The attempted operation is not supported for shadow classes.

Resolution

See "Appendix D Command Reference," and specify a correct command name and class name.

60012 ERROR: *attribute*: invalid attribute name

Explanation

The same *attribute* has already been specified.

Resolution

See "Appendix D Command Reference."

60013 ERROR: *attribute*: attribute name duplicated

Explanation

The same attribute *attribute* has already been specified.

Resolution

You can only specify an attribute once.

60014 ERROR: *attribute=value*: invalid attribute value

Explanation

Attribute value *value* is invalid.

Resolution

See "Appendix D Command Reference."

If "pslice=off: invalid attribute value" is output, the possible cause is that the FJSVsdX-bas package of PRIMECLUSTER GDS may not be installed normally. If this is the cause, reinstall FJSVsdX-bas.

60015 ERROR: *node*: node name duplicated

Explanation

The same node identifier has already been specified.

Resolution

You can only specify a node identifier once.

60016 ERROR: too many nodes in scope

Explanation

There are too many nodes specified in scope.

Resolution

Up to 4 nodes can be specified in a scope.

60017 ERROR: *class*: cannot operate shared objects, cluster control facility not ready

Explanation

Cannot operate shared object since cluster control is not operating.

Resolution

Start cluster control and try again.

60018 ERROR: *node*: remote node cannot be specified for local class

Explanation

A remote node was specified for a local class.

Resolution

Specify the node identifier of the current node for a local class.

60019 ERROR: *node*: remote node cannot be specified for root class

Explanation

A remote node was specified for a root class.

Resolution

Specify the node identifier of the current node for a local class.

60020 ERROR: current node must be specified as scope

Explanation

Current node is not specified for shared class scope.

Resolution

Always specify a node set which includes current node as scope.

60021 ERROR: multi-nodes must be specified for shared class

Explanation

Only one node is specified for shared class scope.

Resolution

Multiple node identifiers must be specified in a scope.

60022 ERROR: *node*: unknown node

Explanation

Node *node* does not exist.

Resolution

First check the cluster system environment. Then, try again after changing the node identifier.

60023 ERROR: *class*: not shared by the current node

Explanation

Scope of *class* does not include the current node.
In the cluster system, *class* created in another node is specified.

Resolution

Execute command on a node sharing *class*. Or, change the specified class name.

60024 ERROR: *node*: current node not active in *class*

Explanation

Current node *node* is a standby node in shared class *class*.

Resolution

Execute the command from another active node.

60025 ERROR: too many disks in *class*

Explanation

There are too many disks registered with *class*.

Resolution

Create a new class.

60026 ERROR: *device*: not connected to *node*

Explanation

The physical disk *device* has not been connected to the node *node*, or *device* has not been registered with the cluster resource database.

Resolution

Check the system configuration and the resource configuration. If *device* has not been connected to *node*, specify a correct physical device name and re-execute the command. If *device* has not been registered with the resource database, perform resource registration to register *device* with the resource database, and re-execute the command.

60027 ERROR: *device*: not a shared device

Explanation

Physical disk *device* is not a shared disk unit, or cluster shared disk definition may not be correct.

Resolution

Check the system configuration and execute the command.

60028 ERROR: *device*: already exists in *class*

Explanation

Physical disk *device* is already registered with *class*.

Resolution

Same physical disk cannot be registered with more than one class. Specify a correct physical disk.

60029 ERROR: *device*: already exists in another class

Explanation

Physical disk *device* is already registered with another class. This class is not shared by current node.

Resolution

Same physical disk cannot be registered with more than one class. Specify a correct physical disk.

60030 ERROR: *device*: no such device

Explanation

There is no physical disk *device*.

Resolution

Specify a correct physical disk name.

60031 ERROR: *object*: physical disk *device* not found

Explanation

The physical disk *device* that composes the object *object* was not found.

Resolution

The physical disk *device* may be connected incorrectly or a disk failure may have occurred. Fix *device* connection or recover from disk failure.
Alternatively, specify an object other than *object*.

60032 ERROR: physical disk *device* not found

Explanation

The physical disk *device* was not found.

Resolution

The physical disk *device* may be connected incorrectly or a disk failure may have occurred. Fix *device* connection or recover from disk failure.

60033 ERROR: *device*: cannot open, errno=*errno*

Explanation

Cannot open physical disk *device*.

Resolution

Confirm that physical disk *device* is operating normally.

60034 ERROR: *device*: not a hard disk

Explanation

Physical disk *device* is not a hard disk.

Resolution

GDS cannot manage devices other than hard disks.

60035 ERROR: *device*: disk driver *driver* not supported

Explanation

Physical disk unit with the driver name *driver* is not supported.

Resolution

There is no resolution.

60036 ERROR: *device*: illegal format

Explanation

Physical disk *device* format is illegal.

Resolution

Check the format status.

60037 ERROR: *object*: device busy

Explanation

Object *object* is in use.

Resolution

Change to unused status and execute the command again.

60038 ERROR: *object*: linked to a cluster service

Explanation

Object *object* is used in a cluster application.

Resolution

Check the cluster environment settings.

60039 ERROR: *device*: configuration information exists in private slice

Explanation

Configuration information exists in the private slice of physical disk *device*, and registering *device* failed. The following events a) to d) are the possible causes.

- a) *device* is already registered with another class.
- b) *device* is registered with a class in another domain.
- c) The entire SDX disk was copied or is being copied to *device* with the disk unit's copy function.
- d) After registering *device* with the class, it was not removed normally, and the private slice and configuration information remain existent in *device* illegally.

Resolution

Check the system configuration and so on, and identify the cause among a) to d). If the cause is a) or b), specify a physical disk not registered with a class. If it is c) or d), collect investigation material and contact your local customer support.

60040 ERROR: *device*: type cannot be specified except undef

Explanation

The class scope includes a node to which physical disk *device* is unconnected and the type can only be set to "undef" for registering *device* to the class.

Resolution

See "Appendix D Command Reference."

60041 ERROR: *class*: *device*: type cannot be specified except undef

Explanation

Class *class* includes a switch group and the type can only be set to "undef" for registering physical disk *device* to *class*.

Resolution

See "Appendix D Command Reference."

60042 ERROR: *device*: write error, *errno*=*errno*

Explanation

A write error occurred in physical disk *device*.

Resolution

A disk failure may have occurred. Identify the cause by referring to disk driver log messages and so on, and recover the disk. If *device* is a write-locked disk that is connected to a switch group, register writable disks first.

60043 ERROR: *device*: disk ID information differs from all disks in specified class

Explanation

The specified class includes no disk with the same disk ID information (class name and disk name) as that stored in the private slice of physical disk *device*. For this reason, *device* cannot be registered with the specified class.

Resolution

Check the system configuration and so on, and specify a physical disk and class name properly.

60044 ERROR: *object*: name already assigned to another object

Explanation

Object name *object* already exists within class. You cannot create multiple objects with the same name.

Resolution

Specify another name and execute the command again.

60045 ERROR: cannot connect to *sdxcld*

Explanation

Connecting to *sdxcld*, a GDS cluster linked daemon was unsuccessful.

Resolution

When unable to identify the cause, collect investigation material and contact your local customer support.

60046 ERROR: physical device driver returned an error, erro=*ermo*

Explanation

Physical disk driver returned an error.

Resolution

Identify the cause by referring to error number and message log.

60047 ERROR: special file operation failed, erro=*ermo*

Explanation

Operating special file was unsuccessful.

Resolution

Identify the cause by referring to error number and GDS log message.

60048 ERROR: sfdsk driver returned an error, erro=*ermo*

Explanation

GDS driver returned an error.

Resolution

Identify the cause by referring to error number, GDS log message, and syslog message.

60049 ERROR: sfdsk driver returned a temporary error, try again for a while

Explanation

GDS driver returned a temporary error.

Resolution

Execute the command again after a while.

60050 ERROR: *class*: class closed down

Explanation

class is closed. You cannot operate objects within a closed class.

Resolution

A number of disk failures may have occurred. Identify the cause by referring to object status, GDS log message, and syslog message. For information on recovery, see "F.1.4 Class Status Abnormality."

60051 ERROR: *class*: class closed down on another node

Explanation

class is closed on another node. You cannot operate objects within a closed class.

Resolution

Recover closed *class*. A number of disk failures may have occurred. Identify the cause by referring to object status, GDS log message, and syslog message. For information on recovery, see "F.1.4 Class Status Abnormality."

60052 ERROR: keep disk cannot be specified for local or shared class

Explanation

Keep disk cannot be registered with local or shared class.

Resolution

Register keep disk with root class.

60053 ERROR: too many keep disks specified

Explanation

The number of specified keep disks exceeds that of undefined disks.

Resolution

When registering multiple keep disks, specify the same number or more as undefined disks at the same time.

60054 ERROR: *class*: already root class exists

Explanation

You tried to create a new class when there is already a root *class*. You can only create one root class within a node.

Resolution

Specify the existing root *class* or a local class.

60056 ERROR: *device* contains overlapping slices

Explanation

The physical disk *device* contains physical slices with overlapping cylinders.

Resolution

Correct the physical slice configuration using the parted(8) command.

60058 ERROR: *device* : one or more free slice numbers are required, delete a slice by parted(8) command

Explanation

The physical disk *device* has no free slice number (number not assigned to a slice).

Resolution

Change the physical slice configuration using the parted(8) command.

60060 ERROR: *device* : no enough unassigned disk space, reserve enough space by parted(8) command

Explanation

The physical disk *device* does not have sufficient free disk space.

Resolution

Use the parted(8) command and create sufficient free disk space. For the size, see "A.2.5 Disk Size."

60061 ERROR: *device* : no unassigned disk space nor swap space, reserve enough space by parted(8) command

Explanation

The physical disk *device* does not have free disk space or swap area.

Resolution

Use the parted(8) command and create sufficient free disk space or swap area. For the size, see "A.2.5 Disk Size."

60062 ERROR: *device*: no enough unassigned disk space nor swap space, reserve enough space by parted(8) command

Explanation

Physical disk *device* does not have sufficient free disk space or swap area.

Resolution

Use the parted(8) command and create sufficient free disk space or swap area. For the size see "A.2.5 Disk Size."

60064 ERROR: *device*: invalid physical device, managed by *driver*

Explanation

Specified physical disk is managed by *driver*, and is therefore invalid.

Resolution

Confirm the I/O and cluster system configuration, and specify a correct physical disk name.

60066 ERROR: *device*: IDE disk cannot be specified as spare disk

Explanation

The *device* is an IDE disk. It cannot be used as a spare disk.

Resolution

For spare disks, use disks other than IDE disks.

60067 ERROR: *class*: no such class

Explanation

Cannot find *class*.

Resolution

Check GDS configuration.

60068 ERROR: *group*: not a group

Explanation

group is not a group name.

Resolution

There is another object with the name *group* within class. Check configuration.

60069 ERROR: *group* is a lower level stripe group

Explanation

Group *group* is a stripe group connected to another group. Disks and groups cannot be connected to or disconnected from *group*.

Resolution

Disconnect *group* from the higher level group as necessary.

60070 ERROR: *group*: connected to a lower level stripe group

Explanation

Group *group* is connected to a lower level stripe group. Disks and groups cannot be connected to or disconnected from *group*.

Resolution

Disconnect the higher level stripe group of *group* from its higher level group as necessary.

60071 ERROR: too many groups in *class*

Explanation

Class *class* already has the maximum number of groups possible.
A maximum of 100 groups can be created within, and a maximum of 1024 groups can be created within a local class or a shared class.

Resolution

Create a new class.

60073 ERROR: too many disks and/or groups are connected to *group*

Explanation

The maximum number of disks or lower level groups are already connected to group *group*.

Resolution

There is no resolution.

60074 ERROR: *object*: smaller than stripe width of group *group*

Explanation

Since the available size of disk or lower level group indicated by *object* is smaller than the stripe width of group *group*, *object* cannot be connected to *group*.

Resolution

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.

60075 ERROR: *class*: three or more nodes exist in class scope

Explanation

The scope of class *class* includes 3 or more nodes. Classes whose scope includes 3 or more nodes cannot include switch groups.

Resolution

Specify a shared class of which scope is 2 nodes.

60076 ERROR: *disk*: *type* disk exists in *class*

Explanation

Class *class* includes disk *disk* of the *type* type. The attempted operation is not supported for classes that include *type* type disks.

Resolution

Delete *disk* from *class* as needed and re-execute the command.

60077 ERROR: *class* includes a group that cannot exist together with a switch group

Explanation

The class *class* includes one of the following groups and a switch group cannot be created.

- Mirror group
- Stripe group
- Concatenation group to which any lower level switch group is not connected

Resolution

Specify a correct class name.

60078 ERROR: active disk must be specified

Explanation

When creating a switch group, specify the active disk in the `-a actdisk=disk` option.

Resolution

Specify the option properly, referring to "Appendix D Command Reference."

60079 ERROR: *disk*: active disk not specified by `-d` option

Explanation

Active disk *disk* was not specified in the `-d` option.

Resolution

When creating a switch group, specify one of the disks specified in the `-d` option of the `sdxdisk -C` command as the active disk.

60080 ERROR: too many disks specified

Explanation

Too many disks specified.

Resolution

Specify a proper number of disks. For the number of disks that can be connected to a group, see "A.1.3 Number of Disks."

60081 ERROR: *disk*: physical scope is not included in class scope

Explanation

The physical scope of disk *disk* is not included in the class scope and *disk* cannot be connected to a switch group.

Resolution

Specify a correct disk name. For the physical scope of a disk, use the `sdxinfo -D` command and check the displayed DEVCONNECT field. For the class scope, use the `sdxinfo -C` command and check the displayed SCOPE field.

60082 ERROR: *disk*: physical scope must be same as class scope

Explanation

The physical scope of the active disk matches the class scope and therefore the physical scope of an inactive disk must match the class scope. The physical scope of disk *disk* does not match the class scope and *disk* cannot be connected to the switch group as the inactive disk.

Resolution

Specify a correct disk name. For the physical scope of a disk, use the `sdxinfo -D` command and check the displayed DEVCONNECT field. For the class scope, use the `sdxinfo -C` command and check the displayed SCOPE field.

60083 ERROR: *disk*: physical scope must include only *node*

Explanation

The physical scope of inactive disk *disk* can include only node *node*.

Resolution

Specify a correct disk name, referring to "Appendix D Command Reference."

60084 ERROR: *disk*: class scope is not included in physical scope

Explanation

The class scope is not included in the physical scope of disk *disk* and *disk* cannot be connected to a group other than a switch group.

Resolution

Specify a correct disk name. For the physical scope of a disk, use the `sdxinfo -D` command and check the displayed DEVCONNECT field. For the class scope, use the `sdxinfo -C` command and check the displayed SCOPE field.

60085 ERROR: *disk*: no such disk

Explanation

There is no *disk*.

Resolution

Check GDS configuration.

60086 ERROR: *object*: already connected to *group*

Explanation

Disk or group indicated by *object*, is already connected to group *group*.

Resolution

Specify a correct disk name or group name.

60087 ERROR: *disk* is a spare disk

Explanation

disk is a spare disk. You cannot connect a spare disk to a group.

Resolution

Change the disk attributes to undefined, and execute the command again.

60088 ERROR: *object* not in *status* status

Explanation

Object status is not *status* status.

Resolution

Confirm that object status is *status*, and execute the command again.

60089 ERROR: *object* too small

Explanation

The size of *object* is too small.

Resolution

Check the necessary object size and specify a larger object.

60090 ERROR: another disk must be connected to *group*

Explanation

You must connect another disk to *group*.

Resolution

Connect another disk and execute the command again.

60091 ERROR: invalid physical slice number, *pslice_num*

Explanation

The specified physical slice number *pslice_num* is invalid.

Resolution

For the physical slice number *pslice_num*, specify an integer from 1 to 15.

60093 ERROR: *pslice* is private slice

Explanation

Physical slice *pslice* is a private slice.

Resolution

If this message is output at the time of `sdxdisk -M` command execution, check whether the disk to which the physical slice *pslice* belongs is registered with a class in another domain. If it is registered, the disk cannot be registered with a class. If not, to register the disk to a class, it is necessary to delete the physical slice *pslice*, for example, with the `parted(8)` command.

If this message is output at the time of `sdxdisk -C` command execution, re-execute the `sdxdisk -C` command. When re-executing, do not specify the slice number of the physical slice *pslice* for the `-v` option.

60094 ERROR: *pslice*: corresponding volume attributes must be specified

Explanation

Volume attributes corresponding to physical slice *pslice* is not specified.

Resolution

See "Appendix D Command Reference."

60095 ERROR: *object*: invalid size

Explanation

Object size of *object* is invalid.

Resolution

See "Appendix D Command Reference."

60096 ERROR: two or more keep disks cannot be connected to a group

Explanation

More than one keep disk cannot be connected to group.

Resolution

See "Appendix D Command Reference."

60097 ERROR: two or more single disks cannot be connected to a group

Explanation

More than one single disk cannot be connected to a group.

Resolution

See "Appendix D Command Reference."

60098 ERROR: both keep and single disks cannot be connected to a group

Explanation

Both keep and single disks cannot be connected to a group at the same time.

Resolution

See "Appendix D Command Reference."

60099 ERROR: *disk*: keep disk cannot be connected to existing group

Explanation

Keep disk *disk* cannot be connected to an existing group.

Resolution

See "Appendix D Command Reference."

60100 ERROR: *disk*: single disk cannot be connected to existing group

Explanation

Single disk *disk* cannot be connected to the existing group.

Resolution

See "Appendix D Command Reference."

60101 ERROR: two or more IDE disks cannot be connected to a group

Explanation

More than one IDE disk cannot be connected to a group.

Resolution

When mirroring an IDE disk, use it with disks other than IDE disks.

60102 ERROR: two or more disks cannot be connected to a group

Explanation

If the FJSVsdx-bas package of PRIMECLUSTER GDS is not installed, multiple disks cannot be connected to a group.

Resolution

If FJSVsdx-bas has not been installed successfully, reinstall it.

60103 ERROR: *group*: a lower level switch group is connected

Explanation

A switch group has been connected to the group *group* and it is impossible to connect a disk to *group* or *group* to another group.

Resolution

Specify a correct higher level group name.

60104 ERROR: *disk*: not a bootable device

Explanation

Disk *disk* cannot be booted.

Resolution

Confirm the disk configuration. When you cannot specify the cause, collect investigation material and contact your local customer support.

60105 ERROR: too few valid configuration database replicas

Explanation

There is not sufficient valid configuration database for class.
This message is displayed when the majority of disks registered with class are unavailable.
Leaving it as is may cause serious problems.

Resolution

For details, see "F.1.4 Class Status Abnormality."

60106 ERROR: *group*: no ENABLE disk in group

Explanation

Group *group* does not have a disk in ENABLE status connected.

Resolution

See "F.1.2 Disk Status Abnormality" and recover disks connected to the group.

60107 ERROR: *msec*: invalid delay value

Explanation

Delay time *msec* is invalid.

Resolution

See "Appendix D Command Reference."

60108 ERROR: *lgroup*: mirror group cannot be connected to another group

Explanation

Since group *lgroup* is a mirror group, it cannot be connected to another group.

Resolution

Specify a correct lower level group name.

60110 ERROR: *lgroup*: same type as higher level group *hgroup*

Explanation

Since the type attributes for group *lgroup* and higher level group *hgroup* are the same, *lgroup* cannot be connected to *hgroup*.

Resolution

Specify a correct group name.

60111 ERROR: *hgroup*: same name as a lower level group

Explanation

For the higher level group name, name *hgroup*, which is the same as the lower level group was indicated.

Resolution

Different names must be indicated for higher level group and lower level group.

60112 ERROR: *hgroup*: any group cannot be connected to a switch group

Explanation

Group *hgroup* is a switch group and any group cannot be connected to.

Resolution

Specify a correct group name.

60114 ERROR: *group*: is a lower level concatenation group

Explanation

The group *group* is a concatenation group connected to another group. Any group cannot be connected to *group*.

Resolution

Disconnect *group* from the higher level group as needed and re-execute the command.

60115 ERROR: *lgroup*: stripe group cannot be connected to concatenation group

Explanation

The group *lgroup* is a stripe group and cannot be connected to a concatenation group.

Resolution

Specify a correct group name.

60116 ERROR: *lgroup*: switch group cannot be connected to mirror group

Explanation

The group *lgroup* is a switch group and cannot be connected to a mirror group.

Resolution

Specify a correct group name.

60117 ERROR: *lgroup*: switch group cannot be connected to stripe group

Explanation

The group *lgroup* is a switch group and cannot be connected to a stripe group.

Resolution

Specify a correct group name.

60118 ERROR: *hgroup*: disk is connected

Explanation

A disk is already connected to concatenation group *hgroup*, and a switch group cannot be connected.

Resolution

Specify a correct group name.

60119 ERROR: *size*: must be more than zero

Explanation

size must be a positive integer.

Resolution

Specify the correct size.

60120 ERROR: *size*: invalid size

Explanation

The indicated size *size* is invalid.

Resolution

See "Appendix D Command Reference."

60121 ERROR: *group*: no such group

Explanation

There is no *group*.

Resolution

Specify a correct group name.

60122 ERROR: *group*: not the highest level group

Explanation

Group *group* is not the highest level group.

Resolution

Disconnect *group* from the higher level group as necessary.

60123 ERROR: too many volumes exist in *object*

Explanation

The maximum number of volumes, already exists in the *object*, that is a class, a group, or a single disk.

Resolution

Create volumes in another group or on a single disk as needed.

60124 ERROR: too many volumes with physical slices in *object*

Explanation

In the *object*, the number of volumes which consist of physical slices has reached the maximum. *object* is a name of group or single disk.

Resolution

Change the physical slice attributes for volumes in *object* from on to off using the `sdxattr -V` command as needed, and execute the command again.

60125 ERROR: *group*: volume with physical slice(s) cannot be created

Explanation

Although a physical slice cannot be created in group *group*, an attempt was made to create a volume with one or more physical slices.

Resolution

Create a volume without any physical slice for *group* either by:

- Turning "off" the [Physical Slice] option in the Volume Configuration dialog box when creating the volume for *group* in GDS Management View
- Using the `-a pslice=off` option when creating the volume for *group* with the `sdxvolume -M` command

If *group* is a mirror group, connect one or more disks to the *group* in order to create a volume with a physical slice.

60126 ERROR: *object*: no enough space

Explanation

There is not enough space in the group or single disk indicated by *object*.

Resolution

Change the size you specify as needed.

60127 ERROR: *disk*: not a single disk

Explanation

The disk *disk* is not a single disk.

Resolution

Specify a correct disk name.

60128 ERROR: *disk: status* disk connected to *group*

Explanation

disk in *status* status is connected to *group*.

Resolution

Restore the status of *disk* and execute the command again.

60129 ERROR: *status* disk exists in *group*

Explanation

The *disk* in *status* status is connected to *group*, or it connected to the lower level group of *group*.

Resolution

Confirm the disk status, and please execute a command again after canceling the *status* status if necessary.

60130 ERROR: *disk: no such* disk in *group*

Explanation

disk is not connected to *group*.

Resolution

Specify a correct disk name and group name.

60131 ERROR: *disk: device: device* busy

Explanation

Disk *disk* (physical disk name is *device*) is in use.

Resolution

Change to unused status, and execute the command again.

60132 ERROR: too many nodes specified

Explanation

Too many nodes were specified.

Resolution

Specify the nodes included in the scope of the class.

60133 ERROR: *option: cannot be specified* for root nor local class

Explanation

Command option *option* cannot be specified for a root nor local class.

Resolution

See "Appendix D Command Reference."

60134 ERROR: *node*: not in scope, class=*class*

Explanation

Node *node* is not included in the scope of the class *class*.

Resolution

Confirm your GDS configuration, specify a node included in the scope of the class, and try the command again.

60135 ERROR: *volume*: cannot start, class closed down, node=*node*

Explanation

Starting volume *volume* was unsuccessful due to the class to which *volume* 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."

60136 ERROR: *volume*: no such volume

Explanation

There is no *volume*.

Resolution

Specify a correct volume name.

60137 ERROR: *object* in *status* status

Explanation

object is in *status* status.

Resolution

Confirm the object status, and cancel the status if necessary.

60138 ERROR: some ACTIVE volumes exist in *class*, node=*node*

Explanation

Volumes of class *class* are already active on node *node*, and volumes of *class* cannot be started or created on this node.

Resolution

Create volumes of class *class* on node *node*. Alternatively, as needed, stop all the volumes within *class* on *node* and then start or create volumes on this node. If a cluster application to which *class* resources are registered is operating or is standby on *node*, before stopping volumes on *node*, stop the cluster application.

60139 ERROR: *volume*: some ACTIVE volumes exist in *class*, node=*node*

Explanation

Volumes of class *class* are already active on node *node*, and volume *volume* cannot be started on this node.

Resolution

Stop all the volumes within class *class* on node *node* as needed, start volume *volume* on this node. If a cluster application to which *class* resources are registered is operating or is standby on *node*, before stopping volumes on *node*, stop the cluster application.

60140 ERROR: *volume*: active disk not connected to *node*

Explanation

The active disk of a switch group including switch volume *volume* is not connected to node *node* and the *volume* cannot be started on the *node*.

Resolution

If an inactive disk is connected to a switch group including *volume*, use the `sdxattr -G` command and switch the inactive disk to the active disk in order to start *volume* on *node*.

60141 ERROR: *volume*: active disk of lower level group *group* is not connected to *node*

Explanation

The active disk of lower level switch group *group* is not connected to node *node*, and volume *volume* cannot be started on *node*. *group* is a lower level switch group that is connected to the highest level group to which *volume* belongs.

Resolution

If an inactive disk is connected to *group*, switch the active disk to the inactive disk with the `sdxattr -G` command in order to enable *volume* startup on *node*.

60142 ERROR: lock is set on volume *volume*, node=*node*

Explanation

The "Lock volume" mode for volume *volume* on node *node* is turned on.

Resolution

Turn off the "Lock volume" mode, or use the `-e unlock` option as necessary.

60143 ERROR: *volume*: cannot stop, class closed down, node=*node*

Explanation

Stopping volume *volume* was unsuccessful due to the class to which *volume* belongs being 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."

60144 ERROR: *object.volume*: status slice exists in *object*

Explanation

Slice *object.volume* in *status* status exists in *object*.

Resolution

Confirm the object status, and cancel the status if necessary.

60145 ERROR: *object* in *status* status, node=*node*

Explanation

object on *node* is in *status* status.

Resolution

Confirm the object status, and cancel the status if necessary.

60146 ERROR: *volume*: stripe type volume cannot be resized

Explanation

volume cannot be resized since it is a volume created in a stripe group.

Resolution

See "Stripe Type Volume and Concatenation Type Volume Expansion" in "A.2.13 Online Volume Expansion."

60147 ERROR: *volume*: concat type volume cannot be resized

Explanation

volume cannot be resized since it is a volume created in a concatenation group.

Resolution

See "Stripe Type Volume and Concatenation Type Volume Expansion" in "A.2.13 Online Volume Expansion."

60148 ERROR: *volume*: consists of multiple mirror slices

Explanation

The volume *volume* cannot be resized since it is a mirror volume that consists of multiple mirror slices.

Resolution

Disconnect disks and lower level groups from the mirror group so that *volume* is composed of only one mirror slice, and try this command again.

60149 ERROR: spare disk connected for *disk*

Explanation

Spare disk is connected instead of *disk*.

Resolution

First recover the *disk* status.

60150 ERROR: *object.volume* is only valid slice

Explanation

object.volume is the only valid slice within volume *volume*. You cannot continue your operation because data within *volume* may be lost if you continue your operation.

Resolution

If *volume* is a mirror volume, you will be able to continue your operation by recovering the mirroring (for example, by connecting a new disk to group). If *volume* is not a mirror volume, remove *volume* as necessary.

60151 ERROR: *object*: the last disk or group in lower level group *group*

Explanation

Object *object* is the only disk or group connected to the lower level group *group*. You cannot disconnect *object* from *group*.

Resolution

Disconnect *group* from the higher level group as necessary.

60152 ERROR: *object*: not connected to the end of concatenation group *group*

Explanation

The object *object* is not the disk or switch group that was last connected to the concatenation group *group*. It is impossible to disconnect *object* from *group*.

Resolution

To disconnect disks and groups from *group*, perform disconnection in inverse order of connection. For the order of connecting disks and groups to *group*, use the `dxinfo -G` command and check the displayed DISKS field.

60153 ERROR: *object*: disk space is assigned to volume *volume*

Explanation

Disk space of disk or lower switch group *object* is assigned to volume *volume*, and *object* cannot be disconnected from the concatenation group.

Resolution

If needed, remove *volume* first.

60154 ERROR: *group*: inactive disk is connected

Explanation

Group *group* has a connected inactive disk and physical disk swapping or disconnection cannot be performed for the active disk of the *group*.

Resolution

According to need, switch the active disk with the `sdxattr -G` command and perform physical disk swapping or disconnection for the previous active disk.

60155 ERROR: *disk*: not inactive disk

Explanation

Disk *disk* is the active disk of a lower level switch group, and physical disk swapping cannot be performed.

Resolution

As needed, switch the active disk with the `sdxattr -G` command, and perform physical disk swapping on the previous active disk.

60156 ERROR: *igroup*: no such group in *hgroup*

Explanation

Group *igroup* is not connected to group *hgroup*.

Resolution

Specify a correct group name.

60157 ERROR: one or more volumes exist in *group*

Explanation

One or more volumes exist in group *group*.

Resolution

First, remove the volume if necessary.

60158 ERROR: *disk* connected to *group*

Explanation

disk is connected to *group*.

Resolution

First, disconnect the disk from group if necessary.

60159 ERROR: *disk*: The last ENABLE disk in class cannot be removed

Explanation

When there is a disk in SWAP or DISABLE status in class, you cannot remove the last disk in ENABLE status.

Resolution

First recover the disk in SWAP or DISABLE status. Or, you can register a new disk with class.

60160 ERROR: *disk*: cannot be removed to avoid class closing down

Explanation

Disk *disk* stores the configuration database and removal of *disk* will result in class closure. For this reason, the disk cannot be removed.

Resolution

Recover or remove a failed disk in the class that includes *disk* and then remove *disk*.

60161 ERROR: one or more volumes exist in *disk*

Explanation

Volume exists in *disk*.

Resolution

First, remove the volume as needed.

60162 ERROR: one or more groups exist in *class*

Explanation

Group exists in *class*.

Resolution

First, remove the group as needed.

60163 ERROR: *disk*: *status* disk exists in *class*

Explanation

disk in *status* status exists in *class*.

Resolution

First, restore the *disk* in *status* status.

60164 ERROR: *volume: status* volume exists in *class*, *node=node*

Explanation

volume in *status* status on *node* exists in *class*.

Resolution

First, change the volume status as needed.

60165 ERROR: *disk: no such disk*

Explanation

There is no disk *disk*.

Resolution

Specify a correct disk name.

60166 ERROR: *volume: not associated with object*

Explanation

You cannot specify slice by combining disk or group indicated by *object*, and volume indicated by *volume*.

Resolution

See "Appendix D Command Reference" and specify a correct disk or group name and volume name.

60167 ERROR: *volume is single volume*

Explanation

Volume is a single volume. Slices of single volumes cannot be detached.

Resolution

Specify a correct volume name.

60168 ERROR: *volume: not a mirror volume*

Explanation

Volume *volume* is not a mirror volume.

Resolution

Specify a correct volume name.

60169 ERROR: *object: not connected to the highest level group*

Explanation

Disk or group indicated by *object* is not connected to the highest-level group.

Resolution

See "Appendix D Command Reference" and specify a correct disk name or group name.

60170 ERROR: *disk*: The last ENABLE disk in class cannot be swapped

Explanation

You cannot swap the last ENABLE status disk in class.

Resolution

Depending on the configuration, use a different method to avoid such situations. For example, you can register a new disk with class.

60171 ERROR: *disk*: keep disk cannot be swapped out

Explanation

Keep disk *disk* cannot be swapped.

Resolution

Change the type attribute of *disk*, and execute the command again.

60172 ERROR: *disk*: volume in status status

Explanation

There is a volume related to disk in *status* status.

Resolution

First, recover the volume in *status* status.

60173 ERROR: *disk*: the highest level group is not a mirror group

Explanation

The highest level group of disk *disk* is not a mirror group. You cannot swap *disk*.

Resolution

Disconnect *disk* from the group, and execute the command in a situation where it is not connected to any group.

60174 ERROR: *disk*: cannot be swapped to avoid class closing down

Explanation

Disk *disk* stores the configuration database and swapping *disk* will close down the class. Therefore, swapping the disk is prevented.

Resolution

Replace a failed disk included in the class to which *disk* is registered and swap *disk*.

60175 ERROR: *disk*: device: device busy on node *node*

Explanation

Disk *disk* (the physical disk name is *device*) is in use on *node*.

Resolution

Change the status to unused and execute the command again.

60176 ERROR: *disk: device: cannot open, erro=erno*

Explanation

Cannot open physical disk *device*.

Resolution

Confirm that physical disk *device* is operating normally.

60177 ERROR: *disk: device: not a hard disk*

Explanation

Physical disk *device* is not a hard disk.

Resolution

Specify a hard disk.

60178 ERROR: *disk: device: illegal format*

Explanation

Physical disk *device* format is illegal.

Resolution

Check the format of the physical disk.

60179 ERROR: *disk busy - /dev/pslice*

Explanation

Physical disk *pslice* is in use.

Resolution

Change to unused status and execute the command again.

60180 ERROR: *disk: device: linked to a cluster service*

Explanation

Physical disk *device* is used in a cluster application.

Resolution

Check the settings for cluster environment.

60181 ERROR: *disk: device: not enough size*

Explanation

Physical disk size is too small.

Resolution

Specify a physical disk with sufficient size.

60182 ERROR: *disk: device: not connected to node*

Explanation

Physical disk *device* is not connected to *node*. Or, cluster shared disk definition may not be correct.

Resolution

Check the system configuration and execute the command.

60183 ERROR: *disk: device: invalid disk, managed by driver*

Explanation

Specified disk (corresponding physical disk name is *device*) is managed by *driver*, and is therefore invalid.

Resolution

Confirm the I/O and cluster system configuration, and specify a correct physical disk name.

60184 ERROR: *object.volume: status slice exists in group*

Explanation

Slice *object.volume* in *status* status already exists in group *group*.

Resolution

Check the object status and cancel *status* status if necessary.

60185 ERROR: *disk : device : invalid disk on node node, managed by driver*

Explanation

On node *node*, physical disk *device* of disk *disk* has been managed by driver *driver* and cannot be handled.

Resolution

Check I/O configurations and cluster system configurations, and specify a correct physical disk name.

60186 ERROR: *disk : device : cannot open on node node, ermo=ermo*

Explanation

Physical disk *device* cannot be opened on node *node*.

Resolution

Check whether or not the physical disk *device* is operating normally.

60187 ERROR: *disk : device : node : not a hard disk*

Explanation

Physical disk *device* on node *node* is not a hard disk.

Resolution

Specify a hard disk.

60188 ERROR: *disk : device : node : illegal format*

Explanation

The format of physical disk *device* on node *node* is incorrect.

Resolution

Check the format.

60189 ERROR: *disk* busy on node *node* - / dev/*pslice*

Explanation

Physical disk *pslice* on node *node* is in use.

Resolution

Change the status to unused and retry the command.

60190 ERROR: *disk* : *device* : *node* : linked to a cluster service

Explanation

Physical disk *device* on node *node* is being used by cluster applications.

Resolution

Check the cluster environment configurations.

60191 ERROR: *disk* : *device* : *node*: not enough size

Explanation

The size of physical disk *device* on node *node* is too small.

Resolution

Replace it with a physical disk of sufficient size.

60192 ERROR: *disk* : *device* : read error, *ermo*=*ermo*

Explanation

A read error occurred on physical disk *device*.

Resolution

The possible cause is a disk failure. Identify the cause based on disk driver log messages and so on and restore the disk.

60193 ERROR: *disk* : *device* : *node*: read error, *ermo*=*ermo*

Explanation

A read error occurred on physical disk *device* on node *node*.

Resolution

The possible cause is a disk failure. Identify the cause based on disk driver log messages and so on and restore the disk.

60194 ERROR: *object*: read error, *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.

60195 ERROR: *object.volume* : no such slice

Explanation

You cannot specify slice by combining disk or group indicated by *object*, and volume indicated by *volume*.

Resolution

See "Appendix D Command Reference" and specify a correct disk or group name and volume name.

60196 ERROR: *disk*: not active disk

Explanation

Disk *disk* is not the active disk.

Resolution

Specify the active disk.

60197 ERROR: *disk*: I/O error not occur

Explanation

An I/O error has not been occurred on the disk *disk*. There is no need to restore *disk* using the `sdxfix -D` command.

Resolution

There is no resolution.

60198 ERROR: *class*: not close down

Explanation

Class *class* is not closed. Recovery *class* with the `sdxfix -C` command is not required.

Resolution

There is no resolution.

60199 ERROR: no valid configuration database

Explanation

No valid configuration database of the class is found. If all (or most of) the disks that are registered with the class are unavailable, this message is output.

Resolution

See "F.1.4 Class Status Abnormality."

60200 ERROR: *class*: closed down on all nodes in class scope

Explanation

During recovery of class *class* from the closed status, *class* was closed on all the nodes within the class scope.

Resolution

Re-execute the `sdxfix -C` command.

60201 ERROR: *disk*: connected to switch group

Explanation

Disk *disk* is connected to a switch group and an I/O error status of the *disk* cannot be removed.

Resolution

Make *disk* switchable with the `sdxswap -O` command or [Operation]:[Swap Physical Disk] in GDS Management View to remove an I/O error status of *disk*. After this procedure, swap the physical disks of *disk* according to need and then make the *disk* usable again with the `sdxswap -I` command or [Operation]:[Restore Physical Disk] in GDS Management View.

60202 ERROR: *volume*: cannot restart to copy, cancel current interrupted copy operation by `sdxcopy` command with `-C` option

Explanation

Copying process within volume could not be resumed.

Resolution

Cancel the interrupted copying process according to need and resume copying.

60203 ERROR: *attribute=value*: cannot modify type attribute of root class

Explanation

The root class type attribute cannot be changed.

Resolution

Specify a correct class name.

60204 ERROR: *class*: class names must be unique within a domain

Explanation

Class *class* with the same name as *class* already exists within the cluster domain.

Resolution

When renaming a class resulted in this message, specify another class name. When expanding the class scope resulted in this message, see "Important Point 2" in "A.2.24 Changing Over from Single Nodes to a Cluster System [Linux2.6]."

60205 ERROR: *class*: volume minor numbers must be unique within a domain

Explanation

A volume with the same minor number as that of a volume in the class *class* was found within the cluster domain.

Resolution

See "Important Point 2" in "A.2.24 Changing Over from Single Nodes to a Cluster System [Linux2.6]."

60206 ERROR: one or more disks not connected to *node*

Explanation

Disks that are not connected to node *node* or for which disk resources are not created yet exist in the class.

Resolution

Check the hardware configuration and the disk configuration within the class. If disk resources are not created yet, create those through the resource registration. For the resource registration, see "Appendix H Shared Disk Unit Resource Registration."

60207 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.

60208 ERROR: *object*: no such object

Explanation

There is no *object*.

Resolution

Specify a correct object name.

60209 ERROR: *class*: shared objects information not yet available, try again for a while

Explanation

Share object information is not available yet.

Resolution

Wait until cluster control is activated, and execute the command again.

60210 ERROR: *node*: node in stopped status

Explanation

You cannot proceed with operation since *node* is in stopped status.

Resolution

Start the *node* and try again.

60211 ERROR: *node*: node in abnormal status

Explanation

node is in an abnormal status and cannot proceed with operation.

Resolution

Confirm the normal startup of *node* and try again.

60212 ERROR: cluster communication failure

Explanation

Cannot proceed with operation since communication failed with cluster.

Resolution

Check that the cluster system and GDS are operating normally.
After recovering, try again.

60213 ERROR: cluster communication failure, *sdxfcrmo=sdxfcrmo*

Explanation

Cannot proceed with operation since communication failed with cluster.

Resolution

Check that the cluster system and GDS are operating normally.
After recovering, try again.

60214 ERROR: cluster communication failure, *remote-node=node, sdxfcrmo=sdxfcrmo*

Explanation

Cluster communication with remote node *node* failed. Operation cannot be performed.

Resolution

Check that the cluster system and GDS are operating normally.
After recovering, try again.

60215 ERROR: *class*: not a root class

Explanation

Class *class* is not a root class.

The possible causes are as follows.

- a. The specified class name is wrong.
- b. An attempt to use a function that is only for a root class was made for a local class or a shared class.
- c. In a system where the FJSVsdx-bas package of PRIMECLUSTER GDS is not installed normally, an attempt to create a local class group or a shared class group was made.
- d. In a system where the FJSVsdx-bss package of PRIMECLUSTER GDS Snapshot is not installed normally, an attempt to perform proxy operation was made for a local class or a shared class.

Resolution

If the cause is a) or b), see "Appendix D Command Reference" and specify a correct class name. If it is c) or d), install FJSVsdx-bas or FJSVsdx-bss normally.

60216 ERROR: *disk*: not a keep disk

Explanation

Disk *disk* is not a keep disk.

Resolution

See "Appendix D Command Reference" and specify a correct disk name.

60217 ERROR: *disk*: not connected to any group

Explanation

Disk *disk* is not connected to a group.

Resolution

See "Appendix D Command Reference" and specify a correct disk name.

60218 ERROR: *volume: status* volume exists in *group*

Explanation

There is *volume* in *status* status in *group*.

Resolution

Recover volume status, and execute the command again.

60219 ERROR: *disk*: not a system disk

Explanation

Disk *disk* is not a system disk.

Resolution

Specify a correct disk name, and execute the command again.

60221 ERROR: *device*: mandatory system disk must be registered to *class*

Explanation

The system disk *device* is not registered with the class *class*. *device* contains the slice currently operating as / (root), /usr, /var, /boot, or /boot/efi, and to perform system disk setting, it is necessary to register *device* with *class*.

Resolution

Complete preparations properly referring to "5.1.1 System Disk Settings [Itanium]" or "Appendix D Command Reference" and execute the command again.

60223 ERROR: *disk*: mandatory system disk must be specified

Explanation

The system disk *disk* was not specified. *disk* contains the slice currently operating as / (root), /usr, /var, /boot, or /boot/efi, and to cancel disk mirroring, it is necessary to specify *disk*.

Resolution

Specify all system disks with the slices currently operating as / (root), /usr, /var, /boot, or /boot/efi, and execute the command again.

60224 ERROR: *disk*: two or more disks connected to *group*

Explanation

group to which *disk* is connected has two or more disks connected.

Resolution

See "Appendix D Command Reference" and complete preparation correctly. Then execute the command again.

60225 ERROR: root file system not mounted on volume

Explanation

Root file system is not mounted on volume.

Resolution

Confirm the configuration, and see "Appendix D Command Reference."

60226 ERROR: illegal slice name

Explanation

Slice name includes a "." (period).

Resolution

Specify the correct slice name.

60227 ERROR: *disk.volume* cannot be operated on the current node, take over by *sdxslice* command with *-T* option

Explanation

Slice *disk.volume* cannot be operated on the current node.

Resolution

Take over the slice by executing *sdxslice -T* command.

60228 ERROR: *volume*: physical slice attribute value is off

Explanation

The physical slice attribute value of volume *volume* is "off." A slice in a volume without physical slices cannot be detached.

Resolution

Retry the command after turning the physical slice attribute of *volume* to be "on" according to need.

60229 ERROR: *object*: device busy on node *node*

Explanation

object is in use on *node*.

Resolution

Change to unused status and execute the command again.

60230 ERROR: *class*: not a shared class

Explanation

class is not a shared class.

Resolution

Specify a shared class.

60231 ERROR: *param*: invalid parameter name

Explanation

Parameter name *param* is invalid.

Resolution

See "Appendix D Command Reference."

60232 ERROR: *param =val*: invalid parameter value

Explanation

Parameter value *val* is invalid

Resolution

See "Appendix D Command Reference."

60233 ERROR: *param*: parameter name duplicated

Explanation

The same parameter name *param* has already been specified.

Resolution

You can only use a parameter name once.

60234 ERROR: *copy_concurrency=val*: value more than or equal to the number of actually running copy operations must be specified

Explanation

Value smaller than the number of copying currently in process has been specified for *copy_concurrency* parameter.

Resolution

Set a value more than or equal to the number of copying currently in process for *copy_concurrency* parameter.

60235 ERROR: *mode=string*: access mode duplicated

Explanation

Multiple access modes have been specified.

Resolution

Specify only one access mode.

60236 ERROR: *mode=string*: invalid access mode

Explanation

Access mode value *string* is invalid.

Resolution

See "Appendix D Command Reference."

60237 ERROR: *volume*: already started with different access mode, node=*node*

Explanation

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

Stop volume *volume* according to need and execute again.

60238 ERROR: *volume*: related to proxy volume *proxy*

Explanation

Volume *volume* is a master volume related to proxy volume *proxy*.

Resolution

Cancel the relationship of master volume *volume* and proxy volume *proxy* according to need and execute again.

60239 ERROR: *volume*: related to master volume *master*

Explanation

Volume *volume* is a proxy volume related to master volume *master*.

Resolution

Cancel the relationship of master volume *master* and proxy volume *volume* according to need and execute again.

60240 ERROR: *volume*: related to proxy volume *proxy* with EC

Explanation

EC session exists between volume *volume* and proxy volume *proxy*.

Resolution

Cancel the relationship of master volume *volume* and proxy volume *proxy* according to need, and execute again.

60241 ERROR: *volume*: related to master volume *master* with EC

Explanation

EC session exists between volume *volume* and master volume *master*.

Resolution

Use the sdxproxyCancel command and cancel the EC session between master volume *volume* and proxy volume *volume* according to need.
Alternatively, as needed, cancel the relationship of master volume *volume* and proxy volume *volume*, and execute again.

60242 ERROR: *volume*: related to proxy volume *proxy* with TimeFinder

Explanation

There is a BCV pair between the volume *volume* and the proxy volume *proxy*.

Resolution

Break the relation between the master volume *volume* and the proxy volume *proxy* according to need, and try the command again.

60243 ERROR: *volume*: related to master volume *master* with TimeFinder

Explanation

There is a BCV pair between the volume *volume* and the master volume *master*.

Resolution

Break the relation between the master volume *volume* and the proxy volume *proxy* according to need, and try the command again.

60244 ERROR: *volume*: related to proxy volume *proxy* with SRDF

Explanation

There is an SRDF pair between the volume *volume* and the proxy volume *proxy*.

Resolution

Break the relation between the master volume *master* and the proxy volume *volume* according to need, and try the command again.

60245 ERROR: *volume*: related to master volume *master* with SRDF

Explanation

There is an SRDF pair between the volume *volume* and the master volume *master*.

Resolution

Break the relation between the master volume *volume* and the proxy volume *proxy* according to need, and try the command again.

60246 ERROR: *proxy* : no parted proxy volume in proxy group

Explanation

No parted proxy volume exists in proxy group *proxy*.

Resolution

According to need, part the proxy and retry the command.

60248 ERROR: *volume*: parted proxy volume

Explanation

Volume *volume* is a proxy volume parted from master volume.

Resolution

Rejoin the proxy volume to the master volume or break the relationship between them according to need and try this command again.

60249 ERROR: *group*: related to proxy group *proxy*

Explanation

Group *group* is a master group related to proxy group *proxy*.

Resolution

Cancel the relationship between master group *group* and proxy group *proxy* according to need, and execute again.

60250 ERROR: *group*: related to master group *master*

Explanation

Group *group* is a proxy group related to master group *master*.

Resolution

Cancel the relationship between master group *master* and proxy group *group* according to need, and execute again.

60251 ERROR: *volume*: related to master or proxy volume

Explanation

Volume *volume* is related to either master volume or proxy volume.

Resolution

Cancel the relationship between master and proxy according to need, and execute again.

60252 ERROR: *volume*: joined to master volume *master*

Explanation

Volume *volume* is a proxy volume joined to master volume *master*.

Resolution

Part volume *volume* from master volume *master*, or cancel the relationship with master volume according to need, and execute again.

60253 ERROR: *volume*: copying from master volume *master*

Explanation

Data is being copied from master volume *master* to volume *volume*.

Resolution

After the copying process is complete, execute again.

60254 ERROR: *volume*: copying from proxy volume *proxy*

Explanation

Data is being copied from proxy volume *proxy* to volume *volume*.

Resolution

After the copying process is complete, execute again.

60255 ERROR: *class* is root class

Explanation

The class *class* is the root class. The attempted operation is not supported for the root class.

Resolution

Specify a correct class name.

60256 ERROR: *object*: not volume nor group

Explanation

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.

60257 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.

60258 ERROR: *object*: same name as master

Explanation

The object name specified for proxy is the same as the master's name, *object*.

Resolution

Different object names must be specified for master and proxy.

60259 ERROR: *group*: not a mirror group

Explanation

Group *group* is not a mirror group.

Resolution

Specify a correct group name.

60261 ERROR: no volume exists in *group*

Explanation

There are no volumes in group *group*.

Resolution

Use `sdxvolume -M` command to create a volume within group *group*, and execute again.

60262 ERROR: too many proxy volumes are related to *master*

Explanation

There are too many proxy volumes related to master object *master*.

Resolution

Check the GDS configuration. For details about the number of proxy volumes, see "A.1.9 Number of Proxy Volumes."

60263 ERROR: *master*: corresponding proxy volume name must be specified

Explanation

There is no corresponding proxy volume name specified for master volume *master*.

Resolution

See "D.14 sdxproxy – Proxy object operations."

60264 ERROR: *proxy*: proxy volume name duplicated

Explanation

The same proxy volume name *proxy* is specified for more than one master volume.

Resolution

Specify a different name, and execute again.

60265 ERROR: *volume*: no such volume in *group*

Explanation

Volume *volume* does not exist.

Resolution

Specify a correct volume name or group name, and execute again.

60266 ERROR: *object*: object name duplicated

Explanation

The duplicate object name *object* was specified, or the group to which the volume *object* belongs and the volume *object* itself were specified at the same time.

Resolution

You can specify an object only once.

60267 ERROR: *proxy*: already parted

Explanation

Proxy volume *proxy* is already parted.

Resolution

Nothing needs to be done.

60268 ERROR: one point copy not available

Explanation

The OPC (One Point Copy) function is unavailable and proxy operations cannot be performed.

Resolution

See "(1) The Advanced Copy function cannot be used in master-proxy copying." in "F.1.7 Proxy Object Abnormality."

60269 ERROR: *proxy*: already joined

Explanation

Proxy object *proxy* is already joined.

Resolution

Nothing needs to be done.

60270 ERROR: *proxy*: not joined to master

Explanation

Proxy volume *proxy* is not joined to master volume.

Resolution

Join the proxy volume *proxy* to master volume according to need, and execute again.

60271 ERROR: *proxy*: no such proxy object

Explanation

Proxy object *proxy* cannot be found.

Resolution

Specify a correct proxy object name.

60272 ERROR: *master*: no such master object

Explanation

Master object *master* cannot be found.

Resolution

Specify a correct master object name.

60273 ERROR: *volume*: exists in proxy group

Explanation

Volume *volume* is a proxy volume in proxy group.

Resolution

Specify a proxy group and execute again.

60274 ERROR: *group*: not a proxy group

Explanation

group is not a proxy group group name.

Resolution

Check the GDS configuration, and specify a proxy group group name.

60275 ERROR: *volume*: copying with EC

Explanation

Volume *volume* is in the process of EC copy.

Resolution

After the copying process is complete, execute again.

60276 ERROR: *volume*: copying with OPC

Explanation

Volume *volume* is in the process of OPC copy.

Resolution

After the copying process is complete, execute again.

60277 ERROR: *volume*: copying with TimeFinder

Explanation

Volume *volume* is being copied with TimeFinder.

Resolution

Execute the command again after the copying process is complete.

60278 ERROR: *volume*: copying with SRDF

Explanation

Volume *volume* is being copied with SRDF.

Resolution

Execute the command again after the copying process is complete.

60279 ERROR: *volume*: related to same master volume as proxy *proxy*, master=*master*

Explanation

Volume *volume* is related to the same master volume *master* as proxy volume *proxy*.

Resolution

Specify a correct volume name. For details, see "Appendix D Command Reference."

60280 ERROR: master and proxy exist in same group *group*

Explanation

The specified master volume and proxy volume exist in the same group *group*.

Resolution

Specify volumes that exist in different groups.

60281 ERROR: *proxy*: joined to master with EC, rejoin them by soft copy and try again

Explanation

Slices cannot be swapped since EC session in process between proxy *proxy* and master.

Resolution

Cancel the EC session by either of the following operations according to need and try this command again.

- Part the master and the proxy once and rejoin them with the `sdxproxy Rejoin -e softcopy` command.
- Break the relationship between the master and the proxy once and rejoin them with the `sdxproxy Join -e softcopy` command.

60282 ERROR: *proxy*: joined to master with TimeFinder, rejoin them by soft copy and try again

Explanation

Slices cannot be swapped because there is a BCV pair between the proxy *proxy* and the master.

Resolution

Cancel the BCV pair by either of the following operations according to need and try this command again.

- Part the master and the proxy once and rejoin them with the `sdxproxy Rejoin -e softcopy` command.
- Break the relationship between the master and the proxy once and rejoin them with the `sdxproxy Join -e softcopy` command.

60283 ERROR: *proxy*: joined to master with SRDF, rejoin them by soft copy and try again

Explanation

Slices cannot be swapped because there is an SRDF pair between the proxy *proxy* and the master.

Resolution

Cancel the SRDF pair by either of the following operations according to need and try this command again.

- Part the master and the proxy once and rejoin them with the `sdxproxy Rejoin -e softcopy` command.
- Break the relationship between the master and the proxy once and rejoin them with the `sdxproxy Join -e softcopy` command.

60284 ERROR: *volume*: proxy volume cannot be specified when using TimeFinder

Explanation

You cannot specify *volume* when you perform parting, rejoining or restoring because there is a BCV pair between the proxy volume *volume* and the master.

Resolution

If you wish to perform parting, rejoining, or restoring, you must specify the group to which *volume* belongs when executing the command.

60285 ERROR: *volume*: proxy volume cannot be specified when using SRDF

Explanation

You cannot specify *volume* when you perform parting, rejoining or restoring because there is an SRDF pair between the proxy volume *volume* and the master.

Resolution

If you wish to perform parting, rejoining, or restoring, you must specify the group to which *volume* belongs when executing the command.

60286 ERROR: *proxy*: failed to start soft copy

Explanation

An error occurred when synchronization copying with the soft copy function was started between the proxy volume *proxy* and the master volume.

Resolution

Collect investigation material and contact your local customer support.

60288 ERROR: OPC not available

Explanation

The OPC function is unavailable.

Resolution

See "(1) The Advanced Copy function cannot be used in master-proxy copying." in "F.1.7 Proxy Object Abnormality."

60289 ERROR: EC not available

Explanation

The EC function is unavailable.

Resolution

See "(1) The Advanced Copy function cannot be used in master-proxy copying." in "F.1.7 Proxy Object Abnormality."

60290 ERROR: *proxy*: too many EC/OPC sessions

Explanation

The number of EC or OPC sessions within the physical disk (LU) or the disk array body has reached the upper limit of supported concurrent sessions. For this reason, a new EC or OPC session cannot be started.

Resolution

To make copying by EC or OPC available, wait until the running session ends, and try this command again. Alternatively, according to need, cancel the running session using the `sdxproxy Cancel` command, the `sdxproxy Break` command, or `[Operation]:[Proxy Operation]:[Break]` and try this command again.

60291 ERROR: *proxy*: offset is different from master volume *master*

Explanation

Master volume *master* and proxy volume *proxy* have different top block (sector) numbers. The top block number is not a physical block number that indicates the offset on a physical disk, but is a logical block number that indicates the offset within a group (or a single disk) to which the volume belongs. The layout (offsets and sizes) of volumes within a master group and a proxy group must be consistent.

Resolution

For the proxy group, choose a group in which the layout of volumes are consistent with those of the master group. For the top block (sector) numbers and the sizes of volumes, use the `sdxinfo` command and check the `1STBLK` field and the `BLOCKS` field of the displayed volume information.

60292 ERROR: *proxy*: number of volumes is different from master group *master*

Explanation

Master group *master* and proxy group *proxy* include different numbers of volumes. The layout (offsets and sizes) of volumes within a master group and a proxy group must be consistent.

Resolution

For the proxy group, choose a group in which the layout of volumes are consistent with those of the master group.

60293 ERROR: proxy: failed to start OPC, source=disk.volume, target=disk.volume, class=class

Explanation

An error occurred between the proxy volume *proxy* and the master volume when OPC started. "source" indicates the copy source slice name, "target" indicates the copy destination slice name, and "class" indicates the class to which the copy source and the copy destination slices belong.

Resolution

Identify the cause by referring to GDS log messages, disk driver log messages, and Fujitsu Storage Systems ETERNUS log messages that were output right before the error occurrence, and recover the status.

60294 ERROR: proxy: failed to start EC session, source=disk.volume, target=disk.volume, class=class

Explanation

An error occurred between the proxy volume *proxy* and the master volume when the EC session started. "source" indicates the copy source slice name, "target" indicates the copy destination slice name, and "class" indicates the class to which the copy source and the copy destination slices belong.

Resolution

Identify the cause by referring to GDS log messages, disk driver log messages, and Fujitsu Storage Systems ETERNUS log messages that were output right before the error occurrence, and recover the status.

60295 ERROR: copy region exceeds volume region, volume=volume, offset=blkno, size=size

Explanation

The region specified with the start block number *blkno* and the size *size* exceeds the region of the volume *volume*.

Resolution

Check the size of the volume *volume* using the `sdxinfo` command, and modify the copy region file.

60296 ERROR: cannot open copy region file *region_file*, errno=errno

Explanation

Opening the copy region file *region_file* failed.

Resolution

Collect investigation material and contact your local customer support.

60297 ERROR: cannot read copy region file *region_file*, errno=errno

Explanation

Reading the copy region file *region_file* failed.

Resolution

Collect investigation material and contact your local customer support.

60298 ERROR: syntax error in copy region file *region_file*, line=*line*

Explanation

The copy region file *region_file* contains an error at the line indicated by *line*.

Resolution

Check contents of *region_file*.

60299 ERROR: too many copy regions in copy region file *region_file*

Explanation

The copy region file *region_file* specifies too many copy regions.

Resolution

Modify *region_file*.

60300 ERROR: copy region file *region_file* contains overlapping regions, line=*line1*, line=*line2*

Explanation

Copy regions specified in the lines indicated by *line1* and *line2* overlap.

Resolution

Check contents of *region_file*.

60301 ERROR: OPC not available with multiple extents

Explanation

The disk unit may not support the OPC copy function with multiple extents or the function may not currently be available for some reason.

Resolution

Check the disk unit's hardware configuration. According to need, specify one copy region at a time and execute the command for every region.

60303 ERROR: cannot get configuration information, *sdxinfo(1)* command failed

Explanation

The *sdxinfo(1)* command failed, and GDS configuration information could not be obtained.

Resolution

Remove the error after identifying the cause by referring to the message for the *sdxinfo(1)* command that was output right before the occurrence of the error, and try the *sdxproxy Root* command again.

60304 ERROR: *volume*: alternative volume *altvol* already specified for volume *curvol*

Explanation

Another alternative volume *volume* was assigned to the volume *curvol* for which the alternative volume *altvol* has already been specified.

Resolution

You can specify only one alternative volume for a volume.

60305 ERROR: *volume*: corresponding volume not specified in /etc/fstab

Explanation

The volume *volume* is not the alternative volume of a volume described in /etc/fstab as a file system or a swap area.

Resolution

Specify a volume that fits one of the following requirements.

- Proxy volume of a master volume described in the /etc/fstab file
- Master volume of a proxy volume described in the /etc/fstab file
- Another proxy volume related to the master volume of a proxy volume described in the /etc/fstab file

60306 ERROR: *volume*: read only volume

Explanation

The access mode of the volume *volume* is ro (read only). Since a read-only volume cannot be mounted, this cannot be configured for an alternative boot environment.

Resolution

Change the access mode attribute of the volume *volume* to rw (read and write) using the `sdxattr -V` command, and try this command again.

60307 ERROR: alternative root volume must be specified

Explanation

The alternative root volume has not been specified.

Resolution

You must specify the alternative root volume.

60308 ERROR: *volume*: file system cannot be checked or repaired, fsck(8) command failed, exit-status=*exitstat*

Explanation

The fsck(8) command for the volume *volume* failed, and a file system on *volume* could not be inspected or repaired. The exit status of the fsck(8) command is *exitstat*. Inconsistency may have arisen in the file system on *volume*.

Resolution

Remove the error after identifying the cause by referring to the message for the fsck(8) command that was output right before the occurrence of the error and to the manual for the fsck(8) command. If required, restore the data on *volume*, for example, using the backed up data, and try the `sdxproxy Root` command again.

60309 ERROR: *volume*: file system cannot be mounted, mount(8) command failed, exit-status=*exitstat* *details*

Explanation

The mount(8) command for the volume *volume* failed, and the file system on *volume* could not be mounted. The exit status of the mount(8) command is *exitstat*. *details* is the error message for the mount(8) command.

Resolution

Remove the error after identifying the cause based on *details*, and try the `sdxproxy Root` command again.

**60311 ERROR: *volume*: cannot change boot disks, efibootmgr(8) command failed,
exit-status=*exitstat*
*details***

Explanation

The efibootmgr(8) command failed and the boot manager setting to boot from the volume *volume* was not set. The exit status of the efibootmgr(8) command is *exitstat*. *details* is the error message for the efibootmgr(8) command.

Resolution

Remove the error after identifying the cause based on *details*, and try the sdxproxy Root command again.

**60312 ERROR: *volume*: file system cannot be unmounted, umount(8) command failed,
exit-status=*exitstat*
*details***

Explanation

The umount(8) command for the volume *volume* failed, and the file system on *volume* temporarily mounted could not be unmounted. The exit status of the umount(8) command is *exitstat*. *details* is the error message for the umount(8) command.

Resolution

Remove the error after identifying the cause based on *details*, unmount *volume*, and try the sdxproxy Root command again.

30313 ERROR: *class*: not a shadow class

Explanation

The class *class* is not a shadow class.

Resolution

See "Appendix D Command Reference," and specify a proper command name and class name.

60314 ERROR: *device*: no configuration information

Explanation

Since no configuration information resides in the private slice on the physical disk *device* or no private slice exists on *device*, *device* cannot be registered with a shadow class. The possible causes are as follows.

- a) *device* is not registered with a class in another domain yet, and the private slice has not been copied from the SDX disk to *device* with the disk unit's copy function. In addition, the disk is not a disk previously removed from a class with the sdxconfig Remove -e keepid command.
- b) *device* is already registered with a class in another domain, but has not been enabled.
- c) The private slice has been copied from the SDX disk to *device* with the disk unit's copy function, but the copy source SDX disk has not been enabled.

Resolution

Check on the system configuration and so on, and identify which cause among a), b) and c) applies. If the cause is a), see "Appendix D Command Reference" and specify a correct command name and physical disk name. If it is b), restore device in another domain. If it is c), restore the copy source SDX disk.

60315 ERROR: *device*: registered with illegal class in another domain

Explanation

While the physical disk *device* is already registered with a different class in another domain, registering it with the same shadow class was attempted.

Resolution

You should register a disk registered with the same class in another domain with one shadow class.

60316 ERROR: *device*: *disk*: not same as disk name *diskname* in another domain

Explanation

The physical disk *device* is already registered as *diskname* with a class in another domain. It cannot be registered with a shadow class as *disk* that is another disk name.

Resolution

Specify the same disk name *diskname* as in another domain.

60317 ERROR: *device*: private slice size not same as another disk in *class*

Explanation

Since the physical disk *device* has the private slice that is unequal in size compared to the private slices on other disks registered with the shadow class *class*, it cannot be registered with *class*.

Resolution

Check on the system configuration and so on, and specify a correct disk name and shadow class name.

60318 ERROR: device is bound to RAW device.

Explanation

The device to be registered has been bound to a RAW device.
With GDS, you cannot register disks bound to a RAW device.

Resolution

Register a disk not bound to a RAW device. Alternatively, cancel the bind to a RAW device using the `raw(8)` command and then register.

60319 ERROR: no license

Explanation

The command cannot be executed. The possible causes are as follows.

- a. In a system where the FJSVsdx-cmd package of PRIMECLUSTER GDS is not installed normally, the `sdxconfig` command was used.
- b. In a system where the FJSVsdx-ss package of PRIMECLUSTER GDS Snapshot is not installed normally, the `sdxshadowdisk` command was used.

Resolution

To use the `sdxconfig` command, install FJSVsdx-cmd normally. To use the `sdxshadowdisk` command, install FJSVsdx-ss normally.

60320 ERROR: output file already exists

Explanation

The specified output file is an existing file.

Resolution

Specify a nonexistent file name. To overwrite an existing file, use the `-e` update option.

60321 ERROR: failed to create configuration file

Explanation

Configuration file creation failed.

Resolution

Check whether the specified path to the configuration file is correct.

60322 ERROR: *class*: failed to get configuration information

Explanation

Acquisition of *class* configuration information failed.

Resolution

Collect investigation material and contact your local customer support.

60323 ERROR: *proxy*: proxy volume exists in *class*

Explanation

Proxy volume *proxy* exists in *class*.

Resolution

Break the relationship of proxy volume *proxy* to the master as needed.

60324 ERROR: *group*: switch group exists in *class*

Explanation

Class *class* includes switch group *group*.

The attempted operation is not supported for classes that include switch groups.

Resolution

Delete the *group* as needed and re-execute the command.

60325 ERROR: failed to output configuration table

Explanation

Sending the configuration table to standard output or to a configuration file failed.

Resolution

Collect investigation material and contact your local customer support.

60326 ERROR: input file not found

Explanation

The specified input file does not exist.

Resolution

Specify a correct file name.

60327 ERROR: *class* : not same as class name *name* in configuration table

Explanation

The specified class name *class* is different from the class name *name* in the configuration table.

Resolution

Specify a correct class name, or change the class name in the configuration table to *class* with the `sdxcfg Convert` command.

60328 ERROR: *disk* : no such disk in configuration table

Explanation

The configuration table does not contain disk *disk*.

Resolution

Specify a correct disk name.

60329 ERROR: *device* : no such device in configuration table

Explanation

The configuration table does not contain physical disk *device*.

Resolution

Specify a correct physical disk name.

60330 ERROR: *group* : no such group in configuration table

Explanation

The configuration table does not contain group *group*.

Resolution

Specify a correct group name.

60331 ERROR: *object* : exists in *type* group in configuration table

Explanation

In the configuration table *object* that is a disk or a lower level group is connected to a *type* type group. Disks connected to concat or stripe type groups and lower level groups connected to stripe type groups cannot be removed from configuration tables.

Resolution

Specify a correct disk name or group name.

60332 ERROR: at least one object must remain in configuration table

Explanation

Removing the specified objects from the configuration table will result in no object in the configuration table. A configuration table needs to contain a minimum of one object.

Resolution

Specify correct object names.

60333 ERROR: *object.volume* is only valid slice in configuration table

Explanation

In the configuration table slice *object.volume* is an only valid slice that comprises mirror volume *volume*. For this reason, *object* that is a disk or a lower level group cannot be removed from the configuration table.

Resolution

Specify a correct disk name or group name.

60334 ERROR: *class* : not a local class

Explanation

Class *class* is not a local class.

Resolution

Specify a local class.

60335 ERROR: file name too long

Explanation

The specified file name contains too many characters.

Resolution

Specify a correct file name.

60336 ERROR: failed to open input file, *errno=errno*

Explanation

Input file open failed.

Resolution

Identify the cause based on the error number *errno*.

60337 ERROR: configuration table corrupted, *sdxfunc=sdxfunc*, *sdxerrno=sdxerrno*

Explanation

Contents of the configuration table are invalid.

Resolution

Collect investigation material and contact your local customer support.

60338 ERROR: *class* : already exists

Explanation

Class *class* already exists.

Resolution

Change the class name in the configuration table with the *sdxconfig Convert* command and re-execute the command.

60339 ERROR: *class* : already exists in another node

Explanation

Class *class* already exists in another node.

Resolution

Change the class name in the configuration file with the `sdxcfg Convert` command and re-execute the command.

60340 ERROR: *device* : assigned to *disk1* and *disk2* in configuration table

Explanation

In the configuration table one physical disk *device* is assigned to disks *disk1* and *disk2*.

Resolution

Change the physical disk assigned to *disk1* or *disk2* in the configuration table with the `sdxcfg Convert` command and re-execute the command.

60341 ERROR: *device* : failed to get physical disk information

Explanation

Acquisition of geometry information, partition table, or the device number of physical disk *device* failed.

Resolution

Check whether physical disk *device* is normally operating.

60342 ERROR: *device* : size must be size blocks

Explanation

The physical disk *device* must have *size* blocks.

Resolution

Replace physical disk *device* with a physical disk that has *size* blocks. Alternatively, change physical disk *device* in the configuration table to another physical disk that has *size* blocks with the `sdxcfg Convert` command.

60343 ERROR: *device* : private slice size must be size blocks

Explanation

The size of the private slice of physical disk *device* is different from that described in the configuration table and therefore the class object configuration cannot be restored. The private slice of *device* must have *size* blocks.

Resolution

Exchange physical disk *device* for a physical disk whose private slice size is *size* blocks. Alternatively, change physical disk *device* in the configuration table to another physical disk whose private slice size is *size* blocks with the `sdxcfg Convert` command.

60344 ERROR: mismatch of class names in private slices on *device1* and *device2*

Explanation

The class names of physical disks *device1* and *device2* stored in the private slices do not match and therefore *device1* and *device2* cannot be registered with one class.

Resolution

Check the system configuration and edit the physical disks described in the configuration table with the `sdxcfg Convert` command.

60345 ERROR: *device* : mismatch of disk names, *disk1* in private slice, *disk2* in configuration table

Explanation

The disk name *disk1* stored in the private slice does not match the disk name *disk2* described in the configuration table and therefore physical disk *device* cannot be registered as disk *disk2* with a class.

Resolution

Check the system configuration and edit the physical disks described in the configuration table with the `sdxconfig Convert` command.

60346 ERROR: *class* : restoration based on configuration file failed

Explanation

Restoration of the *class* configuration based on configuration file *file* failed.

Resolution

Collect investigation material and contact your local customer support.

60347 ERROR: *class*: some node trying to get class master privilege

Explanation

The local node or some other node is trying to obtain master privileges of shared class *class*.

Resolution

If necessary, try the command again after a while.

60348 ERROR: *class*: class master not found

Explanation

The class master of shared class *class* is not found and no operation is possible for *class*.

Resolution

If necessary, try the command again after a while.

E.4.2 Warning Messages (62000-62099)

62000 WARNING: spare disk *disk* too small

Explanation

The size of *disk* is too small and may not function as a spare disk.

Resolution

Specify a larger disk, and execute the command again.

62001 WARNING: *device*: write error, *ermo*=*ermo*

Explanation

A write error occurred in physical disk *device*.

Resolution

If *device* is a write-locked disk that is connected to a switch group, no action is required. Otherwise, remove *device* from the class and check the status of *device*.

62002 WARNING: *group*: free blocks are reduced

Explanation

Free blocks on *group* was reduced.

Resolution

You may not be able to create a volume with sufficient capacity. Execute commands as needed and attempt recovery.

62003 WARNING: another disk must be connected to *group*

Explanation

You must connect another disk to *group*.

Resolution

Connect another disk.

62004 WARNING: *object*: copying not completed successfully

Explanation

Synchronization copying did not complete successfully.

Resolution

Disk failure may have occurred. Identify the cause by referring to GDS log message and syslog message.

62005 WARNING: *object*: gave up wait for the completion of copying by a cancel request

Explanation

Synchronization copying was canceled before completion.

Resolution

Check the status of the object *object*. If the copying is in progress, there is no need to work around. If it is not in progress, re-execute the copying where it is necessary. For details, see "Appendix D Command Reference."

62007 WARNING: *group*: no spare disk available

Explanation

There is no valid spare disk in *group*.

Resolution

Define a spare disk as needed.

62008 WARNING: *object.volume*: cannot attached due to in *status* status

Explanation

Slice *object.volume* could not be attached since it is in *status* status.

Resolution

Check the slice status and cancel *status* status as needed.

62009 WARNING: *volume*: no need to resize volume

Explanation

There was no need to resize the volume *volume*.

Resolution

No particular resolution required.

62010 WARNING: *disk.volume*: special file(s) not found

Explanation

Special file for slice *disk.volume* could not be found.

Resolution

No particular resolution required.

62011 WARNING: *node*: node in stopped *status*

Explanation

node is in STOP status.

Resolution

No particular resolution required. However, promptly activating the node is recommended.

62012 WARNING: *node*: node in abnormal status

Explanation

node is in abnormal status.

Resolution

No particular resolution required. However, promptly recovering the node and activating it normally is recommended.

62013 WARNING: *object*: already in *status* status

Explanation

object is already in *status* status.

Resolution

No particular resolution required.

62014 WARNING: *disk* : *device* : write error, *ermo*=*ermo*

Explanation

A write error occurred on physical disk *device*.

Resolution

If *device* is a write-locked disk that is connected to a switch group, no action is required. In other situations, have *device* switchable and check its status.

62015 WARNING: *disk* : *device* : *node*: write error, *erno*=*erno*

Explanation

A write error occurred on physical disk *device* on node *node*.

Resolution

If *device* is a write-locked disk that is connected to a switch group, no action is required. In other situations, have *device* switchable and check its status.

62016 RNING: *object*: no need to update attribute value

Explanation

There was no need to change *object* attributes.

Resolution

No particular resolution required.

62017 WARNING: *pslice*: entry in */etc/fstab* not updated, unsupported file system type *fstype*

Explanation

Entry in the */etc/fstab* file relevant to the physical slice *pslice* was not updated since *fstype* is a file system type not supported by the *sdxroot* command.

Resolution

Edit the */etc/fstab* file or change the file system configuration according to need.

62018 WARNING: correct */etc/fstab* and file system configuration before rebooting, or system may not be booted

Explanation

The */etc/fstab* file contains an invalid entry. When an entry for a file system, such as */* (root), */usr*, and */var*, essential to system start-up is invalid, it will be impossible to start the system if rebooted.

Resolution

Refer to a GDS's WARNING message output just before this message and pinpoint the invalid entry in the */etc/fstab* file. Be sure to repair the */etc/fstab* file and the file system configuration as needed prior to rebooting the system.

62019 WARNING: *volume*: entry in */etc/fstab* not updated, unsupported file system type *fstype*

Explanation

Since file system type *fstype* is not supported, entry in */etc/fstab* file for *volume* could not be updated.

Resolution

Directly edit */etc/fstab* as necessary.

62020 WARNING: ignored parameter *param* for EC

Explanation

Since EC function will be used for copying, parameter *param* was ignored.

62021 WARNING: ignored parameter *param* for OPC

Explanation

Since OPC function will be used for copying, parameter *param* was ignored.

62022 WARNING: ignored parameter *param* for TimeFinder

Explanation

Parameter *param* was ignored for copying with TimeFinder.

62023 WARNING: ignored parameter *param* for SRDF

Explanation

Parameter *param* was ignored for copying with SRDF.

62024 WARNING: *proxy*: no session

Explanation

No copy session exists between the proxy *proxy* and the master

Resolution

No special action is required.

E.4.3 Information Messages (64000-64099)

64000 INFO: waiting for a response from sdxservd daemon...

Explanation

Awaiting response from sdxservd daemon.

64001 INFO: *class*: created class

Explanation

class was created.

64002 INFO: *disk*: created disk

Explanation

disk was registered.

**64003 INFO: *device*: disabled access to physical special files
*/dev/devices****

Explanation

Physical special files can no longer be accessed.

64004 INFO: *class*: removed class

Explanation

class was removed.

64005 INFO: *group*: created group

Explanation

group was created.

64006 INFO: *disk*: connected disk to group *group*

Explanation

disk was connected to *group*.

64007 INFO: *lgroup*: connected group to another group *hgroup*

Explanation

Group *lgroup* was connected to another group *hgroup*.

64008 INFO: *group*: removed group

Explanation

group was removed.

64009 INFO: *object*: waiting for the completion of copying...

Explanation

Waiting for synchronization copying to complete.

64010 INFO: *object*: copying completed successfully

Explanation

Synchronization copying completed successfully.

64011 INFO: *volume*: created volume

Explanation

volume was created.

**64012 INFO: *volume*: started volume on node *node*
*/dev/sfdsk/class/dsk/volume***

Explanation

volume was started on *node*. You can now access via special file.

64013 INFO: *volume*: stopped volume on node *node*

Explanation

volume was stopped on *node*.

64014 INFO: *volume*: removed volume

Explanation

volume was removed.

64015 INFO: *volume*: resized volume

Explanation

The volume *volume* was resized.

64016 INFO: *disk*: disconnected disk from group *group*

Explanation

disk was disconnected from *group*.

64017 INFO: *lgroup*: disconnected group from group *hgroup*

Explanation

Group *lgroup* was disconnected from *hgroup*.

64018 INFO: *group*: free blocks are increased

Explanation

Free blocks on *group* increased.

64019 INFO: *disk*: removed disk

Explanation

disk was removed from class.

**64020 INFO: *device*: enabled access to physical special files
*/dev/devices****

Explanation

You can now access to physical special files.

64021 INFO: *disk.volume*: temporarily detached slice from volume *volume*

Explanation

Slice *disk.volume* was temporarily disconnected from *volume*.

**64022 INFO: *disk.volume*: enabled slice
/dev/sfdsk/class/dsk/*disk.volume***

Explanation

The slice *disk.volume* can now be accessed independently of the volume.
Use the special file displayed.

64023 INFO: *disk.volume*: disabled slice

Explanation

Slice *disk.volume* cannot be accessed any more.

64024 INFO: *disk.volume*: reattached slice to volume *volume*

Explanation

Slice *disk.volume* is attached to *volume* again.

64025 INFO: *disk: device*: ready for swapping

Explanation

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.

64026 INFO: *disk: device*: ready for swapping on node *node*

Explanation

On node *node*, physical disk *device* of disk *disk* is now switchable. If the disk unit supports hot swap, you may simply demount *device* connected to *node* and perform swapping. If it does not support hot swap, shot down *node* and turn off the power, then demount *device* and perform swapping.

64027 INFO: *disk*: restored disk

Explanation

disk restored its original status.

64028 INFO: *class*: class recovery completed successfully

Explanation

Recovery of the class *class* from the closed status succeeded.

64029 INFO: *volume*: recovered volume with disk *disk*

Explanation

volume was recovered based on *disk*.

64030 INFO: *volume*: recovered volume with lower level group *lgroup*

Explanation

Volume *volume* was recovered based on lower level group *lgroup*.

64031 INFO: *volume* : recovered volume in group *group*

Explanation

Volume *volume* in group *group* has been recovered.

64032 INFO: *disk*: recovered disk

Explanation

disk *disk* was recovered.

64033 INFO: no unsynchronized slice in *volume*

Explanation

There is no slice that needs synchronization copying to be started, or resumed in *volume*.

64034 INFO: *volume*: copying canceled

Explanation

Synchronization copying of *volume* has been cancelled.

64035 INFO: no copying slice exists in *volume*

Explanation

There is no slice in the process of copying in *volume*.

64036 INFO: *volume*: copying interrupted

Explanation

Synchronization copying on *volume* was canceled.

64037 INFO: *class*: updated attribute value of class

Explanation

Attribute value of *class* was updated.

64039 INFO: *disk*: updated attribute value of disk

Explanation

Attribute value of *disk* was updated.

64040 INFO: *group*: updated attribute value of group

Explanation

Attribute value of *group* was updated.

64041 INFO: *volume*: updated attribute value of volume

Explanation

Attribute value of *volume* was updated.

64042 INFO: *slice*: updated attribute value of slice

Explanation

Attribute value of *slice* was updated.

64043 INFO: updated */etc/fstab* file

Explanation

The */etc/fstab* file was updated.

64045 INFO: updated */etc/opt/FJSVsdX/modules/sfdsk.conf* file

Explanation

The */etc/opt/FJSVsdX/modules/sfdsk.conf* file was updated.

64047 INFO: updated */boot/efi/efi/string/elilo.conf* file

Explanation

The */boot/efi/efi/string/elilo.conf* file was updated. *string* is the directory name.

64048 INFO: completed definitions of system disk(s) mirroring

Explanation

Definition for system disk mirroring is complete.

**64049 INFO: system must be rebooted:
shutdown -r now**

Explanation

Reboot the system by executing the following command.
shutdown -r now

64050 INFO: completed definitions of system disk(s) mirroring cancellation

Explanation

Canceling definition for system disk mirroring is complete.

**64051 INFO: *disk.volume*: started slice on node *node*
*/dev/sfdsk/class/dsk/disk.volume***

Explanation

The slice *disk.volume* was activated on *node*.
You may access the slice by using the displayed special file.

64052 INFO: *disk.volume*: stopped slice on node *node*

Explanation

Slice *disk.volume* was stopped on *node*.
You cannot access slice until it is activated again.

64053 INFO: *disk.volume*: disabled slice on remote node *node*

Explanation

Slice *disk.volume* has become unavailable on remote node *node*.

**64054 INFO: *disk.volume*: current node *node* took over and started slice
*/dev/sfdsk/class/dsk/disk.volume***

Explanation

The current node took over the slice *disk.volume* and activated it.
You may access the slice by using the displayed special file.

64055 INFO: *volume*: copying parameter modification completed

Explanation

Modification of parameter related to the copying process of *volume* is complete.

64056 INFO: configuration parameters modification completed

Explanation

Modification of configuration parameter is complete.

64057 INFO: *proxy*: related and joined to master *master*

Explanation

Proxy *proxy* has been related and joined to master *master*.

64058 INFO: *proxy*: created proxy volume

Explanation

Proxy volume *proxy* has been created.

64059 INFO: *proxy*: broken relation with master

Explanation

Relationship between master and proxy *proxy* has been cancelled.

64060 INFO: *proxy*: forced to break relation with master

Explanation

Relationship between master and proxy *proxy* was forcibly cancelled.

64061 INFO: *proxy*: must be recovered by `sdxfix(8)` command

Explanation

In order to restore data of proxy *proxy*, volume must be restored with `sdxfix(8)` command.

**64062 INFO: *volume*: started volume on node *node*,...
*/dev/sfdsk/class/dsk/volume***

Explanation

The volume *volume* started on the nodes *node*,.... It can be accessed with the special file.

64063 INFO: *proxy*: parted from master

Explanation

Proxy *proxy* was parted from master.

64064 INFO: *proxy*: parted from master instantly

Explanation

Proxy *proxy* was immediately parted from master.

64065 INFO: *proxy*: rejoined to master *master*

Explanation

Proxy *proxy* has been rejoined to master *master*.

64066 INFO: *master*: restored data from proxy *proxy*

Explanation

Copying from proxy *proxy* to master *master* started.

64067 INFO: *master*: restored data from proxy *proxy* instantly

Explanation

Copying from proxy *proxy* to master *master* started. Even if the copying is in process, you may start master volumes and access valid data.

64068 INFO: swapped slice components between master and proxy

Explanation

Slice comprising proxy and slice comprising master have been swapped.

64069 INFO: *proxy*: session canceled

Explanation

A session of the disk unit's copy function that resided between proxy *proxy* and the master was canceled.

64070 INFO: *proxy*: related to master *master*

Explanation

Master *master* and proxy *proxy* were related and parted.

64071 INFO: *proxy*: updated data with master *master*

Explanation

Data of proxy *proxy* was updated with data of master *master*.

64072 INFO: *proxy*: updated data with master *master* instantly

Explanation

Data of proxy *proxy* was instantly updated with data of master *master*.

64073 INFO: *proxy*: specified regions updated

Explanation

Data in the specified regions on the proxy *proxy* was updated.

64074 INFO: *master*: specified regions restored from proxy *proxy*

Explanation

Based on the proxy *proxy*, data in the specified regions on the master *master* was restored.

**64075 INFO: *volume*: executing fsck(8) command:
*details***

Explanation

The file system on the alternative root volume *volume* will be inspected and repaired with the fsck(8) command. *details* is a message of the fsck(8) command.

64076 INFO: updated /etc/fstab file on alternative root volume *volume*

Explanation

The /etc/fstab file on the alternative root volume *volume* was updated.

64077 INFO: volume will be used as dump device in alternative boot environment

Explanation

The volume *volume* is used as a dump device in an alternative boot environment. *volume* is made available as a dump device when the alternative boot environment is enabled.

64079 INFO: updated /boot/efi/efi/string/elilo.conf file on alternative root volume *volume*

Explanation

The /boot/efi/efi/string/elilo.conf file on the alternative root volume *volume* was updated. *string* is the directory name.

64081 INFO: updated boot devices:

old=*old-boot-device* ...

new=*new-boot-device* ...

Explanation

The boot devices changed from *old-boot-device* ... to *new-boot-device* ... Be sure to take a copy of *old-boot-device* ... and *new-boot-device* ... To cancel the switching to the alternative boot environment, reverse the boot order using the Change Boot Order option in the EFI boot manager's Boot option maintenance menu.

64082 INFO: completed definitions of switching boot environment

Explanation

The boot environment switching is now defined.

64084 INFO: completed definitions of alternative boot environment:

current-boot-device=*current-boot-device* ...

alternative-boot-device=*alternative-boot-device* ...

Explanation

The alternative boot environment is now defined. The boot device names in the current boot environment are *current-boot-device* ... and those in the alternative boot environment are *alternative-boot-device* ... Be sure to take a copy of *current-boot-device* ... and *alternative-boot-device* ... Reverse the boot order using the Change Boot Order menu in the EFI boot manager's Boot option maintenance menu to switch to the alternative boot environment.

**64088 INFO: *class*: configuration backup created
*file***

Explanation

The object configuration of class *class* was backed up in object table format to configuration file *file*.

**64089 INFO: configuration table converted successfully
file**

Explanation

The configuration table was converted successfully. *file* indicates a configuration file to which the configuration table was output.

64090 INFO: device: disk ID information not removed

Explanation

Physical disk *device* retains unremoved disk identification information.

64091 INFO: class: waiting for restoration...

Explanation

Waiting for restoration of the object configuration of class *class*.

**64092 INFO: class: restored based on configuration file
file**

Explanation

The *class* configuration was restored based on configuration file *file*.

64093 INFO: class: current node node got class master privilege

Explanation

Local node *node* obtained master privileges of shared class *class*.

E.4.4 Fix Messages (66000)

66000 TO FIX: Usage: usage

Resolution

Follow the *usage* syntax and execute the command.

E.4.5 Internal Error Messages (68000)

68000 command: severity: module: internal error, details

Explanation

An internal error occurred. *details* gives you the cause of the error, *command* the command name, *severity* the severity of message, and *module* the module name (usually the internal function name) where the event was detected.

The following message "a" may be output if the `sdxcopy -C` command is executed at node startup or during GDS command execution. If the `sdxcopy -C` command has ended normally (the exit status is 0), there is no problem.

- a. `sdxcopy: ERROR: SDX_COPY_STATUS_WAIT2: internal error, sdxerrno=145`

Resolution

In the event of "a", if the exit status of the `sdxcopy -C` command is 0, the message "a" may be ignored. For the `sdxcopy -C` command's exit status, refer to the following message recorded in the `sdxservd` daemon's log file `/var/opt/FJSVsdx/msglog/sdxservd.log`.

```
SDX:sdxservd: INFO: command exited, exit-status=exitstat:  
sdxcopy -C -c class -v volume,...
```

In other events, collect investigation material and contact your local customer support.

E.5 Operation Management View Messages

This section describes the GDS Management View messages.

E.5.1 Error Messages (5000-5099)

5000 Command line output: *details*

Explanation

An error given in *details* occurred when executing the command.

Resolution

Identify the cause referring to the command line output. When you cannot identify the cause, contact your local customer support.

5001 Executing the command was unsuccessful. Continue?

Explanation

Executing the command was unsuccessful.

Resolution

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

5002 The following exceptions were received: *details*

Explanation

Exceptions given in *details* occurred.

Resolution

Contact your local customer support.

5003 The input object name is already defined. Change the object name.

Explanation

The input object name is already defined.

Resolution

Change the object name.

5004 An error occurred in the log monitoring mechanism on node *node*. Monitoring tasks will stop. Check the status of the node and update to the latest information.

Explanation

An error such as a node crash was detected in the log monitoring mechanism.

Resolution

Verify that the node is in operation, and update to the latest information.

5005 Reading GDS configuration information was unsuccessful on node *node*. Contact the system administrator.

Explanation

Reading GDS configuration information was unsuccessful on node *node*.

Resolution

Contact your local customer support.

5006 An invalid class name was specified. Input another name.

Explanation

An invalid character was specified for the class name.

Resolution

Click <OK> and type a class name again. For information on characters available for class names, see "A.1.1 Object Name."

5007 An invalid group name was specified. Input another name.

Explanation

An invalid character was specified for the group name.

Resolution

Click <OK> and type a group name again. For information on characters available for group names, see "A.1.1 Object Name."

5008 An invalid volume name was specified. Input another name.

Explanation

An invalid character was specified for the volume name.

Resolution

Click <OK> and type a volume name again. For information on characters available for volume names, see "A.1.1 Object Name."

5009 An invalid disk name was specified. Input another name.

Explanation

An invalid character was specified for the disk name.

Resolution

Click <OK> and type a disk name again. For information on characters available for disk names, see "A.1.1 Object Name."

5010 The file system type or mount point has not been set. Select a file system type and/or input a mount point.

Explanation

<Create> was selected before specifying the file system type or mount point.

Resolution

Select a file system type and/or input a mount point, and click <Create>.

5011 The input mount point is already in use. Change the mount point.

Explanation

The input mount point is already in use with a different file system.

Resolution

Change the mount point, and click <Create>.

5012 A system disk cannot be removed from group configuration disks.

Explanation

In system disk setting, deletion of a system disk was attempted.

Resolution

System disks cannot be deleted from group configuration disks.

5013 You cannot add a physical disk. Up to 8 disks can be added to group configuration disks.

Explanation

In system disk setting, adding the ninth disk to group configuration disks was attempted.

Resolution

A maximum of 8-way multiplex mirroring is supported with system disks. You must limit the group configuration disks to eight.

5014 There are no free slices in the system disk. Secure a free slice before setting the system disk.

Explanation

There is no free slice on the system disk for which mirror settings should be applied.

Resolution

Secure a free slice on the system disk by executing the parted(8) command and such. After securing a free slice, set the system disk again.

5015 There are unexpected device ID on node *node*. Exit.

Explanation

Starting GDS Management View failed due to unsupported GDS configuration information on node *node*.

Resolution

Contact your local customer support.

5016 The order of removing disks and groups composing the group is wrong. Remove them in the proper order.

Explanation

A disk or group not removable was selected.

Resolution

When you remove disks and groups composing a group, you must select them in the inverse order against the added order.

5017 The stripe width setting is wrong. Specify an integer that is the power of 2.

Explanation

An invalid character or value was specified for the stripe width.

Resolution

Click <OK>, and input a valid stripe width again.

5018 The specified name is already in use. Input another name.

Explanation

The specified name is in use.

Resolution

Click <OK> and type an unused class name, and then click <Next>.

5019 An invalid volume size was specified. Input a valid size.

Explanation

The specified volume size is invalid.

Resolution

Click <OK> and type a valid volume size. When the maximum size is unknown, you can apply the maximum size to [Volume Size] by clicking <Maximum Size>.

5020 The GDS Management View processing was aborted.

Explanation

The GDS Management View processing was aborted. A command the GDS Management View issued may be timed out.

Resolution

Collect investigation material and contact your local customer support.

5021 The status gathering was aborted.

Explanation

The process to gather object statuses was aborted. The status gathering command the GDS Management View issued may be timed out.

Resolution

Collect investigation material and contact your local customer support.

E.5.2 Warning Messages (5000,5100-5199)

5000 Command line output: *details*

Explanation

An error given in *details* occurred when executing the command.

Resolution

Identify the cause referring to the command line output. When you cannot identify the cause, contact your local customer support.

5100 The input class name is already defined. Register the disk with this class?

Explanation

The input class name is already defined.

Resolution

Click <Yes> to add the disk to the specified class, or <No> to register the disk with another class.

5101 The input group name is already defined. Connect the disk with this group?

Explanation

The input group name is already defined.

Resolution

Click <Yes> to additionally connect the disk with the specified group, or <No> to connect the disk with another group.

5102 If you register a physical disk with this class, all the data contained in the disk will be lost. Continue?

Explanation

All of the disk data will be lost if a physical disk is registered with the class.

Resolution

Click <Yes> to register the physical disk, or <No> to cancel.

5103 If you remove the volume, all the volume data will be lost. Continue?

Explanation

All of the volume data will be lost if you remove the volume.

Resolution

Click <Yes> to remove the volume, or <No> to cancel.

5104 The volume will be recovered. Continue?

Explanation

A volume that cannot be accessed will be recovered.

Resolution

Click <Yes> to recover the volume, or <No> to cancel.

5105 The selected disk is normal. Swap disks?

Explanation

A normal disk was selected for swapping.

Resolution

Click <Yes> to swap the disks, or <No> to cancel.

**5108 The physical disk will be excluded from GDS.
To swap the disks, the relative slices will stop. Detaching the disk will affect data reliability. Continue?**

Explanation

When physical disks are swapped, the relative slices will stop operating until the restore is complete, and the reliability of the volume will fall.

Resolution

Click <Yes> to swap physical disks, or <No> to cancel.

**5109 The physical disk will be included into GDS.
All the data in the physical disk will be lost. Continue?**

Explanation

When a physical disk is restored, data in the physical disk will be lost.

Resolution

Click <Yes> to restore the physical disk, or <No> to cancel.

**5110 The slice will be detached from the volume.
Detaching the slice will affect data reliability. Continue?**

Explanation

Detaching a slice will lower reliability.

Resolution

Click <Yes> to detach the slice, or <No> to cancel.

**5111 The detached slice will be attached to the volume.
All the data in the slice will be lost. Continue?**

Explanation

When a slice is attached, data in the slice will be overwritten.

Resolution

Click <Yes> to attach the slice, or <No> to cancel.

5112 The volume will be stopped. Continue?

Explanation

The volume will be stopped.

Resolution

Click <Yes> to stop volume, or <No> to cancel.

**5113 Removing all disks configuring a class will result in removing the class as well.
Continue?**

Explanation

Removing all the disks in a class will result in removing the class as well.

Resolution

Click <Yes> to remove class, or <No> to cancel.

5114 Removing all disks configuring a group will result in removing the group as well. Continue?

Explanation

Removing all the disks in a group will result in removing the group as well.

Resolution

Click <Yes> to remove group, or <No> to cancel.

5115 A file system will be created on the selected volume. Continue?

Explanation

A file system will be created on the selected volume. If you continue, all the data in the volume will be lost.

Resolution

Click <Yes> to create a file system, or <No> to cancel.

5116 The selected volume will be unmounted and mount information for the volume will be deleted. It will be no longer available as a file system. Continue?

Explanation

The selected volume will be unmounted and mount information for the volume will be deleted. If it is deleted, the volume can no longer be used as a file system.

Resolution

Click <Yes> to remove the file system, or <No> to cancel.

5117 Changes on GDS configuration information were detected. Update to the latest configuration information?

Explanation

Commands other than those from GDS Management View changed the GDS configuration information.

Resolution

Click <Yes> to update the GDS Management View display, or <No> to cancel.

5118 The physical disk will be excluded from GDS. Continue?

Explanation

In order to swap physical disks, an unused disk or a spare disk will be detached. This will not affect the volume.

Resolution

Click <Yes> to swap the physical disks, or <No> to cancel.

5119 A file system has been configured on the selected volume. Remove the file system from the File System Configuration.

Explanation

The volume cannot be removed because of a file system configured on the volume.

Resolution

Remove the file system from the [File System Configuration] screen and then remove the volume.

5120 Physical disks have not been swapped. Verify that swapping took place. Are you sure you want to attach the disk?

Explanation

Physical disks may not have been swapped.

Resolution

Click <Yes> to attach the physical disk, or <No> to cancel.

5121 The status of physical disks connected to the system will be updated. Updating to the latest information may consume time. Continue?

Explanation

The status of physical disks connected to the system will be updated. During the update, operations from GDS Management View are unavailable.

Resolution

Click <Yes> to update the physical disk information, or <No> to cancel.

5122 A file system is created in a volume within the class. You cannot change attributes because it will change the path. Change the attributes after removing the file system.

Explanation

Renaming the class containing a volume with a file system was attempted.

Resolution

Renaming the class will result in changing the access path. First delete the file system.

5123 A file system is created in the volume. You cannot change attributes because it will change the path. Change the attributes after removing the file system.

Explanation

Renaming the volume with a file system was attempted.

Resolution

Renaming the volume will result in changing the access path. First delete the file system.

5124 No disk is available as a system disk after unmirroring. Recover the defective disk before unmirroring the system disk.

Explanation

Unmirroring a system disk was attempted when the mirrored disk was abnormal.

Resolution

First swap the abnormal physical disk for another. After recovering the physical disk, unmirror the system disk.

5125 Unmirroring the system disk will affect system reliability. Are you sure you want to unmirror?

Explanation

<Unmirror> was selected in the [Unmirror System Disk] screen.

Resolution

Click <Yes> to unmirror the system disk, or <No> to cancel.

5126 Since synchronization copying is in process, configuring system disk mirroring is not complete. After synchronization copying is complete, select System Disk Settings again.

Explanation

While configuring system disk mirroring, a volume in process of synchronization copying was detected.

Resolution

Click <OK> and close the message. From the GDS Management View screen, monitor the root class volume in process of synchronization copying. Verify that the volume status changes from "warnings" to "active," and try configuring mirroring for the system disk again.

5127 Configuring system disk mirroring is not complete. You must rewrite the mount information. Proceed with the system disk mirroring configuration?

Explanation

System disk mirroring has not been configured yet.

Resolution

Click <Yes> to rewrite the mount information and finish the configuration, or <No> to cancel.

5128 The slice will become inaccessible. Continue?

Explanation

The detached slice will be made inaccessible.

Resolution

Click <Yes> to restrict access to the slice, or <No> to allow access.

5129 The slice will become accessible. Continue?

Explanation

The slice being inaccessible as a result of stopping the slice will be made accessible.

Resolution

Click <Yes> to allow access to the slice, or <No> to restrict access.

5130 Copying of the slice will be cancelled. It cannot be resumed. Continue?

Explanation

The resynchronization copying of the slice will be cancelled. The copying cannot be resumed from the point where it stops.

Resolution

Click <Yes> to cancel resynchronization copying, or <No> not to cancel.

5131 Copying of the slice will start (resume). Continue?

Explanation

Copying of the slice will be performed to synchronize the data of the slices of the volume. Copying of an interrupted slice will resume as well.

Resolution

Click <Yes> to start (resume) resynchronization copying, or <No> not to start (resume).

5132 Remove a group in the 'root' class?

Explanation

A group in the "root" class is going to be removed.

Resolution

Click <Yes> to remove the group, or <No> to cancel.

5133 An error was detected on *node*. Display *node* information?

Explanation

An object error was detected on another node *node* configuring the cluster system.

Resolution

Click <Yes> to display the information about the node with the error, or <No> not to display.

5134 An error was detected on *node*. Check the node for the status after the current operation ends.

Explanation

An object error was detected on another node *node* configuring the cluster system.

Resolution

Check the node on which the error was detected for the status after the current operation ends.

5139 The proxy will be joined. The joined proxy will lose existing data. Do you want to continue?

Explanation

If the proxy is joined, its contents are overwritten by data of the master volume.

Resolution

Click <Yes> to join the proxy, or <No> to cancel the process.

5140 The proxy will be joined. Do you want to continue?

Explanation

The master group and the proxy group will be joined.

Resolution

Click <Yes> to join the proxy, or <No> to cancel the process.

5141 The proxy will be related. Do you want to continue?

Explanation

The master and the proxy will be related.

Resolution

Click <Yes> to relate the proxy, or <No> to cancel the process.

5142 The relationship between the master and the proxy will be canceled. Do you want to continue?

Explanation

The relationship between the master and the proxy will be canceled.

Resolution

Click <Yes> to break the relationship, or <No> to cancel the process.

5143 The master and the proxy will be rejoined. The rejoined proxy will lose existing data. Do you want to continue?

Explanation

If the proxy is rejoined, its contents are overwritten by data of the master.

Resolution

Click <Yes> to rejoin the proxy, or <No> to cancel the process.

5144 Slices of the master and the proxy will be swapped. Do you want to continue?

Explanation

Slices comprising the master and the proxy are switched.

Resolution

Click <Yes> to switch the slices, or <No> to cancel the process.

5145 Reboot the system to enable the system disk settings.

Explanation

System disk setting ended normally. It is necessary that the system be rebooted to enable system disk settings.

Resolution

Click <OK> to close the message, and reboot the system immediately.

5146 Reboot the system in multi-user mode to complete unmirroring.

Explanation

System disk unmirroring ended normally. It is necessary that the system be rebooted in multi-user mode to complete system disk unmirroring.

Resolution

Click <OK> to close the message, and reboot the system in multi-user mode immediately.

5147 The system disk includes a master or a proxy volume. Break the relationship of the master and proxy before unmirroring the system disk.

Explanation

The group to which the system disk belongs includes a master or a proxy volume.

Resolution

Cancel the proxy related to the system disk and then unmirror the system disk.

E.5.3 Information Messages (5200-5299)

5200 The slice was attached and the volume will be restored.

Explanation

The slice was attached. In order to restore the volume, resynchronization copying will start.

Resolution

Click <OK> and close the message. Check the completion of resynchronization copying in volume information field on the main screen.

For the resynchronization copying progress, check the slice information field in the main screen.

5201 The volume restore is complete.

Explanation

The volume was restored.

Resolution

Click <OK> and close the message.

5202 The slice was detached. The following access path is available.

Access path: */dev/sfdsk/classname/dsk/slicename*

Explanation

The slice was detached to make it accessible.

Resolution

Pay attention to the displayed access path and click <OK>.

5203 The disk was excluded from GDS. Swap physical disks. Ask your local customer support for physical disk swapping.

Explanation

For swapping, a physical disk was excluded from GDS.

Resolution

Click <OK>, and contact your local customer support to have physical disks swapped.

5204 The physical disk was included into GDS. The volume will be restored.

Explanation

The swapped physical disk was placed under GDS.

Resolution

Click <OK> and close the message.

5206 A file system was created.

Explanation

A file system was created and mount information was added.

Resolution

Click <OK> and close the message.

5207 The file system was removed.

Explanation

The file system and mount information was removed.

Resolution

Click <OK> and close the message.

5209 Status monitoring will be updated every *interval* minutes.

Explanation

The object monitoring interval was changed.

Resolution

Click <OK> and close the message.

5210 Exit Global Disk Services?

Explanation

GDS Management View will be exited.

Resolution

Click <Yes> to exit, or <No> to continue operation.

5211 The following files were updated to complete mirroring system disk.

/etc/fstab

/boot/efi/efi/string/elilo.conf

Explanation

The system disk setting process completed mirroring definition. *string* is the directory name.

Resolution

Click <OK> and close the message.

5212 System disk unmirroring is complete.

Explanation

The system disk unmirroring process completed unmirroring.

Resolution

Click <OK> and close the message.

5213 The slice was made inaccessible. Activate the slice in order to gain access.

Explanation

The detached slice was made inaccessible.

Resolution

Click <OK> and close the message.

5214 The slice is now accessible.

Explanation

The detached slice was made accessible.

Resolution

Click <OK> and close the message.

5215 The selected slice is neither temp nor temp-stop status. You cannot perform the operation on the selected slice.

Explanation

The [Stop/Activate Slice] operation was attempted on a slice not in the temp or temp-stop status.

Resolution

Click <OK> and close the message.

5216 Copying of the slice was cancelled.

Explanation

The resynchronization copying of the slice was cancelled.

Resolution

Click <OK> and close the message.

5217 Copying of the slice was started (resumed).

Explanation

The resynchronization copying was started (resumed).

Resolution

Click <OK> and close the message.

5218 The selected volume is not in the copy status. You cannot perform the operation on the selected volume.

Explanation

The copying operation was attempted while the slices in the volume are not in the copy status.

Resolution

Click <OK> and close the message.

5219 The selected slice is neither active nor stop status. You cannot perform the operation on the selected slice.

Explanation

The [Detach Slice] operation was attempted on a slice not in the active or stop status.

Resolution

Click <OK> and close the message.

5220 There is a slice in the copy status in the selected volume. You cannot perform the operation on the selected volume.

Explanation

The [Start Copying] operation was attempted on a volume with a slice in process of synchronization copying.

Resolution

Click <OK> and close the message.

5221 There are no slices in the invalid status in the selected volume. You cannot perform the operation on the selected volume.

Explanation

The [Start Copying] operation was attempted on a volume without a slice in the invalid status.

Resolution

Click <OK> and close the message.

5222 There is no available proxy for the selected master. Create a volume that can be joined first.

Explanation

There is no volume that can be joined as a proxy volume to the selected master volume.

Resolution

Create a volume to be joined to the master volume first conforming to the following rules.

- Belonging to the class to which the master volume belongs
- Belonging to a mirror group or to a single disk
- Not belonging to the group to which the master volume belongs
- Not belonging to the single disk to which the master volume belongs
- Not related to another volume as a master or a proxy
- Equivalent to the master volume in size

For details, see "A.1.8 Exception to Proxy Configuration" and the description of the Join option in "D.14 sdxproxy - Proxy object operations."

5223 There is no available proxy for the selected master. Create a group that can be joined first.

Explanation

There is no group that can be joined as a proxy group to the master group.

Resolution

Create a group to be joined to the master group first conforming to the following rules.

- Belonging to the class to which the master group belongs
- Mirror type
- Not related to another group as a master or a proxy
- Including no volume

For details, see "A.1.8 Exception to Proxy Configuration" and the description of the Join option in "D.14 sdxproxy - Proxy object operations."

5224 Joining the proxy succeeded. Synchronization copying will start.

Explanation

The master and the proxy were joined. Synchronization copying from the master to the proxy will start to synchronize them.

Resolution

Click <OK> to close the message. Check the synchronization copying status in the slice information field in the main screen.

5225 There is no available proxy for the selected master. Create a volume that can be related first.

Explanation

There is no volume that can be related to the selected master volume.

Resolution

Create a volume to be related to the master volume conforming to the following rules.

- Belonging to the class to which the master volume belongs
- Belonging to a mirror group or to a single disk
- Not belonging to the group to which the master volume belongs
- Not belonging to the single disk to which the master volume belongs
- Not related to another volume as a master or a proxy
- Equivalent to the master volume in size

For details, see "A.1.8 Exception to Proxy Configuration" and the description of the Relate option in "D.14 sdxproxy - Proxy object operations."

5226 There is no available proxy for the selected master. Create a group that can be related first.

Explanation

There is no group that can be related as a proxy group to the selected master group.

Resolution

Create a group to be related to the master group conforming to the following rules.

- Belonging to the class to which the master group belongs
- Mirror type
- Not related to another group as a master or a proxy
- Whose volume configuration (offset and size) is same as that of the master group

For details, see "A.1.8 Exception to Proxy Configuration" and the description of the Relate option in "D.14 sdxproxy - Proxy object operations."

5227 Relating the proxy succeeded.

Explanation

The master and the proxy were related.

Resolution

Click the <OK> to close the message.

5228 Breaking the proxy succeeded.

Explanation

The relationship between the master and the proxy was broken.

Resolution

Click the <OK> to close the message.

5229 Parting the proxy succeeded.

Explanation

The proxy was parted from the master.

Resolution

Click <OK> to close the message.

5230 Rejoining the proxy succeeded. The master and the proxy will be resynchronized

Explanation

The master and the proxy were rejoined. Resynchronization copying from the master to the proxy will start to resynchronize them.

Resolution

Click <OK> to close the message. Check the resynchronization copying status in the volume Information field in the main screen. For the resynchronization copying progress, check the slice information field in the main screen.

5231 The update of the proxy started.

Explanation

Copying from the master to the proxy started.

Resolution

Click <OK> to close the message. Check the completion of the resynchronization copying in the volume information field in the main screen. For the copying status, check the slice information field in the main screen.

5232 Updating the proxy succeeded. Start up the proxy to access.

Explanation

Copying from the master to the proxy started. Before the copying ends, you may start the proxy for access. Contents of the proxy are overwritten by the data of the master at the moment when the copying started.

Resolution

Click <OK> to close the message. Check the copying status. in the slice information field on the main screen.

5233 The restoration of the master started.

Explanation

Copying from the proxy to the master started.

Resolution

Click <OK> to close the message. Check the completion of the copying in the volume information field in the main screen. For the copying progress, check the slice information field on the main screen.

5234 Restoring the master succeeded. Start up the master to access.

Explanation

Copying from the proxy to the master started. Before the copying ends, you may start the master for access. Contents of the master are overwritten by the data of the proxy at the moment when the copying started.

Resolution

Click <OK> to close the message. Check the copying status in the volume information field on the main screen.

5235 Rejoining the proxy succeeded, and the restoration of the master started.

Explanation

The master and the proxy were rejoined and synchronization copying from the proxy to the master started.

Resolution

Click <OK> to close the message. Check the completion of the resynchronization copying in the volume information field on the main screen. For the resynchronization copying progress, check the slice information field in the main screen.

5236 Rejoining the proxy and restoring the master succeeded. Start up the master to access.

Explanation

The master and the proxy were rejoined and resynchronization copying from the proxy to the master started. Before the resynchronization copying ends, you may start the master for access. Contents of the master are overwritten by the data of the proxy at the moment when the copying started.

Resolution

Click <OK> to close the message. For the resynchronization copying status, check the slice information field on the main screen.

5237 Swapping the slices of the master and the proxy succeeded.

Explanation

Slices comprising the master and the proxy were switched.

Resolution

Click <OK> to close the message.

Appendix F Troubleshooting

Since GDS automatically attempts recovery when various abnormalities are detected, users do not have to be aware of most abnormalities. However, depending on the type of abnormality, or when multiple abnormalities occur at the same time, automatic recovery may not be possible.

This section describes resolutions to failures on objects and physical disks managed by GDS for systems administrators.



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

If the slice status is one of the following statuses, take action as indicated for the relevant situation.

- Mirror slice configuring the mirror volume is in INVALID status.
- The copy destination slice was made INVALID due to an I/O error generated on the copy source slice during synchronization copying.
- Slice configuring the volume is in TEMP status.
- Slice configuring volume is in TEMP-STOP status.
- Slice configuring the volume is in COPY status.
- Slice configuring the volume is in NOUSE status.

(1) Mirror slice configuring the mirror volume is in INVALID status.

Explanation

You can check the status of the slice configuring the volume as follows.

```
# sdxinfo -S -o Volume1
```

OBJ	CLASS	GROUP	DISK	VOLUME	STATUS
slice	Class1	Group1	Object1	Volume1	ACTIVE
slice	Class1	Group1	Object2	Volume1	INVALID

In this example, among the slices that exist in volume Volume1, the slice within Object2 is in INVALID status, as shown in the STATUS field. Object2 is a disk or lower level group connected to the highest level mirror group Group1.

The following five events could possibly cause the INVALID status of the mirror slice Volume1.Object2.

- An I/O error occurred on the mirror slice Volume1.Object2.
 - (Cause a)
 - A disk component relevant to Object2 failed to operate properly, and an I/O error occurred on the mirror slice Volume1.Object2.
 - (Cause b)
 - A component other than disks relevant to Object2 (such as an I/O adapter, an I/O cable, an I/O controller, a power supply, and a fan) failed to operate properly, and an I/O error occurred on the mirror slice Volume1.Object2.
- An I/O error occurred on the mirror slice Volume1.Object1 during synchronization copying to the mirror slice Volume1.Object2.
 - (Cause a')
 - A disk component relevant to Object1 failed to operate properly during synchronization copying to the mirror slice Volume1.Object2, and an I/O error occurred on the copy source slice Volume1.Object1.
 - (Cause b')
 - A component other than disks relevant to Object1 (such as an I/O adapter, an I/O cable, an I/O controller, a power supply and a fan) failed to operate properly during synchronization copying to the mirror slice Volume1.Object2, and an I/O error occurred on the mirror slice Volume1.Object1 that is the copy source.
- Others
 - (Cause c)
 - Synchronization copying to the mirror slice Volume1.Object2 was cancelled as the result of a cause such as [Cancel Copying] selection from the GDS Management View, sdxcopy command execution, or a power outage.

Resolution

1) Identify the physical disk name of a faulty disk using the sdxinfo command.

(Example A1)

# sdxinfo -G -o Volume1							
OBJ	NAME	CLASS	DISKS	BLKS	FREEBLKS	SPARE	
group	Group1	Class1	Object1:Object2	17596416	17498112	0	

# sdxinfo -D -o Volume1 -e long										
OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	FREEBLKS	DEVCONNECT	STATUS	E
disk	Object1	mirror	Class1	Group1	sda	17596416	*	node1:node2	ENABLE	*
disk	Object2	mirror	Class1	Group1	sdb	17596416	*	node1:node2	ENABLE	*

In this example, Object2 is a disk connected with the highest level group Group1. As indicated in the E field, an I/O error occurred on the disk Object2, and the possible cause is (Cause a) or (Cause b). The physical disk name of the disk Object2 is sdb as shown in the DEVNAM field.

In example A1, if the value 0 is in the E field of the disk Object2 including a slice in the INVALID status and if the value 1 is in the E field of the disk Object1 that is mirrored with the disk Object2, it indicates an I/O error occurred on the disk Object1 and the possible cause is (Cause a') or (Cause b'). In such a case, see "(2) The copy destination slice was made INVALID due to an I/O error generated on the copy source slice during synchronization copying" and perform restoration.

If the E field of any disk does not contain the value 1 in example A1, the possible cause is (Cause c).

(Example B1)

```
# sdxinfo -G -o Volume1
OBJ  NAME  CLASS  DISKS          BLKS  FREEBLKS  SPARE
----  ----  ----  -----          ----  -
group Group1  Class1  Object1:Object2 35127296 35028992  0
group Object1  Class1  Disk1:Disk2     35127296 *          *
group Object2  Class1  Disk3:Disk4     35127296 *          *
```

```
# sdxinfo -D -o Volume1 -e long
OBJ  NAME  TYPE  CLASS  GROUP  DEVNAM  DEVBLKS  FREEBLKS  DEVCONNECT  STATUS  E
----  ----  ----  ----  ----  -----  -
disk  Disk1  stripe  Class1  Object1  sda      17596416 *          node1:node2  ENABLE  0
disk  Disk2  stripe  Class1  Object1  sdb      17596416 *          node1:node2  ENABLE  0
disk  Disk3  stripe  Class1  Object2  sdc      17682084 *          node1:node2  ENABLE  1
disk  Disk4  stripe  Class1  Object2  sdd      17682084 *          node1:node2  ENABLE  0
```

In this example, Object2 is a lower level group connected with the highest level group Group1. As indicated in the E field, an I/O error occurred on the disk Disk3 connected with Object2 and the possible cause is (Cause a) or (Cause b). The physical disk name corresponding to Disk3 is sdc as shown in the DEVNAM field.

In example B1, if the 0 value is in the E field of the disks (Disk3 and Disk4) connected with the lower level group Object2 including a slice in the INVALID status and if the value 1 is in the E field of the disk (Disk1 or Disk2) connected with the lower level group Object1 that is mirrored with the lower level group Object2, it indicates an I/O error occurred on the disk connected with Object1 and the possible cause is (Cause a') or (Cause b'). In such a case, see "(2) The copy destination slice was made INVALID due to an I/O error generated on the copy source slice during synchronization copying" and perform restoration.

In example B1, if there are no disks with "1" in the E field, the possible cause of the INVALID status is (Cause c).

2) Refer to disk driver log messages and check the physical disk abnormalities.

The causes of disk hardware failures can be failures or defects in components such as I/O adapters, I/O cables, I/O controllers, power supplies, and fans other than the disks.

Contact your local customer support and specify which component failed, or might be defective.

If there are no failures or defective components, the possible cause of the INVALID is (Cause c).

The resolution procedure is illustrated below for each of the three causes a, b, and c.

a. For (Cause a)

a1) Perform the following operations before and after disk swapping. For the procedures for swapping disks from Operation Management View, see "5.3.4 Disk Swap."

Before swapping the disks, execute the following command.

(Example A1)

```
# sdxswap -O -c Class1 -d Object2
```

In the example, disk Object2 connected to the highest level group Group1 will be swapped.

(Example B1)

```
# sdxswap -O -c Class1 -d Disk3
```

In the example, disk Disk3 will be swapped. Disk3 is a disk connected to lower level group Object2, which is a lower level group of the highest level group Group1.

a2) Swap disks.

a3) After swapping disks, execute the following command.

(Example A3)

```
# sdxswap -I -c Class1 -d Object2
```

or

(Example B3)

```
# sdxswap -I -c Class1 -d Disk3
```

a4) Check the slice status according to step 3).

b. For (Cause b)

b1) Shut down the system once, repair the disabled part, and reboot the system. Consequently, synchronization copying is performed and the mirroring status is restored.

b2) Check the slice status according to step 3).

c. For (Cause c)

c1) Perform synchronization copying of mirror volume.

```
# sdxcopy -B -c Class1 -v Volume1
```

c2) Check the slice status according to step 3).

3) You can confirm the recovery of the slice configuring the volume, as shown below.

```
# sdxinfo -S -o Volume1
OBJ   CLASS  GROUP  DISK   VOLUME  STATUS
-----
slice Class1  Group1 Object1 Volume1  ACTIVE
slice Class1  Group1 Object2 Volume1  ACTIVE
```

In this example, the slices within Object1 and Object2 are both in ACTIVE status. This indicates that the recovery process was completed successfully.

(2) The copy destination slice was made INVALID due to an I/O error generated on the copy source slice during synchronization copying.

Explanation

When an I/O error occurs on the copy source slice during synchronization copying, the copy destination slice becomes INVALID while the source slice is still ACTIVE.

The following two events are possible causes.

(Cause a')

A disk component of the copy source failed to operate properly, and an I/O error occurred on the copy source slice.

(Cause b')

A component other than the copy source disk (such as an I/O adapter, an I/O cable, an I/O controller, a power supply, and a fan) failed to operate properly during synchronization copying, and an I/O error occurred on the copy source slice.

For details on determining whether the status is relevant to one of these events and identifying the physical disk name of a faulty disk, see [Explanation] and step 1) of [Resolution] described in "(1) Mirror slice configuring the mirror volume is in INVALID status."

Resolution

First examine the physical disk abnormalities referring to disk driver log messages and so on. Then contact your local customer support and locate the disabled or faulty part.

When the possible cause is (Cause b), shut down the system once, repair the disabled part, and reboot the system. Consequently, synchronization copying is performed and the mirroring status is restored. When the possible cause is (Cause a), follow the procedures below and repair the slice. The procedure is illustrated for each of the following three situations.

- A. For /(root), /usr, or /var [Itanium]
- B. For the swap area [Itanium]
- C. For others (other than /(root), /usr, /var, swap)

The following illustrates restoration procedures when the class name is Class1, the volume name is Volume1, the name of a faulty disk of the copy source is Disk1, and the name of a disk of the copy destination is Disk2 as examples.

```
# sdxinfo -S -o Volume1
OBJ CLASS GROUP DISK VOLUME STATUS
----
slice Class1 Group1 Disk1 Volume1 ACTIVE
slice Class1 Group1 Disk2 Volume1 INVALID

# sdxinfo -G -o Volume1
OBJ NAME CLASS DISKS BLKS FREEBLKS SPARE
----
group Group1 Class1 Disk1:Disk2 17596416 17498112 0

# sdxinfo -D -o Volume1 -e long
OBJ NAME TYPE CLASS GROUP DEVNAM DEVBLKS FREEBLKS DEVCONNECT STATUS E
----
disk Disk1 mirror Class1 Group1 sda 17596416 * node1:node2 ENABLE 1
disk Disk2 mirror Class1 Group1 sdb 17596416 * node1:node2 ENABLE 0
```

A. For /(root), /usr, or /var [Itanium]

Contact your local customer support.

B. For the swap area [Itanium]

B.1) Remove the volume from the swap area.

```
# swapoff /dev/sfdsk/Class1/dsk/Volume1
```

Depending on the part or severity of failure in disks that constitute the volume, the swapoff(8) command may fail due to an I/O error. In this event, remove the volume from the swap area performing steps B.1.1) through B.1.2).

B.1.1) Comment out the /dev/sfdsk/Class1/dsk/Volume1 line in the /etc/fstab file to prevent use of the volume as a swap area after the system is rebooted.

```
# vim /etc/fstab
Before edit:
/dev/sfdsk/Class1/dsk/Volume1  swap                swap  defaults  0 0
After edit:
#/dev/sfdsk/Class1/dsk/Volume1  swap                swap  defaults  0 0
```

B.1.2) Reboot the system.

```
# shutdown -r now
```

B.2) Stop the volume.

```
# sdxvolume -F -c Class1 -v Volume1
```

B.3) Restore the status of the copy destination slice in INVALID status.

```
# sdxfix -V -c Class1 -d Disk2 -v Volume1
```

B.4) Verify that the restored copy destination slice is in STOP status and the copy source slice is in INVALID status now.

```
# sdxinfo -S -o Volume1
OBJ   CLASS  GROUP  DISK   VOLUME  STATUS
-----
slice Class1  Group1 Disk1   Volume1 INVALID
slice Class1  Group1 Disk2   Volume1 STOP
```

B.5) Start the volume.

```
# sdxvolume -N -c Class1 -v Volume1
```

B.6) Add the volume to the swap area again.

```
# swapon /dev/sfdsk/Class1/dsk/Volume1
```

When step B.1.1) was performed, undo the edit that was made in the /etc/fstab file.

```
# vim /etc/fstab
Before edit:
#/dev/sfdsk/Class1/dsk/Volume1  swap                swap  defaults  0 0
After edit:
/dev/sfdsk/Class1/dsk/Volume1  swap                swap  defaults  0 0
```

B.7) Remove the faulty copy source disk from GDS management to give it a replaceable status.

```
# sdxswap -O -c Class1 -d Disk1
```

B.8) Swap the faulty copy source disk.

B.9) Put the swapped disk back in control of GDS management to make it available.

```
# sdxswap -l -c Class1 -d Disk1
```

C. For others (other than /(root), /usr, /var, and swap)

C.1) Exit applications using the volume.

C.2) Unmount the file system on the volume when it has been mounted.

In this example, the volume has been used as an ext3 file system.

```
# umount /dev/sfdsk/Class1/dsk/Volume1
```

Depending on the part or the severity of failure in disks that compose the volume, the umount command may fail due to an I/O error. In this event, unmount the file system performing steps C.2.1) through C.2.3).

C.2.1) If class Class1 is registered with a cluster application, remove the cluster application.

C.2.2) Comment out the /dev/sfdsk/Class1/dsk/Volume1 line in the /etc/fstab file to prevent mounting of the volume after the system is rebooted.

C.2.3) Reboot the system.

C.3) Stop the volume.

```
# sdxvolume -F -c Class1 -v Volume1
```

C.4) Restore the status of the copy destination slice in the INVALID status.

```
# sdxfix -V -c Class1 -d Disk2 -v Volume1
```

C.5) Verify that the restored copy destination slice is in the STOP status and the copy source slice is the INVALID status now.

```
# sdxinfo -S -o Volume1
OBJ   CLASS  GROUP  DISK   VOLUME  STATUS
-----
slice Class1  Group1 Disk1   Volume1 INVALID
slice Class1  Group1 Disk2   Volume1 STOP
```

C.6) Start the volume.

```
# sdxvolume -N -c Class1 -v Volume1
```

C.7) The consistency of volume data may have lost. Restore the backup data or perform repair using the fsck(8) command if necessary.



Note

If an I/O error occurred in the copy source slice during just resynchronization copying after the system went down, restoration may possibly be performed with the fsck(8) command.

When step C.2.2) was performed, undo the edit that was made in the /etc/fstab file.

When step C.2.1) was performed, re-create the cluster application removed in step C.2.1).

C.8) Remove the faulty copy source disk from GDS management to make it a replaceable status.

```
# sdxswap -O -c Class1 -d Disk1
```

C.9) Swap the faulty copy source.

C.10) Put the swapped disk back in control of GDS management to make it available.

```
# sdxswap -I -c Class1 -d Disk1
```

(3) Slice configuring the volume is in TEMP status.

Explanation

The slice was not attached after it has been detached with the sdxslice command. Or else, you have not performed [Attach Slice] after performing [Detach Slice] from Operation Management View.

Resolution

Attach the slice again with the sdxslice command, or perform [Attach Slice] from Operation Management View as necessary.

(4) Slice configuring volume is in TEMP-STOP status.

Explanation

The slice was not activated after it has been stopped with the sdxslice command, or the detached node is not current node.

Or else, you have not performed [Stop Slice] after performing [Activate Slice] from Operation Management View.

Resolution

Activate slice or take over slice with the sdxslice command as needed. Or, perform [Activate Slice] from Operation Management View.

(5) Slice configuring the volume is in COPY status.

Explanation

In order to attach a slice, synchronization copying is currently in process.

Or, synchronization copying is in process between master and proxy.

Resolution

Wait until synchronization copying is complete. Note that a slice in the process of synchronization copying will not restrict you from accessing an active volume.

(6) Slice configuring the volume is in NOUSE status.

Explanation

When the status of disk related to slice is either in DISABLE or SWAP status, the slice becomes NOUSE to inhibit slice operation.

Resolution

Recover disk in DISABLE or SWAP status. For details, see "F.1.2 Disk Status Abnormality."

F.1.2 Disk Status Abnormality

If the disk status is one of the following statuses, take action as indicated for the relevant situation.

- Disk is in DISABLE status.
- Disk is in SWAP status.
- Disk is in I/O error status.

(1) Disk is in DISABLE status.

Explanation

Disk ID information is checked when booting the system. If the ID information is regarded invalid at the time, the disk will become DISABLE.

There are three reasons that may cause this DISABLE status.

(Cause a)

You changed the I/O cable connection and booted the system.

(Cause b)

You swapped the disk and booted the system.

(Cause c)

Disk failure occurred.

Resolutions are described for each of the above causes.

Resolution

a) If the possible cause is (Cause a), shut down the system, connect the I/O cables correctly and boot again.

b) If the possible cause is (Cause b), shut down the system, return the original disk and boot again. If swapping the disks is necessary, follow the procedures and swap the disks. For procedures on disk swapping, see "D.8 sdxswap - Swap disk" and "5.3.4 Disk Swap."

c) If the possible cause is (Cause c), follow the procedures and swap the disks.

(2) Disk is in SWAP status.

Explanation

If the sdxswap -O command is executed, or [Swap Physical Disk] is performed from Operation Management View, the disk will be in SWAP status.

Resolution

Complete the disk swap and restore the disk by either executing the sdxswap -I command, or by performing [Restore Physical Disk] from Operation Management View.

(3) Disk is in I/O error status.

Explanation

If an I/O error occurs, 1 (one) is given as a value in the E field displayed with the `sdxinfo -e long` command and the disk icon in GDS Management View turns red. Here, the disk is in ENABLE (enabled) status.

Resolution

Check the statuses of volumes and slices relative to the disk with an I/O error, and perform restoration, referring to "F.1.3 Volume Status Abnormality" or "F.1.1 Slice Status Abnormality". The I/O error status is removed, during the restoration by:

- Executing [Swap Physical Disk] in GDS Management View for the relevant disk
- Executing the `sdxswap -O` command for the relevant disk
- Normally ending synchronization copying of volumes relative to the relevant disk

If none of the operations above are performed during the recovery and the I/O error remains unremoved, after restoring hardware and such, use the `sdxfix -D` command to remove the I/O error status.

```
# sdxfix -D -c class name -d disk name -e online
```

F.1.3 Volume Status Abnormality

If the volume status is one of the following statuses, take action as indicated for the relevant situation.

- Mirror volume is in INVALID status.
- Single volume is in INVALID status.
- Stripe volume or volume in concatenation group is in INVALID status.
- Master volume is in INVALID status.
- Proxy volume is in INVALID status.
- Volume is in STOP status.
- I/O error occurs although mirror volume is in ACTIVE status.
- An I/O error occurs on a single volume.
- An I/O error occurs on a stripe volume or a volume in a concatenation group.

(1) Mirror volume is in INVALID status.

Explanation

You can confirm the status of the volume as shown below.

```
# sdxinfo -V -o Volume1
OBJ   NAME   CLASS  GROUP  SKIP JRM 1STBLK  LASTBLK  BLOCKS  STATUS
-----
volume *      Class1 Group1  *   *      0      65535    65536  PRIVATE
volume Volume1 Class1 Group1  off  on    65536  17596415 17530880 INVALID
```

In this example, volume Volume1 that exists in the highest level group Group1 is in INVALID status, as shown in the STATUS field.

If none of the mirror slices consisting the mirror volume contains valid data (ACTIVE or STOP), the mirror volume becomes INVALID. You cannot start a volume in INVALID status.

There are two reasons that may cause this INVALID status.

(Cause a)

Disk is in DISABLE status.

(Cause b)

Master-proxy relationship was cancelled forcibly while master data was being copied to proxy.

Resolution

- 1) Confirm that there is a disk in DISABLE status within the group with which the volume is associated as follows.

(Example A1)

```
# sdxinfo -G -o Volume1
OBJ  NAME  CLASS  DISKS          BLKS  FREEBLKS SPARE
-----
group Group1 Class1  Disk1:Disk2    17596416  0  0

# sdxinfo -D -o Volume1
OBJ  NAME  TYPE  CLASS  GROUP  DEVNAM  DEVBLKS  DEVCONNECT  STATUS
-----
disk  Disk1  mirror Class1  Group1  sda     17596416  node1:node2  ENABLE
disk  Disk2  mirror Class1  Group1  sdb     17596416  node1:node2  DISABLE
```

In this example, disks Disk1 and Disk2 are connected to the highest level mirror group Group1. The disk Disk2 is in DISABLE status as shown in the STATUS field.

(Example B1)

```
# sdxinfo -G -o Volume1
OBJ  NAME  CLASS  DISKS          BLKS  FREEBLKS SPARE
-----
group Group1 Class1  Group2:Group3  35127296  17530880  0
group Group2 Class1  Disk1:Disk2    35127296  *  0
group Group3 Class1  Disk3:Disk4    35127296  *  0

# sdxinfo -D -o Volume1
OBJ  NAME  TYPE  CLASS  GROUP  DEVNAM  DEVBLKS  DEVCONNECT  STATUS
-----
disk  Disk1  stripe Class1  Group2  sda     17596416  node1:node2  ENABLE
disk  Disk2  stripe Class1  Group2  sdb     17596416  node1:node2  DISABLE
disk  Disk3  stripe Class1  Group3  sdc     17682084  node1:node2  ENABLE
disk  Disk4  stripe Class1  Group3  sdd     17682084  node1:node2  ENABLE
```

In this example, lower level stripe groups, Group2 and Group3 are connected to the highest level mirror group Group1. Disk Disk2 which is connected to Group2 is in DISABLE status as shown in the STATUS field.

- 2) If the possible cause is (Cause a), restore the disk by following the procedures in "F.1.2 Disk Status Abnormality."

- 3) From the disks and lower level groups connected to the highest level mirror group, determine the disk or lower level group to which the slice you will use to recover data belongs. Then, execute the sdxfix command to recover data.

(Example A3)

```
# sdxfix -V -c Class1 -d Disk1 -v Volume1
```

In this example, Volume1 is recovered after a slice in disk Disk1.

(Example B3)

```
# sdxfix -V -c Class1 -g Group3 -v Volume1
```

In this example, Volume1 is recovered after a slice in lower level stripe group Group3.

4) Start the volume.

```
# sdxvolume -N -c Class1 -v Volume1 -e nosync
```

5) Access Volume1 and check its contents. Restore backup data or run fsck to regain data integrity as necessary.

6) Perform synchronization copying on volume.

```
# sdxcopy -B -c Class1 -v Volume1
```

(2) Single volume is in INVALID status.

Explanation

You can confirm the status of the volume as shown below.

```
# sdxinfo -V -o Volume1
```

OBJ	NAME	CLASS	GROUP	SKIP	JRM	1STBLK	LASTBLK	BLOCKS	STATUS
volume	*	Class1	Disk1	*	*	0	32767	32768	PRIVATE
volume	Volume1	Class1	Disk1	off	on	32768	65535	32768	INVALID
volume	*	Class1	Disk1	*	*	65536	8421375	8355840	FREE

In the example, the single volume Volume1 that exists on single disk Disk1 is in INVALID status, as shown in the STATUS field.

You cannot start a volume in INVALID status.

There are two reasons that may cause this INVALID status.

(Cause a)

Single disk is in DISABLE status. In this case, the single slice becomes NOUSE status.

(Cause b)

Master-proxy relationship was cancelled forcibly while master data was being copied to proxy.

Resolution

1) Confirm that the single disk is in DISABLE status as shown below.

```
# sdxinfo -D -o Volume1
```

OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	DEVCONNECT	STATUS
disk	Disk1	single	Class1	*	sda	8355840	node1:node2	DISABLE

In this example, the single disk Disk1 is in DISABLE status, as shown in the STATUS field.

2) If the possible cause is (Cause a), restore the disk by following the procedures given in section "F.1.2 Disk Status Abnormality."

3) Execute the sdxfix command to recover the single volume's data.

```
# sdxfix -V -c Class1 -d Disk1 -v Volume1
```

4) Start the volume.

```
# sdxvolume -N -c Class1 -v Volume1
```

5) Access Volume1 and check its content. Restore backup data or run the fsck command to regain data integrity as necessary.

(3) Stripe volume or volume in concatenation group is in INVALID status.

Explanation

You can confirm the status of the volume as shown below.

```
# sdxinfo -V -o Volume1
```

OBJ	NAME	CLASS	GROUP	SKIP	JRM	1STBLK	LASTBLK	BLOCKS	STATUS
volume *		Class1	Group1	*	*	0	65535	65536	PRIVATE
volume	Volume1	Class1	Group1	off	on	65536	17596415	17530880	INVALID

In this example, volume Volume1 that exists in the highest level group Group1 is in INVALID status, as shown in the STATUS field.

If any of the disks related to volume is in DISABLE status, the slices consisting that volume become NOUSE status, and the volume becomes INVALID. You cannot start a volume in INVALID status.

Resolution

1) You can confirm the status of the disk related to the volume as shown below.

```
# sdxinfo -G -o Volume1 -e long
```

OBJ	NAME	CLASS	DISKS	BLKS	FREEBLKS	SPARE	MASTER	TYPE	WIDTH
group	Group1	Class1	Group2:Group3	70189056	65961984	*	*	stripe	32
group	Group2	Class1	Disk1:Disk2	35127296	*	*	*	concat	*
group	Group3	Class1	Disk3:Disk4	35127296	*	*	*	concat	*


```
# sdxinfo -D -o Volume1
```

OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	DEVCONNECT	STATUS
disk	Disk1	concat	Class1	Group2	sda	17596416	node1:node2	ENABLE
disk	Disk2	concat	Class1	Group2	sdb	17596416	node1:node2	DISABLE
disk	Disk3	concat	Class1	Group3	sdc	17682084	node1:node2	ENABLE
disk	Disk4	concat	Class1	Group3	sdd	17682084	node1:node2	ENABLE

In this example, the lower level concatenation groups Group2 and Group3 are connected to the highest level stripe group Group1, and Disk2 connected to Group2 is in DISABLE status as shown in the STATUS field.

2) Follow the procedures in "F.1.2 Disk Status Abnormality" and restore the disk status.

3) Execute the sdxfix command to recover the volume's data. With -g option, indicate the highest level group name (in this example, Group1).

```
# sdxfix -V -c Class1 -g Group1 -v Volume1
```

4) Start the volume.

```
# sdxvolume -N -c Class1 -v Volume1
```

5) Access Volume1 and check its content. Restore backup data or run the fsck command to regain data integrity as necessary.

(4) Master volume is in INVALID status.

Explanation

If the copying process fails while copying data from the proxy volume to the master volume because of an I/O error or such, the status of the master volume to which the data is being copied becomes INVALID.

Resolution

1) Check if there is a DISABLE status disk in the group to which the volume belongs with the following command.

# sdxinfo -D -o Volume1								
OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	DEVCONNECT	STATUS
disk	Disk1	mirror	Class1	Group1	sda	8421376	*	ENABLE
disk	Disk2	mirror	Class1	Group1	sdb	8421376	*	DISABLE

In this example, Disk2 is in DISABLE status.

If there is a disk in DISABLE status, see section "(1) Disk is in DISABLE status" in "F.1.2 Disk Status Abnormality," and check which of the causes listed in that section applies. If the possible cause is (Cause a) or (Cause b), follow the procedures and restore the disk.

2) Follow the procedures given in section "(1) Mirror slice configuring the mirror volume is in INVALID status" in "F.1.1 Slice Status Abnormality," and check if there is a disk hardware abnormality. If there is, identify the faulty part. When the abnormality is caused by a failed or defective non-disk component, repair the faulty part.

3) Procedures to restore the data for different scenarios are given below.

- Non-disk component failure
 - When recovering data using the proxy volume:
-> Follow steps a) to restore.
 - When recovering data using backup data on media such as tapes:
-> Follow steps b) to restore.
- When caused by a disk component failure:
 - When the disk does not belong to master group:
 - some disks connected to the group have failed
-> Follow steps c) to restore.
 - all disks connected to the group have failed
-> Follow steps d) to restore.
 - When the disk belongs to master group:
-> Follow steps e) to restore.

a) Procedures to recover master volume data using proxy volume.

a1) In order to check if the proxy volume that will be used to recover data is separated from the master volume, execute the sdxinfo -V -e long command, and check the PROXY field.

a2) If the proxy volume is not separated, execute the following command.

```
# sdxproxy Part -c Class1 -p Volume2
```

a3) Exit all applications accessing the proxy volume. When using the proxy volume as a file system, execute umount. When using the proxy volume as a file system, execute umount.

a4) If the proxy volume is started, execute the following command.

```
# sdxvolume -F -c Class1 -v Volume2
```

a5) Recover master volume data using the proxy volume's data.

```
# sdxproxy RejoinRestore -c Class1 -p Volume2
```

b) Procedures to recover data using backup data.

b1) When the volume is in INVALID status, you must first change it to STOP status. Decide on the disk (slice) you wish to use to recover data, and execute the sdxfix command.

```
# sdxfix -V -c Class1 -d Disk1 -v Volume1
```

In this example, Volume 1 is restored after a slice in Disk 1.

b2) When the volume to be restored is stopped, start it with the following command.

```
# sdxvolume -N -c Class1 -v Volume1 -e nosync
```

b3) Access the volume to be restored and check its contents. Restore backup data or run fsck to regain data integrity as necessary.

b4) When mirroring is configured with the volume, perform synchronization copying.

```
# sdxcopy -B -c Class1 -v Volume1
```

c) Procedures to swap some disks connected to the group.

c1) If you restore the INVALID master volume later using data of a proxy volume related to the master volume, or use data of proxy volumes related to the master volume after restoring it, part the proxy volumes using the sdxproxy Part command.

```
# sdxproxy Part -c Class1 -p Volume2
```

c2) When there is a volume in INVALID status in the group, change it to STOP status with the sdxfix -V command. -d option indicates the disk without abnormality.

```
# sdxfix -V -c Class1 -d Disk1 -v Volume1
```

c3) Follow the procedures and swap the disks. For details, see "D.8 sdxswap - Swap disk" and "5.3.4 Disk Swap."

c4) Recover the master volume data. If data will be recovered using the proxy volume, follow procedures described in a). If data will be recovered using backup data on media such as tapes, follow procedures described in b).

d) Procedures to swap all disks connected to the group.

d1) Exit all applications accessing the master volume and the proxy volume that will be used to recover data. When using the proxy volume or the master volume as a file system, execute unmount.

d2) Stop the master volume and proxy volume in d1).

```
# sdxvolume -F -c Class1 -v Volume1  
# sdxvolume -F -c Class1 -v Volume2
```

d3) Execute the sdxproxy RejoinRestore command and restore the master volume data using proxy volume in d1). If the command terminates normally and the master volume is not in INVALID status, restoration process is complete, and you do not need to perform steps d4) and after.

```
# sdxproxy RejoinRestore -c Class1 -p Volume2
```

d4) Execute the sdxproxy Swap command and swap the slices of the master volume with the proxy volume in d1).

```
# sdxproxy Swap -c Class1 -p Volume2
```

d5) By performing step d4), the status of master volume will not be in INVALID status, and the status of the proxy volume becomes INVALID. Follow the procedures given in section "(5) Proxy volume is in INVALID status" in "F.1.3 Volume Status Abnormality," and restore the proxy volume in INVALID status.

d6) Execute the sdxproxy Swap command and swap the slices of the master volume and the proxy volume you swapped in step d4).

```
# sdxproxy Swap -c Class1 -p Volume2
```

e) Procedures to swap disks connected to the master group.

e1) Exit all applications accessing the master group, and volumes in the proxy group that will be used to recover data. When using the volume as a file system, execute umount.

e2) Stop all volumes in the master group and the proxy group in e1).

```
# sdxvolume -F -c Class1 -v Volume1
# sdxvolume -F -c Class1 -v Volume2
```

e3) Execute the sdxproxy RejoinRestore command and restore the master group data using the proxy group in e1). If the command terminates normally and all master volumes are not in INVALID status, restoration process is complete, and you do not need to perform steps e4) and after.

```
# sdxproxy RejoinRestore -c Class1 -p Volume2
```

e4) Execute the sdxproxy Swap command and swap the slices of the master group and the proxy group in e1).

```
# sdxproxy Swap -c Class1 -p Group2
```

e5) By performing step e4), the master volume will not be in INVALID status and the status of the proxy volume becomes INVALID. Follow the procedures given in section "(5) Proxy volume is in INVALID status" in "F.1.3 Volume Status Abnormality," and restore the proxy volume in INVALID status.

e6) Execute the sdxproxy Swap command and swap the slices of the master group and the proxy group you swapped in step e4).

```
# sdxproxy Swap -c Class1 -p Group2
```

(5) Proxy volume is in INVALID status.

Explanation

If the copying process fails while copying data from the master volume to the proxy volume because of an I/O error or such, the status of the proxy volume to which the data is being copied becomes INVALID.

Resolution

1) Check if there is a DISABLE status disk in the group to which the volume belongs with the following command.

```
# sdxinfo -D -o Volume1
```

OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	DEVCONNECT	STATUS
disk	Disk1	mirror	Class1	Group1	sda	8421376	*	ENABLE
disk	Disk2	mirror	Class1	Group1	sdb	8421376	*	DISABLE

In this example, Disk2 is in DISABLE status.

If there is a disk in DISABLE status, see section "(1) Disk is in DISABLE status" in "F.1.2 Disk Status Abnormality," and check which of the causes (There are three causes a, b and c listed.) listed in that section applies. If it is due to (Cause a) or (Cause b), follow the procedures and restore the disk.

2) Follow the procedures given in "(1) Mirror slice configuring the mirror volume is in INVALID status" in "F.1.1 Slice Status Abnormality," and check if there is a disk hardware abnormality. If there is, identify the faulty part. When the abnormality was caused by a failed or defective non-disk component, repair the faulty part.

3) Procedures to restore the data for different scenarios are given below.

- When caused by a non-disk component failure:
-> Follow steps a) to restore.
- When caused by a disk component failure:
 - disk does not belong to proxy group
 - some disks connected to the group have failed
-> Follow steps b) to restore.
 - When all disks connected to the group have failed:
-> Follow steps c) to restore.
 - When the disk belongs to proxy group:
-> Follow steps d) to restore.

a) Procedures to recover proxy volume data using the master volume.

a1) In order to check if the proxy volume is separated from the master volume, execute the `sdxinfo -V -e` long command, and check the PROXY field.

a2) If the proxy volume is not separated, execute the following command.

```
# sdxproxy Part -c Class1 -p Volume2
```

a3) Rejoin the proxy volume with the master volume.

```
# sdxproxy Rejoin -c Class1 -p Volume2
```

b) Procedures to swap some disks connected to the group.

b1) Cancel the relationship with master volume using the `sdxproxy Break` command.

```
# sdxproxy Break -c Class1 -p Volume2
```

b2) Separate the volumes that are in INVALID status in the group with the `sdxfix -V` command, and change them to STOP status. `-d` option indicates the disk without abnormality.

```
# sdxfix -V -c Class1 -d Disk1 -v Volume2
```

b3) Follow the procedures and swap the disks. For details, see "D.8 sdxswap - Swap disk," or section "5.3.4 Disk Swap."

b4) Join the master and the proxy again with the `sdxproxy Join` command.

```
# sdxproxy Join -c Class1 -m Volume1 -p Volume2
```

c) Procedures to swap all disks connected to the group.

c1) Cancel the relationship with the master using the `sdxproxy Break` command.

```
# sdxproxy Break -c Class1 -p Volume2
```

c2) Exit all applications accessing the volume in the group. When using the volume as a file system, execute `umount`.

c3) Stop all volumes in the group.

```
# sdxvolume -F -c Class1 -v Volume2
```

c4) Check the volume configuration of the group (such as volume names and sizes) with the `sdxinfo` command, and keep a note of it.

c5) Remove all volumes in the group.

```
# sdxvolume -R -c Class1 -v Volume2
```

c6) Follow the procedures and swap the disks. For details, see "D.8 sdxswap - Swap disk," or section "5.3.4 Disk Swap."

c7) Create the volume that you removed in step c5) again.

```
# sdxvolume -M -c Class1 -g Group2 -v Volume2 -s size
```

c8) Stop the volume you created in step c7).

```
# sdxvolume -F -c Class1 -v Volume2
```

c9) Join the master volume and the proxy volume again, with the sdxproxy Join command.

```
# sdxproxy Join -c Class1 -m Volume1 -p Volume2
```

d) Procedures to swap disks connected to the proxy group.

d1) Cancel the relationship with the master using the sdxproxy Break command.

```
# sdxproxy Break -c Class1 -p Group2
```

d2) Exit all applications accessing the volume in the group. When using the volume as a file system, execute unmount.

d3) Stop all volumes in the group.

```
# sdxvolume -F -c Class1 -v Volume2
```

d4) Remove all volumes in the group.

```
# sdxvolume -R -c Class1 -v Volume2
```

d5) Follow the procedures and swap the disks. For details, see "D.8 sdxswap - Swap disk," or section "5.3.4 Disk Swap."

d6) Join the master group and the proxy group again with the sdxproxy Join command.

```
# sdxproxy Join -c Class1 -m Group1 -p Group2 -a Volume1=Volume2:on
```

(6) Volume is in STOP status.

Explanation

Normally, volumes automatically start when the system is booted and become ACTIVE. The volume status will change to STOP when the volume is stopped with the Stop Volume menu in the GDS Management View or the sdxvolume -F command.

In a cluster system, among volumes within GDS shared classes registered with cluster applications, volumes other than proxy volumes start or stop according to the cluster application modes. If a cluster application is in Offline mode, volumes other than proxy volumes are in STOP status.

Accessing a volume in STOP status will result in an EIO error (I/O error) or an ENXIO error (No such device or address).



[See](#)

For the problem in a cluster system that volumes in a shared class not registered with a cluster application do not start at node startup, see "(4) The GFS Shared File System is not mounted on node startup" in "F.1.10 Cluster System Related Error."

Resolution

Start the volumes with the Start Volume menu in GDS Management View or the sdxvolume -N as necessary.

To start volumes within a GDS shared class registered with a cluster application, change the cluster application mode to Online.

(7) I/O error occurs although mirror volume is in ACTIVE status.

Explanation

A mirror volume consists of multiple slices, and in an event of an I/O error, the crashed slice will be detached. Therefore, accessing the volume will complete normally.

However, when an I/O error occurs when only one slice is ACTIVE amongst those configuring the volume, accessing the volume will result in an error. At such time, the status of the slice and the volume remains ACTIVE.

Probable situations resulting in such a problem will be described using a two-way multiplex mirroring configuration, where two disks or two lower level groups are connected to a group. As an example, means to circumvent such problems will also be described.

(Situation 1)

One of the slices was detached with the `sdxslice -M` in order to backup volume data. While accessing the volume, an I/O error occurred with the other slice.

(Prevention 1)

Before executing the `sdxslice -M` command, connect a reserved disk and temporarily configure a three-way multiplex mirroring, or make the mirrored volume available for backup.

(Situation 2)

While restoring a slice with an I/O error, an I/O error also occurred on another slice.

(Prevention 2)

By securing a spare disk within the class, effects due to delay in restoring the slice will be avoided to a certain degree

Resolution

Identify the cause of I/O error occurrence in the last ACTIVE slice, by referring to the disk driver log message.

Resolutions are described below assuming the following three circumstances:

- a. Error occurred due to a disk component failure. Will attempt recovery using backup data.
- b. Error occurred due to a disk component failure. Will attempt data recovery from a slice in INVALID status.
- c. Error occurred due to a failed or a defective non-disk component failure.

a. When the error cause is a disk component failure and recovery is performed using backup data

a1) When the error was caused by a disk component failure, no slice with valid data exists. Restore data from the backup data following the procedures given below.

a2) Exit the application accessing the volume. When the volume is used as a file system, execute the unmount command.

When I/O error occurs on the unmount command, execute `-f` option of the unmount command.

a3) Stop the volume with the `sdxvolume` command.

```
# sdxvolume -F -c Class1 -v Volume1
```

a4) If there is a TEMP status slice within the volume, attempt recovery following the procedures given in "F.1.1 Slice Status Abnormality."

a5) If there is a NOUSE status slice within the volume, attempt recovery following the procedures given in "F.1.1 Slice Status Abnormality."

a6) Record the volume size which can be checked as follows.

```
# sdxinfo -V -o Volume1
OBJ   NAME   CLASS  GROUP  SKIP  JRM  1STBLK  LASTBLK  BLOCKS  STATUS
-----
volume *     Class1 Group1  *     *     0       32767    32768   PRIVATE
volume Volume1 Class1  Group1 off    on     32768   4161535 4128768 STOP
```

In this example, the volume size would be 4128768 blocks given in Volume1 BLOCKS field.

a7) Remove the volume with the sdxvolume command.

```
# sdxvolume -R -c Class1 -v Volume1
```

a8) Swap disks following the procedures given in "5.3.4 Disk Swap" and "D.8 sdxswap - Swap disk."

a9) Create a volume with the sdxvolume command again. For the *number_of_blocks*, use the size recorded in a6), in this example, 4128768.

```
# sdxvolume -M -c Class1 -g Group1 -v Volume1 -s number_of_blocks
```

a10) Finally, restore the backup data to Volume1.

b. When the error cause is a disk component failure and data is restored from a slice in INVALID status

b1) When the error was caused by a disk failure, and when no backup data exists, or even if it did, the data is too old, restore data from the detached INVALID status slice, following the procedures given below.

b2) Exit the application accessing the volume. When the volume is used as a file system, execute the unmount command.

When I/O error occurs on the unmount command, execute -f option of the unmount command.

b3) Stop the volume with the sdxvolume command.

```
# sdxvolume -F -c Class1 -v Volume1
```

b4) If there is a TEMP status slice within the volume, attempt recovery following the procedures given in "F.1.1 Slice Status Abnormality."

b5) If there is a NOUSE status slice within the volume, attempt recovery following the procedures given in "F.1.1 Slice Status Abnormality."

b6) Determine the original mirror slice after the volume is recovered. Then, execute the sdxfix command.

(Example 1)

```
# sdxfix -V -c Class1 -d Disk2 -v Volume1
```

In this example, data is recovered from a mirror slice in the disk Disk2 which is connected to the highest level mirror group.

(Example 2)

```
# sdxfix -V -c Class1 -g Group2 -v Volume1
```

In this example, data is recovered from a mirror slice in the lower level group Group2 which is connected to the highest level mirror group.

b7) Start the volume.

```
# sdxvolume -N -c Class1 -v Volume1 -e nosync
```

b8) Create backup of Volume1 and regain data integrity by running fsck as necessary.

b9) Lastly, swap disks following the procedures given in "5.3.4 Disk Swap" and "D.8 sdxswap - Swap disk."

c. When the error cause is a non-disk component failure or defect

The slice with valid data exists within the disk, and shut down the system once, recover the failed component, and then reboot the system. Synchronization copying is automatically performed and the mirroring status will be recovered.

(8) An I/O error occurs on a single volume.

Explanation

Since a single volume consists of only one slice, accessing the volume at the time of an I/O error will result in an error. However, the status of slice and volume will remain ACTIVE.

Resolution

Identify the cause of I/O error occurrence by referring to the disk driver log message.

How to resolve the problem is described in two cases:

- a. When the error cause is a disk component failure and recovery is performed using backup data
- b. When the error cause is a non-disk component failure or defect

a. When the error cause is a disk component failure and recovery is performed using backup data

a1) In the event of a disk component failure, there will be no slice with valid data. Follow the procedures below and restore the data using the backup data. In this example, Disk1 (sda) has a failure.

```
# sdxinfo -D -o Disk1
OBJ  NAME  TYPE  CLASS  GROUP  DEVNAM  DEVBLKS  DEVCONNECT  STATUS
----  ----  ----  ----  ----  ----  ----  ----  ----
disk  Disk1  single  Class1  *      sda      8493876  node1:node2  ENABLE
```

a2) Search the volumes within the faulty disk using the sdxinfo command. And record the volume size,

```
# sdxinfo -V -o Disk1
OBJ  NAME  CLASS  GROUP  SKIP  JRM  1STBLK  LASTBLK  BLOCKS  STATUS
----  ----  ----  ----  ----  ---  ----  ----  ----  ----
volume *      Class1  Disk1  *      *      0      32767  32768  32768  PRIVATE
volume Volume1  Class1  Disk1  off    on    32768  65535  32768  ACTIVE
volume Volume2  Class1  Disk1  off    on    65536  4194303  4128768  ACTIVE
volume *      Class1  Disk1  *      *      4194304  8421375  4227072  FREE
```

In this example, Volume1 and Volume2 are within the faulty Disk1. The size of Volume1 would be 32,768 blocks as shown in the BLOCKS field. The size of Volume2 would be 4,128,768 blocks as shown in the BLOCKS field.

a3) Exit the application accessing the volume. When the volume is used as a file system, execute the unmount command. When I/O error occurs on the unmount command, execute `-f` option of the unmount command.

a4) Stop the volume with the `sdxvolume` command.

```
# sdxvolume -F -c Class1 -v Volume1,Volume2
```

a5) Remove the volumes with the `sdxvolume` command.

```
# sdxvolume -R -c Class1 -v Volume1
# sdxvolume -R -c Class1 -v Volume2
```

a6) Before swapping the disks, execute the following command.

```
# sdxswap -O -c Class1 -d Disk1
```



Note

If the disk is the only remaining disk in the disk class, the command results in an error as shown below. In that event, follow the steps a6'), a7') and a8').

```
SDX:sdxswap: ERROR: Disk1: The last ENABLE disk in class cannot be swapped
```

a7) Swap the disks.

a8) After swapping the disks, execute the following command.

```
# sdxswap -I -c Class1 -d Disk1
```

a6') Before swapping the disks, execute the following command.



Note

If no error is output in a6), the steps a6'), a7'), and a8') are not required.

```
# sdxdisk -R -c Class1 -d Disk1
```

a7') Swap the disks.

a8') After swapping the disks, execute the following command.

```
# sdxdisk -M -c Class1 -d sda=Disk1:single
```

a9) Create volumes with the `sdxvolume` command again. For the `-s` option, use the size recorded in 2a), in this example.

```
# sdxvolume -M -c Class1 -d Disk1 -v Volume1 -s 32768
# sdxvolume -M -c Class1 -d Disk1 -v Volume2 -s 4128768
```

a10) Finally, restore the backup data to Volume1 and Volume2.

b. When the error cause is a non-disk component failure or defect

Shut down the system once, recover the failed component, and then reboot the system. Slice data is valid and there is no need to restore the data.

(9) An I/O error occurs on a stripe volume or a volume in a concatenation group.

Explanation

Since a stripe volume or a volume within a concatenation group consists of only one slice, accessing the volume at the time of an I/O error will also result in an error. However, the status of slice and volume will remain ACTIVE.

Resolution

Identify the cause of I/O error occurrence by referring to the disk driver log message. You can confirm the error status of the disk related to volume and the physical disk name as shown below.

```
# sdxinfo -D -o Volume1 -e long
```

OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	FREEBLKS	DEVCONNECT	STATUS	E
disk	Disk1	concat	Class1	Group2	sda	17596416	*	node1:node2	ENABLE	0
disk	Disk2	concat	Class1	Group2	sdb	17596416	*	node1:node2	ENABLE	1
disk	Disk3	concat	Class1	Group3	sdc	17682084	*	node1:node2	ENABLE	0
disk	Disk4	concat	Class1	Group3	sdd	17682084	*	node1:node2	ENABLE	0

In this example, an I/O error occurs on Disk2, as shown in the E field. The physical disk name corresponding to Disk2 is sdb, as shown in the DEVNAM field.

How to resolve the problem is described in two cases:

- a. When the error cause is a disk component failure and recovery is performed using backup data
- b. When the error cause is a non-disk component failure or defect

a. When the error cause is a disk component failure and recovery is performed using backup data

a1) In the event of a disk component failure, there will be no slices with valid data. Follow the procedures below and restore the data using the backup data.

a2) Record the configuration information of the group that was related to the failed disk using the sdxinfo command.

```
# sdxinfo -G -o Disk2 -e long
```

OBJ	NAME	CLASS	DISKS	BLKS	FREEBLKS	SPARE	MASTER	TYPE	WIDTH
group	Group1	Class1	Group2:Group3	70189056	65961984	*	*	stripe	32
group	Group2	Class1	Disk1:Disk2	35127296	*	*	*	concat	*
group	Group3	Class1	Disk3:Disk4	35127296	*	*	*	concat	*

In this example, the lower level concatenation groups Group2 and Group3 are connected to the highest level stripe group Group1. The disks Disk1 and Disk2 are connected to Group2, and the disks Disk3 and Disk4 are connected to Group3. The stripe width for Group1 is 32 blocks.

a3) Search the volumes that exist in the highest level group that are related to the faulty disk using the `sdxinfo` command.

```
# sdxinfo -V -o Disk2
```

OBJ	NAME	CLASS	GROUP	SKIP	JRM	1STBLK	LASTBLK	BLOCKS	STATUS
volume *		Class1	Group1	*	*	0	65535	65536	PRIVATE
volume	Volume1	Class1	Group1	*	*	65536	98303	32768	ACTIVE
volume	Volume2	Class1	Group1	*	*	98304	4227071	4128768	ACTIVE
volume *		Class1	Group1	*	*	4227072	70189055	65961984	FREE

In this example, Volume1 and Volume2 exist in the highest level group Group1, that is related to the faulty disk Disk2. The size of Volume1 is 32768 blocks, and the size of Volume2 is 4128768 blocks as shown in the BLOCKS field.

a4) Exit the application accessing the volume. When the volume is used as a file system, execute `unmount` command. When I/O error occurs on `unmount` command, execute `-f` option of `unmount` command.

a5) Stop the volume with the `sdxvolume` command.

```
# sdxvolume -F -c Class1 -v Volume1,Volume2
```

a6) Remove the volumes with the `sdxvolume` command.

```
# sdxvolume -R -c Class1 -v Volume1
# sdxvolume -R -c Class1 -v Volume2
```

a7) Disconnect the faulty disk from the group. If the group is in a hierarchical structure, disconnect from the higher group in descending order.

```
# sdxgroup -D -c Class1 -h Group1 -l Group2
# sdxdisk -D -c Class1 -g Group2 -d Disk2
```

In this example, the faulty disk Disk2 is connected to Group2, and Group2 is connected to Group1. Therefore, you should disconnect Group2 first, and then Disk2.

a8) Before swapping the disks, execute the following command.

```
# sdxswap -O -c Class1 -d Disk2
```



Note

If the disk is the only remaining disk in the disk class, the command results in an error as shown below. In that event, follow the steps a8'), a9') and a10').

SDX:sdxswap: ERROR: Disk2: The last ENABLE disk in class cannot be swapped

a9) Swap the disks.

a10) After swapping the disks, execute the following command.

```
# sdxswap -I -c Class1 -d Disk2
```

a8') Before swapping the disks, execute the following command.



Note

If no error is output in a8), the steps a8'), a9'), and a10') are not required.

```
# sdxdisk -R -c Class1 -d Disk2
```

a9') Swap the disks.

a10') After swapping the disks, execute the following command.

```
# sdxdisk -M -c Class1 -d sdb=Disk2
```

a11) Connect the swapped disk to the group, referring to the group information recorded in a2). If the groups were in a hierarchical structure, connect the groups in an ascending order.

```
# sdxdisk -C -c Class1 -g Group2 -d Disk2  
# sdxgroup -C -c Class1 -h Group1 -l Group2 -a type=stripe,width=32
```

a12) Create volumes with the sdxvolume command again. For the -s option, use the size recorded in a3), in this example, 32768 and 4128768.

```
# sdxvolume -M -c Class1 -g Group1 -v Volume1 -s 32768 -a pslice=off  
# sdxvolume -M -c Class1 -g Group1 -v Volume2 -s 4128768 -a pslice=off
```

a13) Finally, restore the backup data to Volume1 and Volume2.

b. When the error cause is a non-disk component failure or defect

Shut down the system once, recover the failed component, and then reboot the system. Slice data is valid and there is no need to restore the data.

F.1.4 Class Status Abnormality

If the class status is one of the following statuses, take action as indicated for the relevant situation.

- Class becomes closed status during operation.

(1) Class becomes closed status during operation.

Explanation

The class becomes closed when the number of configuration databases which store information on object configuration and object status within a class is insufficient, or when the communication error between nodes occurs in a cluster environment.

All objects within a closed class are inaccessible.

An Insufficient number of configuration databases will occur under the following conditions:

1. When there are no disks that can be accessed normally, if there are two or less disks in ENABLE status.
2. When there are one or less disks that can be accessed normally, if there are three to five disks in ENABLE status.
3. When there are two or less disks that can be accessed normally, if there are more than six disks in ENABLE status.

However, in the event of root class, the class will not be closed unless there are no accessible disks.

GDS configuration databases cannot be stored in BCV devices and target (R2) devices since the devices are overwritten by data in copy source disks. Therefore, GDS does not regard BCV devices and target (R2) devices as "disks that can be accessed normally" described in the above conditions.

Resolution

1) You can check whether or not a class was closed during operation as follows. Do not reboot the system or restart sdxservd daemon, as it will make the checking impossible.

# /etc/opt/FJSVsdx/bin/sdxdcdown							
CLASS	DOWN	REASON	NDK	NEN	NDB	NLDB	DEVNAM
Class1	no	-	10	10	8	0	sda:sdb:sd:c:sdd:sde:sdf:sdg:sdh
Class2	yes	Comm	10	10	8	0	sdi:sdj:sdk:sd:l:sdm:sdn:sdo:sdp
Class3	yes	FewDB	10	10	1	7	sdq
Class4	yes	NoDB	10	10	0	8	-

In this example, Class2, Class3, and Class4 with "yes" in the DOWN field are closed. The cause shown in the REASON field are as follows.

(Cause 1)

Comm Communication failure between nodes.

(Cause 2)

FewDB Insufficient number of valid configuration databases.

(Cause 3)

NoDB No valid configuration database.

2) Depending on specific causes, recovery may be difficult.

First, collect the investigation material.

For information on how to collect the investigation material, see "F.2 Collecting Investigation Material."

Resolutions are described for the following two cases:

- a. Closed due to a communication error
- b. Closed due to an insufficient number of configuration databases

3a) In the event of (Cause 1), contact your local customer support.

3b) In the event of (Cause 2) or (Cause 3), all (or the majority) of the disks registered with class have abnormalities.

You can check the disks registered with class as follows.

# sdxinfo -D -c Class3								
OBJ	NAME	TYPE	CLASS	GROUP	DEVNAM	DEVBLKS	DEVCONNECT	STATUS
disk	Disk31	mirror	Class3	Group1	sda	8847360	*	ENABLE
disk	Disk32	mirror	Class3	Group1	sdb	8847360	*	ENABLE
disk	Disk33	mirror	Class3	Group2	sde	8847360	*	ENABLE
disk	Disk34	mirror	Class3	Group2	sdf	8847360	*	ENABLE
disk	Disk35	mirror	Class3	Group3	sdc	17793024	*	ENABLE
disk	Disk36	mirror	Class3	Group3	sdg	17793024	*	ENABLE
disk	Disk37	mirror	Class3	Group4	sdd	17793024	*	ENABLE
disk	Disk38	mirror	Class3	Group4	sdh	17793024	*	ENABLE
disk	Disk39	spare	Class3	Group1	sdr	17793024	*	ENABLE
disk	Disk40	spare	Class3	*	sds	17727488	*	ENABLE

In this example, ten disks from Disk31 to Disk40 are registered with Class3. Physical disk names are shown in the DEVNAM field. Identify the cause of abnormality with these physical disks by referring to disk driver log messages.

The cause of abnormality could be either of the following:

(Failure 1)

Failed or defective non-disk component.

(Failure 2)

Failed disk component.

4b) In the event of (Failure 1), recover the failed or defective non-disk component (such as I/O adapter, I/O cable, I/O controller, power supply, and fan).

5b) For a local class or a shared class, execute the sdxfix command to restore the class status.

```
# sdxfix -C -c Class3
SDX:sdxfix: INFO: Class3: class recovery completed successfully
```

- If the sdxfix command ends normally, skip steps 6b) through 9b) and go on to step 10b).
- If the sdxfix command does not end normally, go on to step 6b).
- For the root class, go on to step 6b).

6b) Open the GDS configuration parameter file with an editor.

```
# vim /etc/opt/FJSVsdx/sdx.cf
```

Add the following one line in the end of the file.

```
SDX_DB_FAIL_NUM=0
```

7b) Reboot the system.

8b) Confirm that objects within the class are accessible.

```
# sdxinfo -c Class3
```

If nothing is displayed, recovery was unsuccessful. You will have to contact your local customer support. If information is displayed normally, proceed with the following procedures.

9b) In the event of (Failure 2), where a disk component has failed, follow the procedures in "5.3.4 Disk Swap," or "D.8 sdxswap - Swap disk," and swap the disks.

10b) After completing the recovery for both (Failure 1) and (Failure 2), check the number of valid configuration databases as described below.

# /etc/opt/FJSVsdx/bin/sdxdcdown										
CLASS	DOWN	REASON	NDK	NEN	NDB	NLDB	DEVNAM			
-----	----	-----	---	---	---	---	-----			
Class1	no	-	10	10	8	0	sda:sdb:sd:c:sdd:sde:sdf:sdg:sdh			
Class2	no	-	10	10	8	0	sdi:sdj:sdk:sd:l:sdm:sdn:sdo:sdp			
Class3	no	-	10	10	8	0	sdq:sdr:sds:sd:t:sdu:sdv:sdw:sdx			
Class4	no	-	10	10	8	0	sdaa:sdab:sdac:sdad:sdae:sdaf:sdag:sdah			

NLDB field gives the insufficient number of configuration databases. If this value is "0," the problem is resolved. If this value is "1" or more, there are still disks that have not been recovered. In the above example, all NLDB fields display "0," indicating the successful recovery. When step 6b) was not performed, the following procedures are not required.

11b) Open the GDS configuration parameter file with an editor.

```
# vim /etc/opt/FJSVsdx/sdx.cf
```

Remove the following one line added in step 6b).

SDX_DB_FAIL_NUM=0

12b) Reboot the system.

If you cannot perform the recovery with the described procedures, contact your local customer support.

F.1.5 System Disk Abnormality [Itanium]

The following description shows resolutions to problems relevant to system disks on which the file systems / (root), /usr, and /var operates.

In one of the following circumstances, take action as indicated for the relevant situation.

- Create backup of root file system.
- System can be booted, but the system disk data is invalid.
- System cannot be booted. (Failure of original boot disk)
- System cannot be booted. (Boot disk data damage)

(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

See "6.1.2 Backing Up."

(2) System can be booted, but the system disk data is invalid.

Explanation

For some reason, the system disk data is invalid. You must restore data from the backup data created in advance.

Resolution

See "6.1.3 Restoring (When the System Can Be Booted)."

(3) System cannot be booted. (Failure of original boot disk)

Explanation

Even when you have the system disk mirrored, you may experience unsuccessful booting. For example, you may encounter such a problem when the file accessed during the booting process in the boot disk is physically damaged, or when a disk is swapped improperly. If the console message of the unsuccessful boot disk indicates such is the case, try booting from the other boot disk, i.e. the mirror disk.

Resolution

1) From boot disk units displayed in the EFI boot manager's boot option selector window, select a mirror disk and boot.
Among the SDX disk names, the top one is the original boot disk unit name and the others are the mirror disk names. When mirroring is configured with three-way or more multiplexing, multiple mirror disks exist. In the following example, Root1 is the original boot disk name and Root2 and Root3 are the mirror disk names.

```
EFI Boot Manager ver 1.10

Please select a boot option

Root1
Root2
Root3
...

Use ↑ and ↓ to change option(s). Use Enter to select an option
```

2) If booting is successful, find out what caused the original disk to fail, and recover the disk by swapping disks.



See

For details on disk swapping, see "D.8 sdxswap - Swap disk" and resolutions described in "5.3.4 Disk Swap."

When the procedures above cannot resolve the problem, causes other than data errors in the original boot disk unit may exist, or data errors may also exist in the mirror disk.
If it would appear that data errors also exist in the mirror disk, see "(4) System cannot be booted. (Boot disk data damage)."

(4) System cannot be booted. (Boot disk data damage)

Explanation

For some reason, the system disk data is invalid. You must restore data from the backup data created in advance. If restoring data from the mirror disk was unsuccessful, follow the procedures shown below.

Resolution

See "6.1.4 Restoring (When the System Cannot Be Booted)."

F.1.6 GDS Management View Abnormality

The following description shows resolutions to problems relevant to Web-Based Admin View and GDS Management View.

In one of the following circumstances, take action as indicated for the relevant situation.

- Web-Based Admin View abnormality
- GDS Management View cannot be started.
- Disk information is not displayed on GDS Management View. Or, the disk sizes for disk units are not displayed correctly.

(1) Web-Based Admin View abnormality

Explanation

The following problems may occur in Web-Based Admin View.

- a. Web-Based Admin View top menu is not displayed. Or, even if it is, the screen is not displayed properly, or messages such as 0001 to 0050 are displayed.
- b. The screen does not change even after confirming 0002, 0003, or 0007 message. Or, the browser freezes when trying to reconnect to a previously-connected management server.
- c. The following message may be displayed on the management server or the Java console screen of the monitor node.
FJSVwvbs:OutOfMemoryError
- d. When trying to access the management server from browser, message 0005 appears and the top menu is not displayed, or the message itself may not appear.
- e. Web-Based Admin View stopped operating after changing the IP address of the public LAN.
- f. When using the PRIMECLUSTER Operation Management View, existing cluster nodes are not displayed on the each product's node list screen.
- g. The screen does not change even after confirming the dialog message.
- h. The following message is displayed on the browser status bar, and top menu is not displayed.
Applet com.fujitsu.webview.base.WebViewMain
error:java.lang.NoClassDefFoundError: com/sun/java/swing/JApplet
Or, the Operation Management View for each product can not be started.
- i. It takes more than a few minutes before the Web-Based Admin View top menu is displayed. Or, even if it is, it freezes or message "0001 Cannot call Management server" may be displayed.
- j. When using a Java Plug-in on a client browser, and using a mouse on the Web-Based Admin View top menu or Operation Management View for each product, the screen has problems such as displaying small flecks in the area where the mouse has been used.

- k. When starting or restarting the Web-Based Admin View, or rebooting the node, the message "FJSVwvbs:webview.cnf abnormal" or "Node process abnormal end." are displayed, and Web-Based Admin View does not start.
- l. When starting or restarting the Web-Based Admin View, message "WARNING:unrelated 'httpip' property." is displayed.
- m. After installing the manual package, the title menu for the manual is not displayed. Or, even after removing the manual package, the title menu for the manual is still displayed.
- n. When trying to start the menu for products that are the target of management from the Web-Based Admin View top menu, the message "0023 This application must be started at the following URL" is displayed.
- o. Start button for products that are the target of management have disappeared from the Web-Based Admin View top menu.
- p. Even after disconnecting from the management server, messages 0002, 0003 or 0007 are not displayed on the client.
- q. When trying to access the management server from the browser, message 0006 appears, and the top menu is not displayed.
- r. When entering the user ID and the password for authentication, error message 0016 appears, and authentication fails.

Resolution

See "Web-Based Admin View Operation Guide" and follow the troubleshooting procedures.

(2) GDS Management View cannot be started.

Resolution

- a. When GDS is not installed.
Install GDS.
- b. When GDS is installed.
Check for the following possible causes:
 - Wrong URL specified.
 - Wrong method caused due to specifying host name.
 - Environment cannot use Swing.
 - Different browser type and level than recommended.

(3) Disk information is not displayed on GDS Management View. Or, the disk sizes for disk units are not displayed correctly.

Explanation

GDS does not recognize disks properly. Possible causes are as follows.

- The power of the disk unit was turned on after the system was booted.
- The disk unit became unavailable for some kind of problem during system operation.

Resolution

Perform the following steps 1) through 2).

1. Use the parted(8) command and check whether the OS recognizes the disk unit. If it does not recognize, identify the cause based on log messages of disk drivers stored in the /var/log/messages file and such.
2. Execute [Update Physical Disk Information] in the [Operation] menu.

F.1.7 Proxy Object Abnormality

For proxy object related errors, in one of the following circumstances, take action as indicated for the relevant situation.

- F.1.7 The Advanced Copy function cannot be used in master-proxy copying.

(1) The Advanced Copy function cannot be used in master-proxy copying.

Explanation

The method handling copying from a master to a proxy can be checked in the CPTYPE field that is displayed with the `sdxinfo` command. When the CPTYPE field displays `soft`, instead of the Advanced Copy function of Fujitsu Storage Systems ETERNUS, the GDS soft copy function is performing the copy operation.

```
# sdxinfo -S -e long -o pv1
OBJ  NAME CLASS GROUP DISK VOLUME JRM  MODE  STATUS COPY CURBLKS  COPYBLKS DLY  CPTYPE CPSOURCE
-----
slice *   Class1 Group1 Disk1pv1  *   *    COPY  run    93952 10027008 0   soft  *
```

The following ten possible causes can create a situation where the Advanced Copy function is not used.

(Cause a)

The disk array unit's Advanced Copy mechanism was not configured properly.

(Cause b)

When a disk constituting the master or the proxy was registered with a class, the disk array unit's Advanced Copy mechanism was not configured properly.

(Cause c)

A disk constituting the master or the proxy was registered with a class in a situation where GDS Snapshot was not installed.

(Cause d)

After GDS Snapshot was installed, the node was not rebooted.

(Cause e)

EC or REC sessions exist between the master and another proxy and OPC or ROPC cannot be used.

(Cause f)

Multiple proxy volumes are related to a master volume, and the number of proxy volumes with EC or REC session settings has reached the upper limit (16 volumes).

(Cause g)

The mirroring configuration of proxy volumes is two-way or more multiplex mirroring.

(Cause h)

The group to which master volumes or proxy volumes belong has a lower level group connected.

(Cause i)

The disk array unit to which the master or the proxy belongs does not support the Advanced Copy function.

(Cause j)

An error occurred in the disk array unit or the Fibre Channel route.

For details, see "A.2.16 Using the Advanced Copy Function in a Proxy Configuration."

Resolution

Check the disk array settings and the GDS object configuration and identify the cause, and take the following actions according to the cause.

- In the event of (Cause a), follow the procedures below.
 - 1) Configure the disk array unit's Advance Copy mechanism.
 - 2) Back up master and proxy volume data if necessary.
 - 3) Delete the disks constituting the master and the proxy from their classes.
 - 4) Register the disks deleted from the classes in step 3) back to the classes.
 - 5) Create master and proxy volumes and restore the data backed up in step 2) as needed.

- In the event of (Cause b) or (Cause c), follow the procedures below.
 - 1) Back up master and proxy volume data if necessary.
 - 2) Among the disks constituting the master and the proxy, delete disks registered with classes before configuring the Advance Copy mechanism or installing GDS Snapshot from their classes.
 - 3) Register the disks deleted from the classes in step 4) back to the classes.
 - 4) Create master and proxy volumes and restore the data backed up in step 1) as needed.

- In the event of (Cause d), follow the procedures below.
 - 1) Reboot the node.
 - 2) Back up master and proxy volume data if necessary.
 - 3) Delete the disks constituting the master and the proxy from their classes.
 - 4) Register the disks deleted from the classes in step 3) back to the classes.
 - 5) Create master and proxy volumes and restore the data backed up in step 2) as needed.

- In the event of (Cause e) or (Cause f), take the following action (Resolution 1) or (Resolution 2).

(Resolution 1)

 - 1) Part a proxy with EC or REC session settings from the target master.
 - 2) Cancel EC or REC sessions of step 1) using the sdxproxy Cancel command.

(Resolution 2)

 - 1) Cancel the relationship between the target master and a proxy with EC or REC session settings.

- In the event of (Cause g) to (Cause i), select another object, or change the object configuration.

- In the event of (Cause j), recover the disk array unit error or the Fibre Channel route error.

F.1.8 EMC Symmetrix Abnormality [Linux2.6]

For EMC Symmetrix related errors, in one of the following circumstances, take action as indicated for the relevant situation.

- SRDF is not used in master-to- proxy synchronization copying.
- The joined proxy volume is in INVALID status.

(1) SRDF is not used in master-to- proxy synchronization copying.

Explanation

The method handling master-to-proxy copying can be checked in the CPTYPE field that is displayed with the `sdxinfo` command. When the CPTYPE field displays `soft` without displaying SRDF, instead of SRDF, the GDS soft copy function is performing the copy operation.

# sdxinfo -S -e long -o pv1														
OBJ	NAME	CLASS	GROUP	DISK	VOLUME	JRM	MODE	STATUS	COPY	CURBLKS	COPYBLKS	DLY	CPTYPE	CPSOURCE
slice	*	Class1	Group1	Disk1	pv1	*	*	COPY	run	93952	10027008	0	soft	*

The following five possible causes can create a situation where SRDF is not used.

(Cause a)

The configuration of the master or proxy group is inappropriate.

The source (R1) and target (R2) devices as the SRDF pair must be connected with the master and proxy groups respectively. With the proxy group, the target (R2) device only must be connected.

(Cause b)

A pair of volumes, instead of a pair of groups, was related as the master and the proxy.

(Cause c)

When joining the master and the proxy groups, the specific was set to create proxy volumes to the proxy group with the physical slice attributes that are different from those of the master volumes.

(Cause d)

Master data was restored using the proxy with the `sdxproxy RejoinRestore` command.

When the `sdxproxy RejoinRestore` command performs master restoration, the SRDF pair is cancelled and SRDF becomes no longer available.

(Cause e)

With the `sdxproxy Join` command or the `sdxproxy Rejoin` command, the `-e` option was specified to use the soft copy function.

When the `-e softcopy` option is specified with the `sdxproxy Join` command, SRDF is not used.

When the `-e softcopy` option is specified with the `sdxproxy Rejoin` command, the SRDF pair is cancelled and SRDF becomes no longer available.

For details, see "A.2.19 Using EMC TimeFinder or EMC SRDF in a Proxy Configuration [Linux2.6]."

Resolution

1) Check whether the possible cause is (Cause a).

1-1) Identify the physical disk names of the disks connected with the master and the proxy groups.

```
# sdxinfo -D -o mg1
OBJ NAME TYPE CLASS GROUP DEVNAM DEVBLKS DEVCONNECT STATUS
-----
disk Disk1 mirror Class1 mg1 emcpoweri 17596416 node1:node2 ENABLE
disk Disk2 mirror Class1 mg1 emcpowerj 17682084 node1:node2 ENABLE

# sdxinfo -D -o pg1
OBJ NAME TYPE CLASS GROUP DEVNAM DEVBLKS DEVCONNECT STATUS
-----
disk Disk3 mirror Class1 pg1 emcpowerp 17596416 node1:node2 ENABLE
```

In this example, the physical disks emcpoweri and emcpowerj have been connected with the master group mg1, and the physical disks emcpowerp has been connected with the proxy group pg1.

If a lower level group is, or multiple disks or lower level groups are, connected with the proxy group, the possible cause is (Cause a).

1-2) Identify the Symmetrix IDs and the Symmetrix device names of the disks connected with the master and the proxy groups using the sympd command provided by SYMCLI.

```
# sympd list
Symmetrix ID: 000183600262

Device Name      Directors      Device
-----
Physical        Sym SA :P DA :IT Config      Attribute Sts      Cap
(MB)
...
/dev/emcpoweri  000 03A:0 02B:C0 RDF1      Grp'd      RW      8592
/dev/emcpowerj  001 03A:0 02B:D0 Unprotected N/Grp'd    RW      8632
/dev/emcpowerk  002 03A:0 01B:D0 RDF2      Grp'd      RW      8632
...

Symmetrix ID: 000282600920

Device Name      Directors      Device
-----
Physical        Sym SA :P DA :IT Config      Attribute Sts      Cap
(MB)
...
/dev/emcpowerp  005 03A:0 02B:C0 RDF2      Grp'd      RW      8592
/dev/emcpowerq  006 03A:0 01B:D0 RDF1      Grp'd      RW      8632
/dev/emcpowerr  007 03A:0 01A:D0 RDF1      Grp'd      RW      8632
...
```

In this example, as to the physical disks emcpoweri and emcpowerj, the Symmetrix ID is 000183600262 and the Symmetrix device names are 000 and 001 respectively. As for the physical disk emcpowerp, the Symmetrix ID is 000282600920 and the Symmetrix device name is 005.



See

For details on the sympd command, see the Symmetrix Manager manual.

1-3) Check the source (R1) device and the target (R2) device as the SRDF pair whether they are connected with the master and the proxy groups respectively using the symrdf command provided by SYMCLI.

The SymDev field and the RDev field of the symrdf list command output display the Symmetrix device names of the source (R1) device and the corresponding target (R2) device respectively.

```
# symrdf list
Symmetrix ID: 000183600262

Local Device View
-----
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).



See

For details on the symrdf command, see the Symmetrix Manager manual.

1-4) If the possible cause is (Cause e), release the master-proxy relationship once, and re-configure properly.

For example, according to the symrdf command results in step 1.3), you will see that the source (R1) device 006 in the Symmetrix frame 000282600920 and the target (R2) device in the Symmetrix frame 000183600262 are also the SRDF pair. The symprd command results in step 1.2) are indicating that the physical disk names of the Symmetrix devices 006 and 002 are emcpowerq and emcpowerk respectively. The proxy configuration associated with SRDF can be created by registering emcpowerq and emcpowerk with a class, connecting them with respective groups, and relating the groups as the master and proxy.

In addition, the registration of the source (R1) device with a device group can be done after registering the disk with a class, but it must be carried out before joining the master and the proxy.

2) Check whether the possible cause is (Cause b).

```
# sdxinfo -GV -e long -o pv1
```

OBJ	NAME	CLASS	DISKS	BLKS	FREEBLKS	SPARE	MASTER	TYPE	WIDTH
group	pg1	Class1	Disk3	17596416	7503872	0	*	mirror	*

OBJ	NAME	TYPE	CLASS	GROUP	DISK	MASTER	PROXY	SKIP	JRM	MODE	CMODE	LOCK	1STBLK	LASTBLK	BLOCKS	STATUS	PSLICE	SNUM
volume	*	mirror	Class1	pg1	*	*	*	*	*	*	*	*	0	65535	65536	PRIVATE	*	*
volume	pv1	mirror	Class1	pg1	*	mv1	*	off	on	ro	*	on	65536	10092543	10027008	STOP	on	1
volume	*	mirror	Class1	pg1	*	*	*	*	*	*	*	*	10092544	17596415	7503872	FREE	*	*

In this example, the MASTER field of the group pg1 to which the proxy volume pv1 belongs displays the asterisk (*) and the MASTER field of pv1 displays the master volume name mv1, and the possible cause is (Cause b).

If the possible cause is (Cause b), the action as follows should be implemented.

2-1) Break the relationship between the master and proxy volumes.

```
# sdxproxy Break -c Class1 -p pv1
```

2-2) Join a pair of groups instead of a pair of volumes as the master and the proxy. In this example, the master group name is mg1.

```
# sdxproxy Join -c Class1 -m mg1 -p pg1 -a mv1=pv1:on
```

3) Check whether the possible cause is (Cause c).

```
# sdxinfo -V -e long -o mg1
```

OBJ	NAME	TYPE	CLASS	GROUP	DISK	MASTER	PROXY	SKIP	JRM	MODE	CMODE	LOCK	1STBLK	LASTBLK	BLOCKS	STATUS	PSLICE	SNUM
volume	*	mirror	Class1	mg1	*	*	*	*	*	*	*	*	0	65535	65536	PRIVATE	*	*
volume	mv1	mirror	Class1	mg1	*	*	*	off	on	rw	rw	off	65536	10092543	10027008	ACTIVE	on	1
volume	*	mirror	Class1	mg1	*	*	*	*	*	*	*	*	0092544	17596415	7503872	FREE	*	*


```
# sdxinfo -V -e long -o pg1
```

OBJ	NAME	TYPE	CLASS	GROUP	DISK	MASTER	PROXY	SKIP	JRM	MODE	CMODE	LOCK	1STBLK	LASTBLK	BLOCKS	STATUS	PSLICE	SNUM
volume	*	mirror	Class1	pg1	*	*	*	*	*	*	*	*	0	65535	65536	PRIVATE	*	*
volume	pv1	mirror	Class1	pg1	*	mv1	*	off	on	ro	*	on	65536	10092543	10027008	STOP	off	1
volume	*	mirror	Class1	pg1	*	*	*	*	*	*	*	*	10092544	17596415	7503872	FREE	*	*

In this example, the attributes of the physical slices of the master volume mv1 and the proxy volume pv1 are on and off respectively, and the possible cause is (Cause c).

If the possible cause is (Cause c), the action as follows should be implemented.

3-1) Break the master-proxy relationship.

```
# sdxproxy Break -c Class1 -p pg1
```

3-2) Relate the master and the proxy again. With the -a option for the sdxproxy Join command, if the physical slice attribute is not specified, the physical slice attribute of a created proxy volume conforms to that of the master volume.

```
# sdxproxy Join -c Class1 -m mg1 -p pg1 -a mv1=pv1:on
```

4) Whether the possible cause is (Cause d) or (Cause e) can be determined by checking the command action history referencing the GDS log file /etc/opt/FJSVsdx/msglog/sdxservd.log.
If the possible cause is (Cause d) or (Cause e), the master and the proxy should be disjoined once and then rejoined.

(2) The joined proxy volume is in INVALID status.

Explanation

TimeFinder and SRDF may stop operating because the exclusive lock on a Symmetrix device is not released due to an event such as a node down. In this event, copying by TimeFinder or SRDF does not operate normally, and the proxy volume status may become INVALID.

The exclusive locks for Symmetrix devices are divided into two types. Whether respective exclusive locks are left not released can be checked using the symdev and symcfg commands provided by SYMCLI.

```
# symdev -lock list  
# symcfg -lock list
```

If an exclusive lock has remained, the lock number is displayed.

Resolution

Release the exclusive lock on the Symmetrix device with the following method.

As to an exclusive lock indicated by the symdev -lock list command, use the symdev command to release it. Substitute the lock number displayed with the symdev -lock list command for *number*.

```
# symdev -lock number release
```

As to an exclusive lock indicated by the symcfg -lock list command, use the symcfg command to release it. Substitute the lock number displayed with the symcfg -lock list command for *number*.

```
# symcfg -lockn number release
```



See

For details on the symdev command and the symcfg command, see the Symmetrix Manager manual.

F.1.9 Cluster System Related Error

For cluster system related errors, in one of the following circumstances, take action as indicated for the relevant situation.

- The error message "ERROR: class: cannot operate in cluster environment, ..." is output, and the operation cannot be conducted on the class class. [Linux2.6]
- The PRIMECLUSTER CF clinitreset(8) command ends abnormally outputting an error message # 6675.
- Cluster applications become "Inconsistent".
- The GFS Shared File System is not mounted on node startup.

(1) The error message "ERROR: *class*: cannot operate in cluster environment, ..." is output, and the operation cannot be conducted on the class *class*. [Linux2.6]

Explanation

The local class created when cluster control facility was inactive cannot directly be used in a cluster system. When the cluster control facility is activated, the following message is output to the system log and the GDS daemon log file, and the local class becomes nonoperational.

ERROR: *class*: cannot operate in cluster environment, created when cluster control facility not ready

This error message will be output when:

- The cluster initial configuration was executed after the local class *class* had been created on a node on which that configuration was incomplete.
- The local class *class* was created in single user mode.
- The single node on which the local class *class* was created was changed over to a cluster system.

Resolution

Make the local class available in a cluster system by method a) or method b) as follows. Generally method a) should be used, but method b) should be used in order to prevent the volume data from being backed up and restored.

Method a) Re-creating the local class in the cluster system:

- 1) Activate the node in single user mode.
- 2) Back up volume data if necessary.
- 3) Delete the class.
- 4) Re-activate the cluster control facility on the node booted in multi-user mode.
- 5) Re-create the class and volumes deleted in step 3).
- 6) Restore the volume data backed up in step 2) as needed.

Method b) Converting the local class to one for a cluster system:

Convert the local class to a class for a cluster system taking the following procedures. The following illustrates the procedures when the class name is Class1.

- 1) Activate the node in single user mode.

- 2) Stop the GDS management daemon, `sdxservd`.

```
# /etc/opt/FJSVsdx/bin/sdx_stop -S
sfdsk: received shutdown request
sfdsk: volume status log updated successfully, class=0x40000004
#
```

Confirm that the `sdxservd` daemon was stopped (information on `sdxservd` daemon processes is not displayed) in the following manner.

```
# ps -e | grep sdxservd
#
```

3) Back up the configuration database for the local class Class1.

```
# rm -rf /var/opt/FJSVsdx/backup/DB/Class1
# /etc/opt/FJSVsdx/bin/sdxcitrandb -B -c Class1
sdxsavedb: INFO: /dev/sda1: backup succeeded
sdxsavedb: INFO: /dev/sdb1: backup succeeded
sdxsavedb: INFO: Class1: backup succeeded

# cd /var/opt/FJSVsdx/backup/DB/Class1
# ls -l
-rw-r--r--  1 root    other   14164992 May  6 09:00 sda1
-rw-r--r--  1 root    other   14164992 May  6 09:00 sdb1
```



Note

- Verify that there is free space equal to or larger than 150 [MB] under /var/opt/FJSVsdx/backup/DB, and if it is insufficient expand it.
- If an error occurs, you should not go on to the following procedures but should re-create the local class Class1 according to method a).

4) Convert the configuration database for the local class Class1 to that for a cluster system.

```
# /etc/opt/FJSVsdx/bin/sdxcitrandb -C -c Class1
sdxcvertdb: INFO: /dev/sda1: conversion succeeded
sdxcvertdb: INFO: /dev/sdb1: conversion succeeded
sdxcvertdb: INFO: Class1: conversion succeeded
```



Note

If an error occurs, you should restore the configuration database taking steps from 7-3) onward and re-create the local class Class1 according to method a).

5) Re-activate the cluster control facility by making the node be in multi-user mode.



Note

The following messages are output during shutdown, but there are no problems.

```
SDX:sdxcshutdown: INFO: waiting for a response from sdxcservd daemon...
SDX:sdxcshutdown: ERROR: connection timeout
```

6) Verify that the configuration database for the local class Class1 was converted successfully.

```
# sdxcinfo -C -c Class1
OBJ  NAME  TYPE  SCOPE  SPARE
----  ----  ----  ----  ----
class Class1 local  Node1  0
```

Confirm that the node identifier is displayed properly in the SCOPE field. If it is displayed properly, the process is finished.



Note

If the SCOPE field statement is improper, it means that the Class1 configuration database was not converted successfully. If that happens, you should restore the configuration database taking steps from 7-1) onward and re-create the local class Class1 according to method a).

7) Restore the configuration database backed up in step 3) when an error occurred in step 4) or 6).

7-1) Activate the node in the single user mode.

7-2) Stop the GDS management daemon, sdxservd.

```
# /etc/opt/FJSVsdx/bin/sdx_stop -S
sfdsk: received shutdown request
sfdsk: volume status log updated successfully, class=0x40000004
#
```

Confirm that the sdxservd daemon was stopped (information on sdxservd daemon processes is not displayed) in the following manner.

```
# ps -e | grep sdxservd
#
```

7-3) Restore the configuration database for the local class Class1.

```
# /etc/opt/FJSVsdx/bin/sdxcltrandb -R -c Class1
sdxrestoredb: INFO: /dev/sda1: restore succeeded
sdxrestoredb: INFO: /dev/sdb1: restore succeeded
sdxrestoredb: INFO: Class1: restore succeeded
```

7-4) Re-activate the node in single user mode.



Note

The following messages are output during shutdown, but there are no problems.

```
SDX:sdxshutdown: INFO: waiting for a response from sdxservd daemon...
SDX:sdxshutdown: ERROR: connection timeout
```

7-5) Verify that the configuration database for the local class Class1 was restored normally.

```
# sdxinfo -C -c Class1
OBJ NAME TYPE SCOPE SPARE
-----
class Class1 local Node1 0
```

Confirm that the node identifier is displayed properly in the SCOPE field. If it is displayed properly, the restoration is finished.

(2) The PRIMECLUSTER CF clinitreset(8) command ends abnormally outputting an error message # 6675.

Explanation

When a class exists in a cluster system, initializing the PRIMECLUSTER resource database with the PRIMECLUSTER CF clinitreset command results in that the clinitreset command fails outputting the following error message.

```
FJSVcluster: ERROR: clinitreset: 6675: Cannot run this command
because Global Disk Services has already been set up.
```

When a node containing a shadow class is rebooted because of an event such as shutdown and panic, the shadow class is deleted, but the `/dev/sfdsk/Class Name` directory is not deleted. If the clinitreset command is executed here, the command also fails outputting the error message as above.

Resolution

1. On all nodes in the cluster system, view the configuration of objects and delete a class if any exists. If a class is deleted, volume data will be lost. If necessary, back up volume data in advance.



See

- For using GDS Management View, see "5.5 Removals."
 - For using commands, see "Appendix D Command Reference."
2. On all nodes in the cluster system, check whether a class directory exists in the /dev/sfdsk directory, and delete a class directory if any exists. The following shows an example when a directory of class Class1 exists.
_adm and _diag are special files used by GDS and cannot be deleted.

```
# cd /dev/sfdsk
# ls
_adm  _diag  Class1
# rm -rf Class1
```

(3) Cluster applications become "Inconsistent".

Explanation

If a shared class is not to be used as an RMS resource, volumes included in the class are started on node startup. If a cluster application that uses those volumes are started there, the cluster application becomes "Inconsistent" because the volumes are already active. By default, classes are not to be used as RMS resources. Classes can be made available as RMS resources by:

- Specifying them and using the hvgdsetup -a command

Resolution

Make the shared class available as an RMS resource with one of the following methods. After performing the procedures, restart the cluster application.

- Execute the following command.

```
# /usr/opt/reliant/bin/hvgdsetup -a Class Name
...
Do you want to continue with these processes ? y
```

(4) The GFS Shared File System is not mounted on node startup.

Explanation

If a shared class is to be used as an RMS resource, volumes included in the class are not started on node startup. Therefore, the GFS Shared File System on those volumes is not mounted on node startup. By default, classes are not to be used as RMS resources, but they are made available as RMS resources by:

- Specifying them and using the hvgdsetup -a command

Resolution

Take one of the following actions.

- a) When using the shared class as an RMS resource, do not create the GFS Shared File System on volumes in the class, but create it on volumes in a difference class.
- b) When not using the shared class as an RMS resource, make the class unavailable as an RMS resource again with one of the following methods. After performing the procedures, reboot the system.

- Execute the following command.

```
# /usr/opt/reliant/bin/hvgsdsetup -d Class Name
...
Do you want to continue with these processes ? y
...
Do you need to start volumes in the specified disk class ? n
```

F.1.10 Shadow Object Errors

For shadow object related errors, in one of the following circumstances, take action as indicated for the relevant situation.

- A shadow volume became inaccessible.

(1) A shadow volume became inaccessible.

Explanation

The configuration information of a shadow volume is not saved on the private slice, but managed in the memory. For this reason, the shadow volume configuration is cleared when the node on which the shadow volume is defined is rebooted, and thus the shadow volume becomes inaccessible.

Resolution

Delete the device special file according to the procedure described in "Rebooting a Node" in "A.2.26 Shadow Volume," and re-create a shadow volume.

F.1.11 Disk Unit Error



Note Identifying a Failed Disk Unit

GDS Management View shows disks with I/O errors by indicating them with red disk icons. In addition, if an I/O error occurs, 1 (one) is given as a value in the E field displayed with the `sdxinfo -e` long command.

However, for hardware-related errors, identify the failed parts based on log messages of disk drivers stored in the `/var/log/messages` file and such.

If GDS Management View and the `sdxinfo` command indicate the invalid status of slice data (invalid or INVALID), the slice data is invalid. It does not necessarily indicate that that slice has a hardware error.

For example, if an I/O error occurs on the copy source slice due to a hardware error during the synchronization copying of a mirror volume, data on the copy destination slice becomes invalid, but the copy source slice is still indicated as normal (active, stop, ACTIVE, or STOP). For details, see "F.1.1 Slice Status Abnormality."



See

For the methods of swapping disks, see "1.2.3 Hot Swap," "5.3.4 Disk Swap," and "D.8 sdxswap - Swap disk."

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.
- In a cluster system, collect investigation material on all the nodes.

F.2.1 Collecting with pclsnap Command

The pclsnap command, a Linux system information collection tool, is provided in the FJSVpclsnap package. When a problem occurs with GDS, this command allows you to collect investigation material to identify the cause.

How to execute the pclsnap command is as follows.

1. Log in with superuser authority.
2. Execute the pclsnap command as follows.

```
# /opt/FJSVpclsnap/bin/pclsnap -h output
```

or

```
# /opt/FJSVpclsnap/bin/pclsnap -a output
```

output indicates either the special file name (example: /dev/st0) of the output media which stores the investigation material collected with the pclsnap command, or the output file name.

F.2.2 Collecting Initial Investigation Material with sdxsnap.sh Script

Since the amount of information collected with the pclsnap 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 pclsnap command, and e-mail this initial investigation material first.

Execute the sdxsnap.sh script by following the procedures below.

1. Login with the superuser authority.
2. Execute the sdxsnap.sh script.

```
# /etc/opt/FJSVsdx/bin/sdxsnap.sh -s [-d dir]
# cd dir
# tar cvf snap.tar SDXSNAP
```

Use the -d option to save the material in the *dir* directory. If the -d option is not used, the material will be saved in the /tmp directory.

Send us the snap.tar file, which will be used as initial investigation material.

Appendix G Frequently Asked Questions (FAQ)

In this section, frequently asked questions (FAQ) have been compiled.

G.1 Operation Design

Questions regarding system operation design are explained.

- Is operation possible without using GUI (GDS Management View)?
- Is it possible to perform GDS Snapshot operations through the GUI (GDS Management View)?
- Can disks with different performance specifications, such as size and revolution speed, be mirrored?
- Do disks have to have equivalent attributes, such as size and revolution speed when I relate them as master and proxy?
- Can I combine mirroring and striping?
- Do I need to specify certain settings in order to perform a hot swap?
- What is the purpose of single volume?
- Can the messages output by GDS driver, daemon and command be displayed in a language other than English?
- Are there any restrictions on the LUN (Logical Unit) configuration within the disk array when managing the disk array using GDS?
- Are the mirrored disks differentiated as "primary" and "secondary" disks?
- When reading the mirror volume, which one of the slices in the volume will the I/O request be issued to?
- 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?
- Which of the GDS messages should be monitored?
- I want to know GDS resident daemons to monitor resident processes.
- When configuring a 24-hour operation system, would files of incremental size such as message log file have any effect?

Q: Is operation possible without using GUI (GDS Management View)?

A

Yes. Two types of operation management interfaces are available: the GUI (Graphical User Interface) and the CLI (Command Line Interface).

Q: Is it possible to perform GDS Snapshot operations through the GUI (GDS Management View)?

A

Proxy object operations are possible. To operate shadow objects, use the CLI (Command Line Interface).

Q: Can disks with different performance specifications, such as size and revolution speed, be mirrored?

A

Yes. However, when mirroring disks of different sizes, you will only be able to use the capacity of the smaller disk. For example, if a 4 GB disk and a 9 GB disk are mirrored, only 4 GB out of the 9 GB disk will be accessible.

Also, when mirroring disks with different specifications such as revolution speed, the read performance becomes unbalanced and the write performance will depend on the slower disk's performance. Therefore, it is recommended to mirror disks with the same specifications.

See "A.3.1 Guidelines for Mirroring" also.

Q: Do disks have to have equivalent attributes, such as size and revolution speed when I relate them as master and proxy?

A

You can relate disks as master and proxy even if they have different attributes, such as size and revolution speed. However, it will result in that the read performance becomes unequal and the write performance depends on the slower disk's performance. Therefore, it is recommended to relate disks with the same attributes wherever possible.
For details see also "A.1.8 Exception to Proxy Configuration."

Q: Can I combine mirroring and striping?

A

You can mirror stripe groups. However, you cannot stripe across multiple mirror groups, using them as stripe columns.

Q: Do I need to specify certain settings in order to perform a hot swap?

A

No. As long as the disk you are using supports hot swap, no special settings are required.
For procedures, see "5.3.4 Disk Swap" and "D.8 sdxswap – Swap disk."

Q: What is the purpose of single volume?

A

GDS is 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 "1.3.2 Centralized Disk Management."

For example, if availability is already realized by using disk arrays, the user will benefit by using the physical disk as a single volume excluded from mirroring, to improve its manageability.

Using GDS to manage the volume as a single volume not only ensures consistent management, but also allows you to include the volume in the mirroring configuration without stopping the service application, when the need arises.

Q: Can the messages output by GDS driver, daemon and command be displayed in a language other than English?

A

No. The messages are in English only.

Q: Are there any restrictions on the LUN (Logical Unit) configuration within the disk array when managing the disk array using GDS?

A

There are no restrictions on LUN RAID level or size.

Q: Are the mirrored disks differentiated as "primary" and "secondary" disks?

A

No. Disks mirrored with GDS are handled in the same way. Both the user and the application do not differentiate between the two.
However, the only exception is the system disk. The system disks may be differentiated as the primary and the secondary disks to initially be booted with EFI. For details, see "F.1.5 System Disk Abnormality [Itanium]."

Q: When reading the mirror volume, which one of the slices in the volume will the I/O request be issued to?

A

When a read is issued to a volume, the last access block number of each disk in the volume will be checked. The slice with the value closest to the read block number will be read. In other words, disk with the shortest seek distance will be read first.

When a write is issued to a volume, it will write to all slices in the volume (except from slice in ACTIVE status) and will return the write results after completing the write request on all slices.

Q: When mirroring a disk array that supports the hot spare function, which is preferable; using the spare disk of the disk array, or that of GDS?

A

Generally, using the spare disk of the disk array is recommended.

Q: Which of the GDS messages should be monitored?

A

Generally, the following messages given in the "GDS Messages" should be monitored.

- PANIC or WARNING level messages and Internal error messages shown in "E.2 Driver Messages"
- HALT, ERROR, or WARNING level messages and Internal error messages shown in "E.3 Daemon Messages"

For details, see "Appendix E GDS Messages."

Q: I want to know GDS resident daemons to monitor resident processes.

A

The following table contains GDS resident daemons and the description.

Name	ps -ef command's CMD	Description
sdxmond	/usr/sbin/sdxmond	Monitors GDS daemons
sdxservd	/usr/sbin/sdxservd	Required daemon
sdxlogd	/usr/sbin/sdxlogd	Required daemon
sdxexd	/usr/sbin/sdxexd	Required daemon
sdxcld	/usr/sbin/sdxcld -W	Cooperates with PRIMECLUSTER CF (Only in a cluster system)
sdxcld	/usr/sbin/sdxcld -W	Cooperates with PRIMECLUSTER CF (Only in a cluster system)
sdxcle	/usr/sbin/sdxcle -f <i>n</i>	Cooperates with PRIMECLUSTER CF (Only in a cluster system) <i>n</i> is a natural number.

If any resident daemon other than sdxmond terminates abnormally, sdxmond automatically restarts the daemon and outputs a message as follows to the GDS log file /var/opt/FJSVsdx/msglog/sdxservd.log.

When succeeded in restart

```
SDX:sdxmond: WARNING: respawned daemon daemon successfully
```

When failed in restart

```
SDX:sdxmond: HALT: failed to respawn daemon daemon, osfunc=osfunc, errno=errno
```

If sdxmond terminates abnormally, the OS init(8) process automatically restarts the daemon. In that event, no message is output.

As described above, the GDS resident daemons are automatically restarted in the event of abnormal termination, and it is not necessary to monitor them. In addition, GDS Snapshot has no resident daemon.



Note

If the Linux kernel version is 2.4, sdxmond is not automatically rebooted.



See

For the meanings of the messages and the resolutions, see "E.3 Daemon Messages."

Q: When configuring a 24-hour operation system, would files of incremental size such as message log file have any effect?

A

Since GDS's message log file is managed so that it will not exceed a certain size, there are no effects that need to be considered.

G.2 Environment Configuration

Questions regarding environment configuration are explained.

- How do I make a backup of a system disk that is managed with GDS? [Itanium]
- How do I make a backup of a disk (other than a system disk) that is managed with GDS?
- I cannot execute the command parted(8) on a physical disk managed by GDS (e.g. sda). How can I solve this problem?

Q: How do I make a backup of a system disk that is managed with GDS? [Itanium]

A

For procedures on how to make backups of file systems /(root), /usr and /var, which cannot be unmounted while the system is operating, see "6.1 Backing Up and Restoring a System Disk [Itanium]" and "6.2 Backing Up and Restoring a System Disk through an Alternative Boot Environment [Itanium]."

Q: How do I make a backup of a disk (other than a system disk) that is managed with GDS?

A

The backup of GDS's logical volume can be made in the same way as a physical slice. However, make sure you make a backup using the logical volume. For example, do not boot from the CD-ROM and make a backup using the physical slice, since correct data may not be saved.

You can also efficiently create a backup with the mirroring function. For details, see "1.3.8 Snapshots by Slice Detachment" "D.5 sdxslice - Slice operations", "5.3.2 Backup" and "6.3 Backing Up and Restoring Local Disks and Shared Disks."

By making use of GDS Snapshot proxy volumes, you can also configure backup administration that minimizes operation stop time. For details, see "1.5 Proxy Volume," "D.14 sdxproxy - Proxy object operations," "D.19 Snapshot Creation Using Command," "5.3.2.2 Backup (by Synchronization)," "5.3.2.3 Backup (by OPC)" and "6.4 Online Backup and Instant Restore through Proxy Volume."

Furthermore, backing up data through a server outside the domain running the service using GDS snapshot shadow volumes can hold the effects on the service to a minimum. For details, see "1.6 Shadow Volume" and "6.6 Backing Up and Restoring through an External Server."

Q: I cannot execute the command parted(8) on a physical disk managed by GDS (e.g. sda). How can I solve this problem?

A

Since access to physical disks is restrained by GDS, you cannot perform this operation. For details, see "A.2.2 Restraining Access to Physical Special File."

If all you need to do with the commands parted(8) is to refer to the disk contents, you may do so by booting from the CD-ROM, and specifying the physical disk. However, you must exercise care since updating the data on the physical disk will damage the mirrored data. For details, see "A.2.3 Booting from a CD-ROM Device."

G.3 Operation

Questions regarding system operation are explained.

- Do I need to exit all applications to perform synchronization copying on the volume? If it is not necessary, what effects should I be aware of?
- How do I change the volume size?
- I cannot boot the system. How do I restore it? [Itanium]
- When parting a proxy volume in a shared class, I want to lock the proxy volume not to start up on certain nodes.
- Is there a possibility that joined proxy volumes are parted automatically?

Q: Do I need to exit all applications to perform synchronization copying on the volume? If it is not necessary, what effects should I be aware of?

A

You do not need to exit the applications, but there could be some affect on the performance. If the application response is seriously affected, you can either suspend the synchronization copying, or adjust the load caused by the copying.

Interrupting the copying:

```
# sdxcopy -I -c class name -v volume name
```

Resuming the copying:

```
# sdxcopy -B -c class name -v volume name
```

Adjusting the copying load:

```
# sdxcopy -P -c class name -v volume name -e delay=delay time
```

Checking the copying process:

```
# sdxinfo -S -c class name -o volume name -e long
```



See

For details, see "D.10 sdxcopy - Synchronization copying operation," and "D.6 sdxinfo - Display object configuration and status information."

Q: How do I change the volume size?

A

Volumes can be expanded with the `sdxvolume -S` command during the service. For details, see "1.3.7 Online Volume Expansion."

The procedure for changing the volume size while suspending access to the volume is as follows.

1. Exit the application that is accessing the volume
2. Backup the volume data as necessary.
3. Remove the volume.
4. Check if you can create a volume with the necessary size in the group in which the volume existed. You can check this as follows.

```
# sdxinfo -V -c class name -o group name
```

The volume which indicates FREE in the STATUS field will be available. The BLOCKS field indicates the volume size. When there is more than one available volume, the maximum size you can create would be the size of the largest available volume.
5. When you cannot secure sufficient size within the group by performing procedure 4, backup a different volume and remove it, or create another group.
6. Specify a new size and create the volume. When creating the volume, specifying the same name as the volume you removed in procedure 3 will reduce the amount of information you have to input when changing the device definition information in the following procedure.
7. Update the device definition information such as `/etc/fstab`, as necessary.
8. Restore the backup data you collected in procedure 2.

Q: I cannot boot the system. How do I restore it? [Itanium]

A

See "F.1.5 System Disk Abnormality [Itanium]."

Q: When parting a proxy volume in a shared class, I want to lock the proxy volume not to start up on certain nodes.

A

Before parting a proxy volume, on nodes where you want to restrain the proxy volume from startup, turn on the proxy volume lock mode with the `sdxattr -V` command. For details on the lock mode, see "Lock Volume" and "D.7 sdxattr - Set objects attributes."

Q: Is there a possibility that joined proxy volumes are parted automatically?

A

In the situation where the master volume has only one ACTIVE slice, if an I/O error occurs in that slice, the proxy volume in STOP or INVALID status that is joined with the master volume through soft copy is parted automatically.

Appendix H Shared Disk Unit Resource Registration

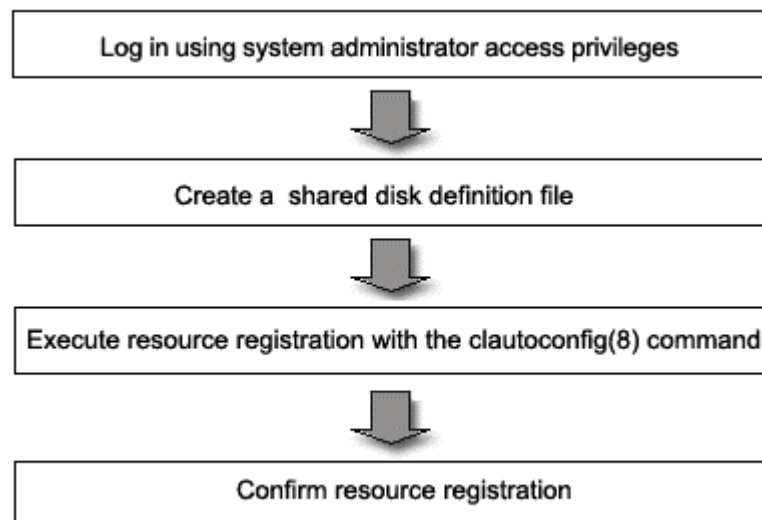
H.1 What Is Shared Disk Resource Registration?

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

H.2 The Flow of Shared Disk Resource Registration

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

Figure: The Flow of Shared Disk Unit Resource Registration



H.3 Preconditions and Important Points

- Before performing resource registration, make sure the following particulars.
 - 1) CF, CIP, and resource database settings are complete on all nodes.



See

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

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

[Recovery procedure]

- 1) Resetting the resource database
Reset the resource database using the clinitreset(8) command.



See

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

- 2) Rebooting the node
Reboot the node.



See

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

H.4 Shared Disk Unit Resource Registration Procedure

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

1. Logging in with system administrator access privileges

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

2. Creating a shared disk definition file

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

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

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

Specify values as follows for "Resource name", "Device name", and "Node identification name".

Resource key name

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

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

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

Device name

For a device name, specify the device name of a shared disk unit. The device name is specified in "/dev/device" format.

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

Node identification name

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



See

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

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

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

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

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

3. Resource registration with the clautoconfig(8) command

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

[Format]

```
/etc/opt/FJSVcluster/bin/clautoconfig -f Full path of the shared disk definition file
```

[Example]

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



Note

If the clautoconfig(8) command terminates abnormally, take corrective actions as indicated in a displayed message. For clautoconfig(8) command messages, see "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide."

4. Confirming resource registration

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

[Example]

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

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

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



Note

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

[Recovery procedure]

- 1) Resetting the resource database
Reset the resource database using the clinitreset(8) command.



See

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

- 2) Rebooting the node
Reboot the node.

H.5 Command Reference

This section explains the following commands.

- **clautoconfig(8)**
- **cldelrsc(8)**

H.5.1 **clautoconfig(8)-Execute resource registration**

SYNOPSIS

```
/etc/opt/FJSVcluster/bin/clautoconfig -f diskfile  
/etc/opt/FJSVcluster/bin/clautoconfig -r -n
```

DESCRIPTION

This command enters information of the following hardware connected to a cluster system in the resource database.

- **Shared disk devices**
Information of shared disk devices specified in the shared disk definition file is entered in the resource database.
- **Network devices**
Network devices are automatically detected and entered in the resource database.

You can execute this command on any single node in the cluster system with the Administrator's authority.

While this command is running, you may not execute this command on the given node or another node.

OPTIONS

-f *diskfile*

Registers the shared disk unit in the resource database. Specify the full path to a predefined shared disk definition file for *diskfile*.

-n

Registers the network device in the resource database. The detected network device is active. The inactive devices are not detected. Execute the `ifconfig(8)` command to confirm if the network device is running.

-r

Registers the resource automatically.

EXIT STATUS

0:

Normal termination.

Other than 0:

Abnormal termination.

NOTES

You need to check the following before executing the command:

- CF, CIP, and resource database settings are complete on all nodes
- All nodes and disk units are connected.
- All nodes are up in multi-user mode.

H.5.2 cldelrsc(8)-Delete resource

SYNOPSIS

`/etc/opt/FJSVcluster/bin/cldelrsc -r rid`

DESCRIPTION

The `"/etc/opt/FJSVcluster/bin/cldelrsc -r rid"` command deletes resources from the resource database. Resources of the following resource classes can be deleted.

Resource class name	Meaning
DISK	Resource of the shared disk unit
SHD_DISK	Shared resource of the shared disk unit
Ethernet	Resource of the network device

Use the `clgettree(1)` command to check the resources to be deleted. Among shared disk unit resources, if resources in the DISK class are deleted and shared resources in the SHD_DISK class remain, it is also necessary to delete the shared resources. Execute this command on any single node in the cluster system using system administrator access privileges.

OPTIONS

`-r rid`

Specify *rid* for the resource ID of the resource to be deleted. Check the resource ID executing the `clgettree(1)` command.

EXIT STATUS

0:

Normal termination.

Other than 0:

Abnormal termination.

NOTES

If a shared disk of which resources are to be deleted is registered with a GDS class, delete the shared disk from the GDS class, and then delete the resources.

EXAMPLES

The following example shows the procedure for deleting resources of a shared disk unit "/dev/sdc" shared between node1 and node2.

1. Check the resource IDs of /dev/sdc resources using the `clgettree(1)` command. In the following example, the /dev/sdc resource ID on node1 is 18, and the /dev/sdc resource ID on node2 is 19.

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

2. Delete the resources of the shared disk unit "/dev/sdc" using the `cldelrsc(8)` command.

```
# /etc/opt/FJSVcluster/bin/cldelrsc -r 18
# /etc/opt/FJSVcluster/bin/cldelrsc -r 19
```


Glossary

[A]

active disk

The operating disk between disks that are connected to a switch group.

active slice

A switch slice.

Advanced Copy function

The copy function of Fujitsu Storage Systems ETERNUS. A generic term for EC (Equivalent Copy) and OPC (One Point Copy). For details, see the Fujitsu Storage Systems ETERNUS manual.

alternative root volume [Itanium]

The alternative volume of the root volume.

alternative volume [Itanium]

A volume that is directly or indirectly related to the master or the proxy. The alternative volume of a master volume is the proxy volume. The alternative volume of a proxy volume is the master volume or another proxy volume related to the master volume.

array disk

See "disk array."

attach slice

To add the detached slice to mirror volume making it accessible again as a component of the mirror volume.

available size of a disk

The capacity on an SDX disk that is available for creating volumes. It equals the physical disk size rounded to a cylinder boundary, minus the private slice size.

available size of a group

The capacity that is available for a volume in a group.

[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.

backup object configuration [Linux2.6]

To output the configuration of objects within a class to a configuration file in configuration table format.

block

A unit for disk capacity or data transmission. Usually, one block is 512 bytes. Sector is used as a synonym.

boot disk

A bootable disk with /boot/efi installed.

break

To cancel the master-proxy relationship, and make independent again.

[C]**class**

1. A disk class.
2. The generic term for disk classes and shadow classes.

close down class

To make all objects within a class unavailable for reference.

close down volume

To make the volume unavailable (invalid).

cluster application

A resource categorized as a userApplication used to group resources into a logical collection. For details, see the PRIMECLUSTER manual.

cluster control

A component contained in PRIMECLUSTER responsible for detecting various abnormalities and state transition. For details, see the PRIMECLUSTER manual.

cluster system

A highly available and reliable system with multiple nodes connected via a special network. For details, see the PRIMECLUSTER manual.

concatenate

To connect more than one physical disk. It is helpful when using multiple disks as one virtual mass disk.

concatenation

To concatenate.

concatenation group

A group whose type attribute is "concat." It is also the set of disks that will be concatenated.

configuration database

Database storing the configuration data of the objects managed by GDS. The configuration database is stored on a private slice.

configuration file [Linux2.6]

A text file that contains a configuration table.

configuration table [Linux2.6]

Text data that tabulates the object configuration of a class.

connect disk

To add an SDX disk to a disk group. User can also connect a keep disk, a single disk, or an undefined disk. Spare disk will be connected automatically by GDS when required.

connect group

To add a group to another group.

The former group is called a "lower level group," and the latter, a "higher level group."

For example, when more than one stripe group is connected to a mirror group that is the highest level group, the mirror volume that is created within the highest level group will be a logical device unit that will be both mirrored and striped.

convert configuration table [Linux2.6]

To change a configuration table according to a restoration environment in order to restore the object configuration of the class in the environment that is different from the original.

current root volume [Itanium]

The volume currently running as the root file system.

current volume [Itanium]

The master volume or the proxy volume that is described as a file system or a swap area in the /etc/fstab file.

[D]**detach slice**

To temporarily detach a mirror slice from a mirror volume, making it inaccessible as a component of the mirror volume.

Other than by user command, GDS will automatically detach slices to maintain synchronization when an I/O error of a mirrored slice is detected.

disconnect disk

To remove a disk from a disk group.

disconnect group

To remove the lower level group from a higher level group.

disk

1. An SDX disk.
2. The generic term for SDX disks and shadow disks.

disk array

A disk unit which combines multiple physical disks to be handled as one virtual disk. Spreading the access across multiple physical disks greatly improves performance and reliability. There are several access methods which are referred to as RAID (Redundant Array of Inexpensive Disks) levels. Main RAID level includes RAID 0 (striping), RAID 1 (mirroring), and RAID 5 (striping with parity).

disk class

A collection of SDX objects. Disk classes may be called "classes." Physical disks can be registered with a class and accordingly are managed by GDS. Classes are classified into root classes, local classes, and shared classes.

disk group

A collection of SDX disks and lower level groups. Disk groups may be called "groups." SDX disks and lower groups that belong to the same group become the unit of mirroring, striping, concatenating, or switching according to the type attribute of the group (mirror, stripe, concatenation, switch).

disk switch

To switch the roles of the active and inactive disks.

domain

One cluster system built by use of PRIMECLUSTER.

A server (computer) that does not belong to a cluster system is also referred to as a domain.

[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 functions provided by Fujitsu Storage Systems ETERNUS. It keeps duplicating disk data within a disk unit. By performing detaching control with software, a replicate of the disk will be created. For details, see the Fujitsu Storage Systems ETERNUS manual.

EC session

EC request from the host (server) to Fujitsu Storage Systems ETERNUS. For details, see the Fujitsu Storage Systems ETERNUS manual.

emcpower disk

A virtual device that provides multiple access paths with the emcpower driver. Being able to use multiple paths enables spreading out of the load when accessing the disk, and increases the availability at the time of path failure.

ETERNUS3000 series

Disk array unit manufactured by Fujitsu.
For details see ETERNUS3000 series manuals.

ETERNUS6000 series

Disk array unit manufactured by Fujitsu.
For details see ETERNUS6000 series manuals.

[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.

geometry

Physical properties of a disk such as the numbers and sizes of cylinders, heads, and sectors.

GFS

Abbreviation of Global File Services. It is software that manages a GFS shared file system. For details see "PRIMECLUSTER Global File Services Configuration and Administration Guide."

GFS Shared File System

Shared file system that features simultaneous accessibility that maintains consistency between multiple nodes to which a shared disk unit is connected and enables the continuation of processing on another node even if a certain node fails. For details, see "PRIMECLUSTER Global File Services Configuration and Administration Guide."

GR700 series

Disk array unit manufactured by Fujitsu.
For details see GR700 series manuals.

group

1. A disk group.
2. The generic term for disk groups and shadow groups.

[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.

higher level shadow group

A shadow group to which another shadow group belongs.

hot spare

A feature realizing automatic mirroring recovery while maintaining data accessibility at times of mirror disk failure.

hot swap

A feature allowing users to exchange the faulty physical disk components without stopping the application in use. Also referred to as active exchange.

[I]**inactive disk**

A standby disk between disks that are connected to a switch group.

inactive slice

A standby slice that will be the operating slice when the active disk of the switch group is changed through a disk switch. The standby slice is a gathering of physically continuous blocks in the standby disk that belongs to the switch group."

instant restore

To instantly restore the master's data, by overwriting the master using the parted proxy data. Instant restore can be created by using the `-e` instant option with the `sdxproxy RejoinRestore` command or the `-e` instant option with the `sdxproxy Restore` command.

instant snapshot

To create a snapshot of a master instantly by cooperation with the OPC (One Point Copy) function provided by Fujitsu Storage Systems ETERNUS. Also it is a created snapshot. Snapshots can be created by using the `-e` instant option with the `sdxproxy Part` command.

instant snapshot by OPC

See "instant snapshot."

[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 and a proxy and synchronize the pair.

joined

The state when master and proxy are joined. Proxy volume in joined status cannot be accessed.

JRM

An abbreviation for Just Resynchronization Mechanism.

just resynchronization 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 [Itanium]**

A disk that retains the disk format and data when being registered with a class or being connected to a group.

[L]**local class [Linux2.6]**

A "local" type class. Objects that belong to a local class can be used only on the local node. With the local class, disks connected to the local node, excluding system disks, can be registered.

local object [Linux2.6]

An object that belongs to a "local" class.

logical slice

The component of a logical volume. A logical slice consists of contiguous blocks of one physical disk or multiple blocks of multiple physical disks. A mirror volume consists of one or more logical slices that are mirrored. A single volume, a stripe volume, a volume in a concatenation group, and a switch volume consist of a single logical slice. Logical slices may be called "slices."

logical volume

A logical disk device that provides the user with a block device interface [Linux2.4] as well as a physical slice (disk partition). The user can access data on a physical disk through accessing the logical volume. Logical volumes are classified into single volumes, mirror volumes, stripe volumes, volumes in concatenation groups, and switch volumes. Logical volumes may be called "volumes."

lower level group

A group that belongs to another group. A volume can not be created within the lower level group.

lower level shadow group

A shadow group that belongs to another shadow group. A shadow volume cannot be created in the lower level shadow group.

[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 volume from which the proxy volume is copied.
Volumes within master group are also called master volumes.
The term "master" can refer to both master volume and master group.

mirror group

A group whose type attribute is "mirror." It is the set of disks and lower level groups, which will be mirrored each other.

mirror slice

A logical slice that composes a mirror volume. A mirror volume is created by mirroring more than one logical slice. For each mirror volume, there will be one logical slice in each disk and each lower level group belonging to the mirror group.

mirror volume

A volume that is created within a mirror group. Its data will be made redundant by mirroring.

mirroring

To maintain redundancy by allocating the same data to multiple slices. A feature which allows you to continue access to a volume if one or more valid slices exist in the volume, even if some failures occur on the other slices.

MPHD disk

A virtual device that provides multiple access paths with the mphd driver. The user will be able to continue operation, even if an access path failure occurs.

MPLB disk

A virtual device that provides multiple access paths with the mplb driver. Being able to use multiple paths enables spreading out of the load when accessing the disk, and increases the availability at the time of path failure.

[N]**node**

One computer (with one operating system). For details, see the PRIMECLUSTER manual.

node identifier

A general name for each node within a PRIMECLUSTER system.

node name

The name returned with the `uname -n` command

NVURM

An abbreviation for Non-Volatile Update Region Map. This is information of the updated volume region stored on a private slice for just resynchronization.

[O]**object**

1. An SDX object. Objects are the generic term for classes, groups, disks, volumes, and slices.
2. The generic term for SDX objects and shadow objects.

offline backup

To create backups from a volume after stopping applications that use the volume.

online backup

To create volume backup while an application is accessing the volume.

online disk migration

To change from one disk to another without stopping applications accessing the disk data.

online snapshot

To create a volume snapshot while an application is accessing the volume.

OPC

Copying process using the OPC (One Point Copy) function.

OPC function

An abbreviation for One Point Copy function. It is one of the Advanced Copy function provided by Fujitsu Storage Systems ETERNUS. Disk data within the disk unit at a certain point will be instantly copied to separate area in the disk unit. The logical copying process is completed immediately after the request to start copying is sent from the server. The actual copying process will be performed later in the background.

For details, see the Fujitsu Storage Systems ETERNUS manual.

OPC session

OPC request from the host (server) to the Fujitsu Storage Systems ETERNUS.

For details, see the Fujitsu Storage Systems ETERNUS manual.

[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 was parted. Although parted, the master-proxy relationship is maintained.

parted

The status of a proxy isolated from the master. When a proxy volume is parted, access to the proxy that is independent of access to the master is available.

physical disk

A storage media possible of handling random input and output of data. Mainly used to refer to a set of physical slices.

physical scope

A group of nodes to which physical disks are connected.

physical slice

A part of physical disk and a series of physically continued blocks. A disk partition. The starting sector number and the number of sectors occupied by the physical slice is registered with the partition table.

physical special file

A device special file of a physical disk, such as `/dev/sd[a-z]* [1-4]*`.

PRIMECLUSTER

Software to configure the Cluster System. For details, see the PRIMECLUSTER manual.

PRIMECLUSTER system

A cluster system configured using PRIMECLUSTER. For details, see the PRIMECLUSTER manual.

private slice

A special slice used by GDS for control. Part of the SDX disk. Cannot be used as a logical slice.

proxy

See "proxy object."

proxy group

Replication group of master group.

It has the same volume configuration as the master group, and each volume is the proxy volume of the master volume that is within the master group.

The term "proxy" can refer to both proxy group and proxy volume.

proxy object

A generic term for proxy volume and proxy group.

Also simply referred to as "proxy."

proxy volume

A replication volume of a master volume. This may be called an alternative volume. Proxy volumes and proxy groups may be called "proxies" collectively.

[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 Fujitsu Storage Systems ETERNUS. REC is EC function available between Fujitsu Storage Systems ETERNUS boxes. For details, see the Fujitsu Storage Systems ETERNUS manual.

register disk

To add a physical disk to a certain class. The registered disk is managed by GDS subsequently.

rejoin

To rejoin the parted proxy with master.

relate

To relate a pair of volumes or a pair of groups as a master and a proxy and leave them parted.

remove disk

To remove an SDX disk from disk class.

resource

The concept describing the elements managed by PRIMECLUSTER among hardware and software that compose the PRIMECLUSTER system. In a PRIMECLUSTER system, shared disk units and GDS disk classes are managed as resources.

For details, see the PRIMECLUSTER manual.

resource database

The database that manages resource information in a PRIMECLUSTER system. The resource database is managed by the PRIMECLUSTER CF cluster resource management facility. For details, see the PRIMECLUSTER manual.

resource registration

The function that registers hardware units with the resource database in a PRIMECLUSTER system. This function is also referred to as automatic resource registration. To perform resource registration, use the PRIMECLUSTER CF clautoconfig command. For details, see "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide" and "Appendix H Shared Disk Unit Resource Registration."

restore

To restore the master's data, by overwriting the master using the parted proxy data.

restore object configuration [Linux2.6]

To create object configuration of a class according to the configuration table backed up to a configuration file.

resynchronization copying

To copy data between mirror slices to recover mirroring state when rebooting after a system failure, after a cluster failover, or after reattaching a detached slice to a volume.
Or, to copy data between master and proxy to match their data when rejoining a parted proxy to a master.

root class [Itanium]

A "root" type class. Objects that belong to a root class can be used only on the local node. With the root class, disks connected to the local node, including system disks, can be registered.

root disk [Itanium]

A disk to which the root file system will be mounted.

root object [Itanium]

An object that belongs to a "root" class.

root volume [Itanium]

The volume mounted as the root file system.

ROPC

Copying process by ROPC (Remote One Point Copy) function.

ROPC function

An abbreviation Remote One Point Copy function. It is one of the Advanced Copy functions provided by Fujitsu Storage Systems ETERNUS. ROPC is OPC function available between Fujitsu Storage Systems ETERNUS boxes. For details, see the Fujitsu Storage Systems ETERNUS manual.

[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.

SAN boot

The function that boots the OS from a disk array unit connected to the SAN instead of an internal disk.

scope

One of the class attributes. The scope indicates a group of nodes on which objects of the class are available.

SDX disk

A disk registered with a disk class. SDX disks may be called single disks, keep disks, spare disks, and undefined disks according to their usage. SDX disks may be called "disks."

SDX object

A virtual resource managed by GDS. SDX objects are the generic term for classes, groups, disks, volumes, and slices. SDX objects may be called "objects."

sector

See "block."

shadow class

A collection of shadow objects. Disk classes and shadow classes may be called "classes" collectively. Physical disks managed as SDX disks in another domain and those as the hard copy destinations of SDX disk data copied with disk unit copy functions can be registered with a shadow class and accordingly can be managed and accessed as shadow objects. Objects that belong to a shadow class can be used only on the local node.

shadow disk

A disk registered with a shadow class. SDX disks and shadow disks may be called "disks" collectively.

shadow group

A group created in a shadow class. The shadow group is a collection of shadow disks and lower shadow groups. Disk groups and shadow groups may be called "groups" collectively.

shadow object

A virtual resource managed by GDS. Shadow objects are the generic term for shadow classes, shadow groups, shadow disks, shadow volumes, and shadow slices. These shadow objects are available on a node where GDS Snapshot has been implemented. SDX objects and shadow objects may be called "objects" collectively.

shadow slice

The slice that composes a shadow volume. Logical slices and shadow slices may be called "slices" collectively.

shadow volume

A volume created in a shadow group. Logical volumes and shadow volumes may be called "volumes" collectively.

shared class

A "shared" type class. Objects that belong to a shared class are available on multiple nodes. With the shared class, disks connected to multiple nodes, including the local node, can be registered. Shared classes are also the units of resources used by cluster applications.

shared disk definition file

A file that declares physical connection configuration of nodes and shared disk units in a PRIMECLUSTER system. Specify the full path to a shared disk definition file with the -f option when executing the clautoconfig(8) command, and disk unit information in the shared disk definition file is registered with the PRIMECLUSTER resource database. For details, see "PRIMECLUSTER Cluster Foundation (CF) Configuration and Administration Guide" and "Appendix H Shared Disk Unit Resource Registration."

shared object

An object that belongs to a "shared" class.

single disk

An SDX disk which does not belong to any group and is capable of creating single volumes.

single node

A node that is not Cluster System configuration.

single slice

A logical slice that consists a single volume. Single volume consists of one single slice. It is also a series of physically continued block in a single disk.

single volume

A volume that is created within a single disk that does not belong to a group. Its data is not made redundant.

size of a group

The total of available size of a group and the size of one private slice.

slice

1. A logical slice.
2. The generic term for logical slices and shadow slices.

slice number

A partition number.

snapshot

To create a replica of a volume at a certain moment in another. It is also a created replica. GDS provides snapshot functions by slice detachment and by proxy volumes. Through use of proxy volumes, snapshot by synchronization and instant snapshot by OPC are available.

snapshot by slice detachment

To create a snapshot of a volume at a moment to a slice by detaching a slice from a mirror volume temporarily. Also it is a created snapshot.

snapshot by synchronization

To create a snapshot of a master at a moment to a proxy by parting the master and proxy which were previously joined and synchronized. Also it is a created snapshot.

soft copy

To copy disk data, using the GDS 's driver.

spare disk

A reserved disk to recover mirroring when a mirrored disk fails. When an I/O error is detected in a disk belonging to a mirror group or its lower level group, this spare disk will automatically be connected to recover normal mirroring.

SRDF

EMC's software. An abbreviation for Symmetrix Remote Data Facility. Using Symmetrix hardware functions, disk data is copied between Symmetrix cases. For details see the SRDF manual.

start volume

To make volume accessible.

stop volume

To make volume inaccessible.

Storage Area Network

See "SAN."

stripe column

Disks and lower level groups by which striping is conducted. In other words, disks and lower level groups belonging to a stripe group.

stripe group

A disk group whose type attribute is "stripe." This is a group of disks and lower level groups by which striping is conducted.

stripe volume

A volume that is created within a stripe group. Stripe volume can distribute the I/O load among multiple disks with its striping function. Its data will not be redundant.

stripe width

The size of divided data when striping.

striping

To divide data into a certain size and then allocate the divided data alternately to multiple slices. Striping distributes the I/O load across multiple physical disks.

swap

To swap the slices comprising master and proxy that are joined and synchronized.

switch

Disk switch.

switch group

A group of the switch type. This is a gathering of an active disk and an inactive disk that are switched.

switch slice

An operating logical slice comprising a switch volume. A switch volume consists of one switch slice. The switch slice is a gathering of physically continuous blocks in the active disk that belongs to the switch group.

switch volume

A volume created in a switch group. With the disk switch function, the roles of the active and inactive disks can be switched. Data will not be redundant.

Symmetrix

Storage devices from EMC. For details, refer to a guide to Symmetrix.

synchronization

The state when mirroring is operating normally.
Or, a state when the master and proxy data are the same.

synchronization copying

To copy data from a volume to an attached slice in order to make a mirrored state or to increase the multiplicity of the mirroring configuration, or to copy data between a master and a proxy in order to make their contents equivalent.

synchronized

The state when synchronization is maintained.

system disk

The disk that contains the Linux operating system running on the system. It refers to the entire disk containing slice operating as one of the following file systems (or a swap area).
`/`, `/usr`, `/var`, `/boot`, `/boot/efi`, or a swap area

system volume [Itanium]

The volume currently running as one of the following file systems (or a swap area).
`/`, `/usr`, `/var`, `/boot`, `/boot/efi`, or a swap area

[T]

the highest level group

A group that does not belong to any other group. Volumes can be created within the highest level group.

the highest level shadow group

A shadow group that does not belong to another shadow group. A shadow volume can be created in the highest level shadow group.

TimeFinder

EMC's software. Using Symmetrix hardware functions, disk data is copied within a Symmetrix case. For details see the TimeFinder manual.

[U]

undefined disk

An SDX disk that is neither a single disk, keep disk or a spare disk. Its usage is not specified yet.

update

To copy (overwrite) data from a master to a parted proxy. The updated proxy is a snapshot that retains an image of the master at a certain moment.

[V]

volume

1. A logical volume.
2. The generic term for logical volumes and shadow volumes.

Web-Based Admin View

Software for system operation and management using the Web screen. For details, see "Web-Based Admin View Operation Guide."