Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054
U.S.A.

Part No: 819–4245–10
December 2006, Revision A

060830@15490

Contents

Preface

Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy provides procedures for administering Hitachi TrueCopy data replication with Sun™ Cluster Geographic Edition software. This document is intended for experienced system administrators with extensive knowledge of Sun software and hardware. This document is not to be used as a planning or presales guide.

The instructions in this book assume knowledge of the Solaris™ Operating System (Solaris OS) and expertise with the volume manager software that is used with Sun Cluster software.

Related Books

Information about related Sun Cluster Geographic Edition topics is available in the documentation that is listed in the following table. All Sun Cluster Geographic Edition documentation is available at `http://docs.sun.com`.

Topic	Documentation
Overview	*Sun Cluster Geographic Edition Overview*
Glossary	*Sun Java Enterprise System Glossary*
Hardware administration	Individual hardware administration guides
Software installation	*Sun Cluster Geographic Edition Installation Guide*
System administration	*Sun Cluster Geographic Edition System Administration Guide*
	Sun Cluster Geographic Edition Data Replication Guide for Sun StorEdge Availability Suite
	Sun Cluster Geographic Edition Data Replication Guide for Hitachi TrueCopy
	Sun Cluster Geographic Edition Data Replication Guide for EMC Symmetrix Remote Data Facility
Command and function references	*Sun Cluster Geographic Edition Reference Manual*

For a complete list of Sun Cluster documentation, see the release notes for your Sun Cluster software at `http://docs.sun.com`.

Using UNIX Commands

This document contains information about commands that are used to install, configure, or administer a Sun Cluster Geographic Edition configuration. This document might not contain complete information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices.

See one or more of the following sources for this information:

- Online documentation for the Solaris software system
- Other software documentation that you received with your system
- Solaris OS man pages

Related Third-Party Web Site References

Third-party URLs are referenced in this document and provide additional, related information.

Note – Sun is not responsible for the availability of third-party web sites mentioned in this document. Sun does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. Sun will not be responsible or liable for any actual or alleged damage or loss caused or alleged to be caused by or in connection with use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

Documentation, Support, and Training

The Sun web site provides information about the following additional resources:

- Documentation (`http://www.sun.com/documentation/`)
- Support (`http://www.sun.com/support/`)
- Training (`http://www.sun.com/training/`)

Typographic Conventions

The following table describes the typographic conventions that are used in this book.

TABLE P–1 Typographic Conventions

Typeface	Meaning	Example
`AaBbCc123`	The names of commands, files, and directories, and onscreen computer output	Edit your `.login` file. Use `ls -a` to list all files. `machine_name% you have mail.`
`AaBbCc123`	What you type, contrasted with onscreen computer output	`machine_name%` **`su`** `Password:`
aabbcc123	Placeholder: replace with a real name or value	The command to remove a file is `rm` *filename*.
AaBbCc123	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the *User's Guide*. A *cache* is a copy that is stored locally. Do *not* save the file. **Note:** Some emphasized items appear bold online.

Shell Prompts in Command Examples

The following table shows the default UNIX system prompt and superuser prompt for the C shell, Bourne shell, and Korn shell.

TABLE P–2 Shell Prompts

Shell	Prompt
C shell	`machine_name%`
C shell for superuser	`machine_name#`
Bourne shell and Korn shell	`$`
Bourne shell and Korn shell for superuser	`#`

◆ ◆ ◆ CHAPTER 1

Replicating Data With Hitachi TrueCopy Software

During data replication, data from a primary cluster is copied to a backup or secondary cluster. The secondary cluster can be located at a geographically separated site from the primary cluster. This distance depends on the distance support that is available from your data replication product.

The Sun Cluster Geographic Edition software supports the use of Hitachi TrueCopy software for data replication. Before you start replicating data with Hitachi TrueCopy software, you must be familiar with the Hitachi TrueCopy documentation, have the Hitachi TrueCopy product, and have the latest Hitachi TrueCopy patches installed on your system. For information about installing the Hitachi TrueCopy software, see the Hitachi TrueCopy product documentation.

This chapter contains the procedures for configuring and administering data replication with Hitachi TrueCopy software. The chapter contains the following sections:

For information about creating and deleting data replication device groups, see "Administering Hitachi TrueCopy Data Replication Device Groups" on page 40. For information about obtaining a global and a detailed runtime status of replication, see "Checking the Runtime Status of Hitachi TrueCopy Data Replication" on page 59.

Administering Data Replication in a Hitachi TrueCopy Protection Group

This section summarizes the steps for configuring Hitachi TrueCopy data replication in a protection group.

TABLE 1–1 Administration Tasks for Hitachi TrueCopy Data Replication

Task	Description
Perform an initial configuration of the Hitachi TrueCopy software.	See "Initial Configuration of Hitachi TrueCopy Software" on page 12.
Create a protection group that is configured for Hitachi TrueCopy data replication.	See "How to Create and Configure a Hitachi TrueCopy Protection Group That Does Not Use Oracle Real Application Clusters" on page 28.
Add a device group that is controlled by Hitachi TrueCopy.	See "How to Add a Data Replication Device Group to a Hitachi TrueCopy Protection Group" on page 40.
Add an application resource group to the protection group.	See "How to Add an Application Resource Group to a Hitachi TrueCopy Protection Group" on page 37.
Replicate the protection group configuration to a secondary cluster.	See "How to Replicate the Hitachi TrueCopy Protection Group Configuration to a Secondary Cluster" on page 48.
Test the configured partnership and protection groups to validate the setup.	Perform a trial switchover or takeover and test some simple failure scenarios. See Chapter 3.
Activate the protection group.	See "How to Activate a Hitachi TrueCopy Protection Group" on page 52.
Check the runtime status of replication.	See "Checking the Runtime Status of Hitachi TrueCopy Data Replication" on page 59.
Detect failure.	See "Detecting Cluster Failure on a System That Uses Hitachi TrueCopy Data Replication" on page 63.
Migrate services by using a switchover.	See "Migrating Services That Use Hitachi TrueCopy Data Replication With a Switchover" on page 64.
Migrate services by using a takeover.	See "Forcing a Takeover on a System That Uses Hitachi TrueCopy Data Replication" on page 67.
Recover data after forcing a takeover.	See "Recovering Services to a Cluster on a System That Uses Hitachi TrueCopy Replication" on page 70.
Detect and recover from a data replication error.	See "Recovering From a Hitachi TrueCopy Data Replication Error" on page 81.

Initial Configuration of Hitachi TrueCopy Software

This section describes how to configure Hitachi TrueCopy software on the primary and secondary cluster. It also includes information about the preconditions for creating Hitachi TrueCopy protection groups.

Initial configuration of the primary and secondary clusters includes the following:

- Configuring a Hitachi TrueCopy device group, `devgroup1`, with the required number of disks

- Configuring the VERITAS Volume Manager disk group, `oradg1`
- Configuring the VERITAS Volume Manager volume, `vol1`
- Configuring the file system, which includes creating the file system, creating mount points, and adding entries to the `/etc/vfstab` file
- Creating an application resource group, `apprg1`, which contains a `HAStoragePlus` resource

If you use the Hitachi TrueCopy Command Control Interface (CCI) for data replication, you must use RAID Manager. For information about which version you should use, see the *Sun Cluster Geographic Edition Installation Guide*.

Note – This model requires specific hardware configurations with Sun StorEdge 9970/9980 Array or Hitachi Lightning 9900 Series Storage. Contact your Sun service representative for information about Sun Cluster configurations that are currently supported.

Sun Cluster Geographic Edition software supports the hardware configurations that are supported by the Sun Cluster software. Contact your Sun service representative for information about current supported Sun Cluster configurations.

Caution – If you are using storage-based replication, do not configure a replicated volume as a quorum device. The Sun Cluster Geographic Edition software does not support Hitachi TrueCopy S-VOL and Command Device as a Sun Cluster quorum device. See "Using Storage-Based Data Replication" in *Sun Cluster 3.0-3.1 Hardware Administration Manual for Solaris OS* for more information.

Configuring Data Replication With Hitachi TrueCopy Software on the Primary Cluster

This section describes the steps you must perform on the primary cluster before you can configure Hitachi TrueCopy data replication in Sun Cluster Geographic Edition software. To illustrate each step, this section uses an example of two disks, or LUNs, that are called `d1` and `d2`. These disks are in a Hitachi TrueCopy array that holds data for an application that is called `apprg1`.

Configuring the `/etc/horcm.conf` File

Configure Hitachi TrueCopy device groups on shared disks in the primary cluster by editing the `/etc/horcm.conf` file on each node of the cluster that has access to the Hitachi array. Disks `d1` and `d2` are configured to belong to a Hitachi TrueCopy device group, `devgroup1`. The application, `apprg1`, can run on all nodes that have Hitachi TrueCopy device groups configured.

For more information about how to configure the `/etc/horcm.conf` file, see the *Sun StorEdge SE 9900 V Series Command and Control Interface User and Reference Guide*.

The following table describes the configuration information from our example that is found in the `/etc/horcm.conf` file.

TABLE 1–2 Example Section of the `/etc/horcm.conf` File on the Primary Cluster

dev_group	dev_name	port number	TargetID	LU number	MU number
devgroup1	pair1	CL1–A	0	1	
devgroup1	pair2	CL1–A	0	2	

The configuration information in the table indicates that the Hitachi TrueCopy device group, `devgroup1`, contains two pairs. The first pair, `pair1`, is from the `d1` disk, which is identified by the tuple `<CL1–A , 0, 1>`. The second pair, `pair2`, is from the `d2` disk and is identified by the tuple `<CL1–A, 0, 2>`. The replicas of disks `d1` and `d2` are located in a geographically separated Hitachi TrueCopy array. The remote Hitachi TrueCopy is connected to the partner cluster.

▼ How to Configure the Volumes for Use With Hitachi TrueCopy Replication

Hitachi TrueCopy supports VERITAS Volume Manager volumes. You must configure VERITAS Volume Manager volumes on disks `d1` and `d2`.

Caution – If you are using storage-based replication, do not configure a replicated volume as a quorum device. The Sun Cluster Geographic Edition software does not support Hitachi TrueCopy S-VOL and Command Device as a Sun Cluster quorum device. See "Using Storage-Based Data Replication" in *Sun Cluster 3.0-3.1 Hardware Administration Manual for Solaris OS* for more information.

1. **Create VERITAS Volume Manager disk groups on shared disks in `cluster-paris`.**

   For example, the `d1` and `d2` disks are configured as part of a VERITAS Volume Manager disk group, which is called `oradg1`, by using commands, such as `vxdiskadm` and `vxdg`.

2. **After configuration is complete, verify that the disk group was created by using the `vxdg list` command.**

   This command should list `oradg1` as a disk group.

3. **Create the VERITAS Volume Manager volume.**

   For example, a volume that is called `vol1` is created in the `oradg1` disk group. The appropriate VERITAS Volume Manager commands, such as `vxassist`, are used to configure the volume.

▼ How to Configure the Sun Cluster Device Group That Is Controlled by Hitachi TrueCopy Software

Before You Begin If you are using storage-based replication, do not configure a replicated volume as a quorum device. The Sun Cluster Geographic Edition software does not support Hitachi TrueCopy S-VOL and Command Device as a Sun Cluster quorum device. See "Using Storage-Based Data Replication" in *Sun Cluster 3.0-3.1 Hardware Administration Manual for Solaris OS* for more information.

1. **Register the VERITAS Volume Manager disk group that you configured in the previous procedure.**

   Use the Sun Cluster commands, `scsetup` or `scconf`.

   For more information about these commands, refer to the `scsetup`(1M) or the `scconf`(1M) man page.

2. **Synchronize the VERITAS Volume Manager configuration with Sun Cluster software, again by using the `scsetup` or `scconf` commands.**

3. **After configuration is complete, verify the disk group registration.**

   ```
   # scstat -D
   ```

   The VERITAS Volume Manager disk group, `oradg1`, should be displayed in the output.

   For more information about the `scstat` command, see the `scstat`(1M) man page.

▼ How to Configure a Highly Available File System for Hitachi TrueCopy Replication

Before You Begin Before you configure the file system on `cluster-paris`, ensure that the Sun Cluster entities you require, such as application resource groups, device groups, and mount points, have already been configured.

If you are using storage-based replication, do not configure a replicated volume as a quorum device. The Sun Cluster Geographic Edition software does not support Hitachi TrueCopy S-VOL and Command Device as a Sun Cluster quorum device. See "Using Storage-Based Data Replication" in *Sun Cluster 3.0-3.1 Hardware Administration Manual for Solaris OS* for more information.

1. **Create the required file system on the `vol1` volume at the command line.**

2. **Add an entry to the `/etc/vfstab` file that contains information such as the mount location.**

   Whether the file system is to be mounted locally or globally depends on various factors, such as your performance requirements, or the type of application resource group you are using.

   Note – You must set the `mount at boot` field in this file to `no`. This value prevents the file system from mounting on the secondary cluster at cluster startup. Instead, the Sun Cluster software and the Sun Cluster Geographic Edition framework handle mounting the file system by using the `HAStoragePlus` resource when the application is brought online on the primary cluster. Data must not be mounted on the secondary cluster or data on the primary will not be replicated to the secondary cluster. Otherwise, the data will not be replicated from the primary cluster to the secondary cluster.

   ---

3. **Add the `HAStoragePlus` resource to the application resource group, `apprg1`.**

   Adding the resource to the application resource group ensures that the necessary file systems are remounted before the application is brought online.

   For more information about the `HAStoragePlus` resource type, refer to the *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

Example 1–1 Configuring a Highly Available Cluster Global File System

This example assumes that the `apprg1` resource group already exists.

1. Create a UNIX file system (UFS).

   ```
   phys-paris-1# newfs dev/vx/dsk/oradg1/vol1
   ```

   The following entry is created in the `/etc/vfstab` file:

   ```
   # /dev/vs/dsk/oradg1/vol1 /dev/vx/rdsk/oradg1/vol1 /mounts/sample \
   ufs 2 no logging
   ```

2. Add the `HAStoragePlus` resource type.

   ```
   phys-paris-1# scrgadm -a -j rs-hasp -g apprg1 -t SUNW.HAStoragePlus \
   -x FilesystemMountPoints=/mounts/sample -x AffinityOn=TRUE \
   -x GlobalDevicePaths=oradg1
   ```

Configuring Data Replication With Hitachi TrueCopy Software on the Secondary Cluster

This section describes the steps you must complete on the secondary cluster before you can configure Hitachi TrueCopy data replication in Sun Cluster Geographic Edition software.

Configuring the `/etc/horcm.conf` File

You must configure the Hitachi TrueCopy device group on shared disks in the secondary cluster as you did on the primary cluster by editing the `/etc/horcm.conf` file on each node of the cluster that has access to the Hitachi array. Disks `d1` and `d2` are configured to belong to a Hitachi TrueCopy device group that is called `devgroup1`. The application, `apprg1`, can run on all nodes that have Hitachi TrueCopy device groups configured.

For more information about how to configure the `/etc/horcm.conf` file, see the *Sun StorEdge SE 9900 V Series Command and Control Interface User and Reference Guide*.

The following table describes the configuration information from the example that is found in the `/etc/horcm.conf` file.

TABLE 1–3 Example Section of the `/etc/horcm.conf` File on the Secondary Cluster

dev_group	dev_name	port number	TargetID	LU number	MU number
devgroup1	pair1	CL1–C	0	20	
devgroup1	pair2	CL1–C	0	21	

The configuration information in the table indicates that the Hitachi TrueCopy device group, `devgroup1`, contains two pairs. The first pair, `pair1`, is from the `d1` disk, which is identified by the tuple `<CL1–C , 0, 20>`. The second pair, `pair2`, is from the `d2` disk and is identified by the tuple `<CL1–C, 0, 21>`.

After you have configured the `/etc/horcm.conf` file on the secondary cluster, you can see the status of the pairs by using the `pairdisplay` command as follows:

```
phys-paris-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-A , 0, 1) 54321 1..  SMPL ----  ------,----- ----  -
devgroup1 pair1(R) (CL1-C , 0, 20)12345 609..SMPL ----  ------,----- ----  -
devgroup1 pair2(L) (CL1-A , 0, 2) 54321 2..  SMPL ----  ------,----- ----  -
devgroup1 pair2(R) (CL1-C , 0, 21)12345 610..SMPL ----  ------,----- ----  -
```

Configuring the Other Entities on the Secondary Cluster

Next, you need to configure the volume manager, the Sun Cluster device groups, and the highly available cluster global file system. You can configure these entities in two ways:

- By replicating the volume manager information from `cluster-paris`
- By creating a copy of the volume manager configuration on the LUNs of `cluster-newyork` by using the VERITAS Volume Manager commands `vxdiskadm` and `vxassist`

Each of these methods is described in the following procedures.

▼ How to Replicate the Volume Manager Configuration Information From the Primary Cluster

Before You Begin If you are using storage-based replication, do not configure a replicated volume as a quorum device. The Sun Cluster Geographic Edition software does not support Hitachi TrueCopy S-VOL and Command Device as a Sun Cluster quorum device. See "Using Storage-Based Data Replication" in *Sun Cluster 3.0-3.1 Hardware Administration Manual for Solaris OS* for more information.

1 Start replication for the `devgroup1` device group.

```
phys-paris-1# paircreate -g devgroup1 -vl -f async

phys-paris-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-A , 0, 1) 54321   1..P-VOL COPY ASYNC ,12345 609   -
devgroup1 pair1(R) (CL1-C , 0, 20)12345 609..S-VOL COPY ASYNC ,-----   1   -
devgroup1 pair2(L) (CL1-A , 0, 2) 54321   2..P-VOL COPY ASYNC ,12345 610   -
devgroup1 pair2(R) (CL1-C , 0, 21)12345 610..S-VOL COPY ASYNC ,-----   2   -
```

2 Wait for the state of the pair to become `PAIR` on the secondary cluster.

```
phys-newyork-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-C , 0, 20)12345 609..S-VOL PAIR ASYNC,-----, 1     -
```

```
devgroup1 pair1(R) (CL1-A , 0, 1) 54321   1..P-VOL PAIR ASYNC,12345, 609   -
devgroup1 pair2(L) (CL1-C , 0, 21)12345 610..S-VOL PAIR ASYNC,-----, 2     -
devgroup1 pair2(R) (CL1-A , 0, 2)54321    2..P-VOL PAIR ASYNC,12345, 610   -
```

3 Split the pair by using the `pairsplit` command and confirm that the secondary volumes on `cluster-newyork` are writable by using the `-rw` option.

```
phys-newyork-1# pairsplit -g devgroup1 -rw
phys-newyork-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-C , 0, 20)12345 609..S-VOL SSUS ASYNC, -----  1    -
devgroup1 pair1(R) (CL1-A , 0, 1) 54321   1..P-VOL PSUS ASYNC,12345  609   W
devgroup1 pair2(L) (CL1-C , 0,21) 12345 610..S-VOL SSUS ASYNC,-----   2    -
devgroup1 pair2(R) (CL1-A , 0, 2) 54321   2..P-VOL PSUS ASYNC,12345  610   W
```

4 Import the VERITAS Volume Manager disk group, `oradg1`.

```
phys-newyork-1# vxdg -C import oradg1
```

5 Verify that the VERITAS Volume Manager disk group was successfully imported.

```
phys-newyork-1# vxdg list
```

6 Enable the VERITAS Volume Manager volume.

```
phys-newyork-1# /usr/sbin/vxrecover -g oradg1 -s -b
```

7 Verify that the VERITAS Volume Manager volumes are recognized and enabled.

```
phys-newyork-1# vxprint
```

8 Register the VERITAS Volume Manager disk group, `oradg1`, in Sun Cluster.

```
phys-newyork-1# scconf -a -D type=vxvm, name=oradg1, \
nodelist=phys-newyork-1:phys-newyork-2
```

9 Synchronize the volume manager information with the Sun Cluster device group and verify the output.

```
phys-newyork-1# scconf -c -D name=oradg1,sync
phys-newyork-1# scstat -D
```

10 Add an entry to the `/etc/vfstab` file on `phys-newyork-1`.

```
phys-newyork-1# /dev/vx/dsk/oradg1/vol1 /dev/vx/rdsk/oradg1/vol1 \
/mounts/sample ufs 2 no logging
```

11 Create a mount directory on `phys-newyork-1`.

```
phys-newyork-1# mkdir -p /mounts/sample
```

12 Create an application resource group, `apprg1`, by using the `scrgadm` command.

```
phys-newyork-1# scrgadm -a -g apprg1
```

13 Create the `HAStoragePlus` resource in `apprg1`.

```
phys-newyork-1# scrgadm -a -j rs-hasp -g apprg1 -t SUNW.HAStoragePlus \
-x FilesystemMountPoints=/mounts/sample -x AffinityOn=TRUE \
-x GlobalDevicePaths=oradg1 \
```

14 If necessary, confirm that the application resource group is correctly configured by bringing it online and taking it offline again.

```
phys-newyork-1# scswitch -z -g apprg1 -h phys-newyork-1
phs-newyork-1# scswitch -F -g apprg1
```

15 Unmount the file system.

```
phys-newyork-1# umount /mounts/sample
```

16 Take the Sun Cluster device group offline.

```
phys-newyork-1# scswitch -F -D oradg1
```

17 Verify that the VERITAS Volume Manager disk group was deported.

```
phys-newyork-1# vxdg list
```

18 Reestablish the Hitachi TrueCopy pair.

```
phys-newyork-1# pairresync -g devgroup1
phys-newyork-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-C , 0, 20)12345 609..S-VOL PAIR ASYNC,-----   1    -
devgroup1 pair1(R) (CL1-A , 0, 1) 54321   1..P-VOL PAIR ASYNC,12345  609   W
devgroup1 pair2(L) (CL1-C , 0,21) 12345 610..S-VOL PAIR ASYNC,-----   2    -
devgroup1 pair2(R) (CL1-A , 0, 2) 54321   2..P-VOL PAIR ASYNC,12345  610   W
```

Initial configuration on the secondary cluster is now complete.

▼ How to Create a Copy of the Volume Manager Configuration

This task copies the volume manager configuration from the primary cluster, `cluster-paris`, to LUNs of the secondary cluster, `cluster-newyork`, by using the VERITAS Volume Manager commands `vxdiskadm` and `vxassist` command.

Note – The device group, `devgroup1`, must be in the `SMPL` state throughout this procedure.

---

1 Confirm that the pair is in the `SMPL` state.

```
phys-newyork-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-C , 0, 20)12345 609..SMPL ---- ------,----- ----   -
devgroup1 pair1(R) (CL1-A , 0, 1) 54321   1..SMPL ---- ------,----- ----   -
devgroup1 pair2(L) (CL1-C , 0, 21)12345 610..SMPL ---- ------,----- ----   -
devgroup1 pair2(R) (CL1-A, 0, 2) 54321    2..SMPL ---- ------,----- ----   -
```

2 **Create VERITAS Volume Manager disk groups on shared disks in `cluster-paris`.**

For example, the `d1` and `d2` disks are configured as part of a VERITAS Volume Manager disk group, which is called `oradg1`, by using commands, such as `vxdiskadm` and `vxdg`.

3 **After configuration is complete, verify that the disk group was created by using the `vxdg list` command.**

This command should list `oradg1` as a disk group.

4 **Create the VERITAS Volume Manager volume.**

For example, a volume that is called `vol1` is created in the `oradg1` disk group. The appropriate VERITAS Volume Manager commands, such as `vxassist`, are used to configure the volume.

5 **Import the VERITAS Volume Manager disk group.**

```
phys-newyork-1# vxdg -C import oradg1
```

6 **Verify that the VERITAS Volume Manager disk group was successfully imported.**

```
phys-newyork-1# vxdg list
```

7 **Enable the VERITAS Volume Manager volume.**

```
phys-newyork-1# /usr/sbin/vxrecover -g oradg1 -s -b
```

8 **Verify that the VERITAS Volume Manager volumes are recognized and enabled.**

```
phys-newyork-1# vxprint
```

9 **Register the VERITAS Volume Manager disk group, `oradg1`, in Sun Cluster.**

```
phys-newyork-1# scconf -a -D type=vxvm, name=oradg1, \
nodelist=phys-newyork-1:phys-newyork-2
```

10 **Synchronize the VERITAS Volume Manager information with the Sun Cluster device group and verify the output.**

```
phys-newyork-1# scconf -c -D name=oradg1, sync
phys-newyork-1# scstat -D
```

11 **Create a UNIX file system.**

```
phys-newyork-1# newfs dev/vx/dsk/oradg1/vol1
```

12 **Add an entry to the `/etc/vfstab` file on `phys-newyork-1`.**

```
phys-newyork-1# /dev/vx/dsk/oradg1/vol1 /dev/vx/rdsk/oradg1/vol1 /mounts/sample \
ufs 2 no logging
```

13 **Create a mount directory on `phys-newyork-1`.**

```
phys-newyork-1# mkdir -p /mounts/sample
```

14 Create an application resource group, `apprg1`, by using the `scrgadm` command.

```
phys-newyork-1# scrgadm -a -g apprg1
```

15 Create the `HAStoragePlus` resource in `apprg1`.

```
phys-newyork-1# scrgadm -a -j rs-hasp -g apprg1 -t SUNW.HAStoragePlus \
-x FilesystemMountPoints=/mounts/sample -x AffinityOn=TRUE \
-x GlobalDevicePaths=oradg1 \
```

16 If necessary, confirm that the application resource group is correctly configured by bringing it online and taking it offline again.

```
phys-newyork-1# scswitch -z -g apprg1 -h phys-newyork-1
phs-newyork-1# scswitch -F -g apprg1
```

17 Unmount the file system.

```
phys-newyork-1# umount /mounts/sample
```

18 Take the Sun Cluster device group offline.

```
phys-newyork-1# scswitch -F -D oradg1
```

19 Verify that the VERITAS Volume Manager disk group was deported.

```
phys-newyork-1# vxdg list
```

20 Verify that the pair is still in the `SMPL` state.

```
phys-newyork-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-C , 0, 20)12345 609..SMPL ---- ------,-----  ----  -
devgroup1 pair1(R) (CL1-A , 0, 1) 54321   1..SMPL ---- ------,-----  ----  -
devgroup1 pair2(L) (CL1-C , 0, 21)12345 610..SMPL ---- ------,-----  ----  -
devgroup1 pair2(R) (CL1-A, 0, 2)  54321   2..SMPL ---- ------,-----  ----  -
```

◆ ◆ ◆ CHAPTER 2

Administering Hitachi TrueCopy Protection Groups

This chapter contains the procedures for configuring and administering data replication with Hitachi TrueCopy software. The chapter contains the following sections:

Strategies for Creating Hitachi TrueCopy Protection Groups

Before you begin creating protection groups, consider the following strategies:

- Taking the application offline before creating the protection group.

  This strategy is the most straightforward because you use a single command to create the protection group on one cluster, retrieve the information on the other cluster, and start the protection group. However, because the protection group is not brought online until the end of the process, you must take the application resource group offline to add it to the protection group.

- Creating the protection group while the application remains online.

  While this strategy allows you to create a protection group without any application outage, it requires issuing more commands.

The following sections describe the steps for each strategy.

Creating a Protection Group While the Application Is Offline

To create a protection group while the application resource group is offline, complete the following steps.

- Create the protection group from a cluster node.

  For more information, see "How to Create and Configure a Hitachi TrueCopy Protection Group That Does Not Use Oracle Real Application Clusters" on page 28.
- Add the data replication device group to the protection group.

  For more information, see "How to Add a Data Replication Device Group to a Hitachi TrueCopy Protection Group" on page 40.
- Take the application resource group offline.
- Add the application resource group to the protection group.

  For more information, see "How to Add an Application Resource Group to a Hitachi TrueCopy Protection Group" on page 37.
- On the other cluster, retrieve the protection group configuration.

  For more information, see "How to Replicate the Hitachi TrueCopy Protection Group Configuration to a Secondary Cluster" on page 48.
- From either cluster, start the protection group globally.

  For more information, see "How to Activate a Hitachi TrueCopy Protection Group" on page 52.

Creating a Protection Group While the Application Is Online

To add an existing application resource group to a new protection group without taking the application offline, complete the following steps on the cluster where the application resource group is online.

- Create the protection group from a cluster node.

  For more information, see "How to Create and Configure a Hitachi TrueCopy Protection Group That Does Not Use Oracle Real Application Clusters" on page 28.
- Add the data replication device group to the protection group.

  For more information, see "How to Add a Data Replication Device Group to a Hitachi TrueCopy Protection Group" on page 40.
- Start the protection group locally.

  For more information, see "How to Activate a Hitachi TrueCopy Protection Group" on page 52.
- Add the application resource group to the protection group.

For more information, see "How to Add an Application Resource Group to a Hitachi TrueCopy Protection Group" on page 37.

Complete the following steps on the other cluster.

- Retrieve the protection group configuration.

  For more information, see "How to Replicate the Hitachi TrueCopy Protection Group Configuration to a Secondary Cluster" on page 48.
- Activate the protection group locally.

  For more information, see "How to Activate a Hitachi TrueCopy Protection Group" on page 52.

EXAMPLE 2–1 Creating a Hitachi TrueCopy Protection Group While the Application Remains Online

This example creates a protection group without taking the application offline.

In this example, the `apprg1` resource group is online on the `cluster-paris` cluster.

1. Create the protection group on `cluster-paris`.

```
phys-paris-1# geopg create -d truecopy -p Nodelist=phys-paris-1,phys-paris-2 \
-o Primary -s paris-newyork-ps tcpg
Protection group "tcpg" has been successfully created
```

2. Add the device group, `tcdg`, to the protection group.

```
phys-paris-1# geopg add-device-group -p fence_level=async tcdg tcpg
```

3. Activate the protection group locally.

```
phys-paris-1# geopg start -e local tcpg
Processing operation.... this may take a while....
Protection group "tcpg" successfully started.
```

4. Add to the protection group an application resource group that is already online.

```
phys-paris-1# geopg add-resource-group apprg1 tcpg
Following resource groups were successfully inserted:
            "apprg1"
```

5. Verify that the application resource group was added successfully.

```
phys-paris-1# geoadm status
Cluster: cluster-paris

   Partnership "paris-newyork-ps"    : OK
      Partner clusters               : newyork
      Synchronization                : OK
      ICRM Connection                : OK

      Heartbeat "hb_cluster-paris~cluster-newyork" monitoring \
```

EXAMPLE 2–1 Creating a Hitachi TrueCopy Protection Group While the Application Remains Online *(Continued)*

```
"paris-newyork-ps" OK
        Plug-in "ping-plugin"       : Inactive
        Plug-in "tcp_udp_plugin"    : OK

   Protection group "tcpg"          : Degraded
      Partnership                   : paris-newyork-ps
      Synchronization               : OK

      Cluster cluster-paris         : Degraded
         Role                       : Primary
         Configuration              : OK
         Data replication           : Degraded
         Resource groups            : OK

      Cluster cluster-newyork       : Unknown
         Role                       : Unknown
         Configuration              : Unknown
         Data Replication           : Unknown
         Resource Groups            : Unknown
```

6. On a node of the partner cluster, retrieve the protection group.

```
phys-newyork-1# geopg get -s paris-newyork-ps tcpg
Protection group "tcpg" has been successfully created.
```

7. Activate the protection group locally on the partner cluster.

```
phys-newyork-1# geopg start -e local tcpg
Processing operation.... this may take a while....
Protection group "tcpg" successfully started.
```

8. Verify that the protection group was successfully created and activated.

   Running the `geoadm status` command on `cluster-paris` produces the following output:

```
phys-paris-1# geoadm status
Cluster: cluster-paris

   Partnership "paris-newyork-ps"       : OK
      Partner clusters                  : newyork
      Synchronization                   : OK
      ICRM Connection                   : OK

      Heartbeat "hb_cluster-paris~cluster-newyork" monitoring \
"paris-newyork-ps": OK
        Plug-in "ping-plugin"           : Inactive
        Plug-in "tcp_udp_plugin"        : OK
```

EXAMPLE 2–1 Creating a Hitachi TrueCopy Protection Group While the Application Remains Online *(Continued)*

```
Protection group "tcpg"              : Degraded
   Partnership                       : paris-newyork-ps
   Synchronization                   : OK

   Cluster cluster-paris             : Degraded
      Role                           : Primary
      Configuration                  : OK
      Data replication               : Degraded
      Resource groups                : OK

   Cluster cluster-newyork           : Degraded
      Role                           : Secondary
      Configuration                  : OK
      Data Replication               : Degraded
      Resource Groups                : OK
```

Creating, Modifying, Validating, and Deleting a Hitachi TrueCopy Protection Group

This section contains procedures for the following tasks:

- "How to Create and Configure a Hitachi TrueCopy Protection Group That Does Not Use Oracle Real Application Clusters" on page 28
- "How to Create a Protection Group for Oracle Real Application Clusters" on page 30
- "How the Data Replication Subsystem Validates the Device Group" on page 33
- "How to Modify a Hitachi TrueCopy Protection Group" on page 33
- "Validating a Hitachi TrueCopy Protection Group" on page 34
- "How to Delete a Hitachi TrueCopy Protection Group" on page 35

Note – You can create protection groups that are not configured to use data replication. To create a protection group that does not use a data replication subsystem, omit the `-d` *datareplicationtype* option when you use the `geopg` command. The `geoadm status` command shows a state for these protection groups of `Degraded`.

For more information, see "Creating a Protection Group That Does Not Require Data Replication" in *Sun Cluster Geographic Edition System Administration Guide*.

▼ How to Create and Configure a Hitachi TrueCopy Protection Group That Does Not Use Oracle Real Application Clusters

Use the steps in this task to create and configure a Hitachi TrueCopy protection group. If you want to use Oracle Real Application Clusters, see "How to Create a Protection Group for Oracle Real Application Clusters" on page 30.

Before You Begin Before you create a protection group, ensure that the following conditions are met:

- The local cluster is a member of a partnership.
- The protection group you are creating does not already exist.

Note – Protection group names are unique in the global Sun Cluster Geographic Edition namespace. You cannot use the same protection group name in two partnerships on the same system.

You can also replicate the existing configuration of a protection group from a remote cluster to the local cluster. For more information, see "Replicating the Hitachi TrueCopy Protection Group Configuration to a Secondary Cluster" on page 48.

1 Log in to a cluster node.

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see "Sun Cluster Geographic Edition Software and RBAC" in *Sun Cluster Geographic Edition System Administration Guide*.

2 Create a new protection group by using the `geopg create` command.

This command creates a protection group on all nodes of the local cluster.

```
# geopg create -s partnershipname -o localrole -d truecopy [-p property [-p...]] \
protectiongroupname
```

`-s` *partnershipname*	Specifies the name of the partnership.
`-o` *localrole*	Specifies the role of this protection group on the local cluster as either `primary` or `secondary`.
`-d truecopy`	Specifies that the protection group data is replicated by the Hitachi TrueCopy software.
`-p` *propertysetting*	Specifies the properties of the protection group. You can specify the following properties: ■ `Description` – Describes the protection group. ■ `Timeout` – Specifies the timeout period for the protection group in seconds.

- `Nodelist` – Lists the host names of the machines that can be primary for the replication subsystem.
- `Cluster_dgs` – Lists the device groups where the data is written.

For more information about the properties you can set, see Appendix A, "Standard Sun Cluster Geographic Edition Properties," in *Sun Cluster Geographic Edition System Administration Guide*.

protectiongroupname Specifies the name of the protection group.

For information about the names and values that are supported by Sun Cluster Geographic Edition software, see Appendix B, "Legal Names and Values of Sun Cluster Geographic Edition Entities," in *Sun Cluster Geographic Edition System Administration Guide*.

For more information about the `geopg` command, refer to the `geopg`(1M) man page.

Example 2–2 Creating and Configuring a Hitachi TrueCopy Protection Group

This example creates a Hitachi TrueCopy protection group on `cluster-paris`, which is set as the primary cluster.

```
# geopg create -s paris-newyork-ps -o primary -d truecopy \
-p Nodelist=phys-paris-1,phys-paris-2 tcpg
```

Example 2–3 Creating a Hitachi TrueCopy Protection Group for Application Resource Groups That Are Online

This example creates a Hitachi TrueCopy protection group, `tcpg`, for an application resource group, `resourcegroup1`, that is currently online on `cluster-newyork`.

1. Create the protection group without the application resource group.

   ```
   # geopg create -s paris-newyork-ps -o primary -d truecopy \
    -p nodelist=phys-paris-1,phys-paris-2 tcpg
   ```

2. Activate the protection group.

   ```
   # geopg start -e local tcpg
   ```

3. Add the application resource group.

   ```
   # geopg add-resource-group resourcegroup1 tcpg
   ```

▼ How to Create a Protection Group for Oracle Real Application Clusters

Before You Begin Before you create a protection group for Oracle Real Application Clusters, ensure that the following conditions are met:

- The nodelist of the protection group must be the same as the nodelist of Oracle Real Application Clusters framework resource group.
- If one cluster is running Oracle Real Application Clusters on a different number of nodes than another cluster, ensure that all nodes on both clusters have the same resource groups defined.
- All Oracle Real Application Clusters server resource groups and all Oracle Real Application Clusters listener resource groups must belong to the same protection group.
- *If you are using cluster volume manager to manage data*, you must specify the cluster volume manager disk group and Sun Cluster device groups for other data volumes in the `cluster_dgs` property.

When a cluster and the cluster volume manager software restart, the Oracle Real Application Clusters framework automatically tries to import all cluster volume manager device groups that were imported already before cluster went down. Therefore, the attempt to import the device groups to the original primary fails.

1 Log in to a cluster node on the primary cluster.

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see "Sun Cluster Geographic Edition Software and RBAC" in *Sun Cluster Geographic Edition System Administration Guide*.

2 Create a new protection group with a cluster volume manager disk group by using the `geopg create` command.

This command creates a protection group on all nodes of the local cluster.

```
# geopg create -s partnershipname -o localrole -d truecopy [-p property [-p...]] \
protectiongroupname
```

`-s` *partnershipname*	Specifies the name of the partnership.
`-o` *localrole*	Specifies the role of this protection group on the local cluster as `primary`.
`-d truecopy`	Specifies that the protection group data is replicated by the Hitachi TrueCopy software.
`-p` *propertysetting*	Specifies the properties of the protection group. You can specify the following properties: ■ `Description` – Describes the protection group. ■ `Timeout` – Specifies the timeout period for the protection group in seconds.

- `Nodelist` – Lists the host names of the machines that can be primary for the replication subsystem.
- `Cluster_dgs` – Specifies the cluster volume manager disk group where the data is written.

For more information about the properties you can set, see Appendix A, "Standard Sun Cluster Geographic Edition Properties," in *Sun Cluster Geographic Edition System Administration Guide*.

protectiongroupname Specifies the name of the protection group.

For information about the names and values that are supported by Sun Cluster Geographic Edition software, see Appendix B, "Legal Names and Values of Sun Cluster Geographic Edition Entities," in *Sun Cluster Geographic Edition System Administration Guide*.

For more information about the `geopg` command, refer to the `geopg`(1M) man page.

3 Add a Hitachi TrueCopy device group to the protection group.

```
# geopg add-device-group [-p property [-p...]] \
protectiongroupname
```

`-p` *propertysetting* Specifies the properties of the protection group.

You can specify the `Fence_level` properties which defines the fence level that is used by the disk device group. The fence level determines the level of consistency among the primary and secondary volumes for that disk device group. You must set this to `never`.

Caution – To avoid application failure on the primary cluster, specify a `Fence_level` of `never` or `async`. If the `Fence_level` parameter is not set to `never` or `async`, data replication might not function properly when the secondary site goes down.

If you specify a `Fence_level` of `never`, the data replication roles do not change after you perform a takeover.

Do not use programs that would prevent the `Fence_level` parameter from being set to `data` or `status` because these values might be required in special circumstances.

If you have special requirements to use a `Fence_level` of `data` or `status`, consult your Sun representative.

For more information about the properties you can set, see Appendix A, "Standard Sun Cluster Geographic Edition Properties," in *Sun Cluster Geographic Edition System Administration Guide*.

protectiongroupname Specifies the name of the protection group.

4 Add the Oracle Real Application Clusters framework resource group, all Oracle Real Application Clusters server resource groups, and all Oracle Real Application Clusters listener resource groups to the protection group.

```
# geopg add-resource-group resourcegroup protectiongroupname
```

resourcegroup Specifies a comma-separated list of resource groups to add to or delete from the protection group. The specified resource groups must already be defined.

The protection group must be online before you add a resource group. The `geopg add-resource-group` command fails when a protection group is offline and the resource group that is being added is online.

Note – If a protection group has already been started at the time that you add a resource group, the resource group remains unmanaged. You must start the resource group manually by running the `geopg start` command.

protectiongroupname Specifies the name of the protection group.

Example 2–4 Creating a Protection Group for Oracle Real Application Clusters

This example creates the protection group `pg1` which uses Oracle Real Application Clusters and the cluster volume manager.

A cluster volume manager disk group `oracle-dg` controls the data which is replicated by the Hitachi TrueCopy device group `VG01`. The nodelist of the Oracle Real Application Clusters framework resource group is set to all nodes of the cluster.

1. Create the protection group on the primary cluster with the cluster volume manager disk group `oracle-dg`.

```
# geopg create -s pts1 -o PRIMARY -d Truecopy -p cluster_dgs=oracle-dg pg1
Protection group "pg1" successfully created.
```

2. Add the Hitachi TrueCopy device group `VG01` to protection group `pg1`.

```
# geopg add-device-group --property fence_level=never VG01 pg1
Device group "VG01" successfully added to the protection group "pg1".
```

3. Add the Oracle Real Application Clusters framework resource group `rac-framework-rg`, all Oracle Real Application Clusters server resource groups, and all Oracle Real Application Clusters listener resource groups to the protection group.

```
# geopg add-resource-group rac-framework-rg,rac-server-rg1,\
rac-listener-rg1,rac-server-rg2,rac-listener-rg2 pg1
```

How the Data Replication Subsystem Validates the Device Group

Before creating the protection group, the data replication layer validates that the `horcmd` daemon is running.

The data replication layer validates that the `horcmd` daemon is running on at least one node that is specified in the `Nodelist` property. For more information about the `horcmd` daemon, see the *Sun StorEdge SE 9900 V Series Command and Control Interface User and Reference Guide*.

If the `Cluster_dgs` property is specified, then the data replication layer verifies that the device group specified is a valid Sun Cluster device group. The data replication layer also verifies that the device group is of a valid type.

Note – The device groups that are specified in the `Cluster_dgs` property must be written to only by applications that belong to the protection group. This property must not specify device groups that receive information from applications outside the protection group.

A Sun Cluster resource group is automatically created when the protection group is created.

This resource in this resource group monitors data replication. The name of the Hitachi TrueCopy data replication resource group is `rg-tc-`*protectiongroupname*.

Caution – These automatically created replication resource groups are for Sun Cluster Geographic Edition internal implementation purposes only. Use caution when you modify these resource groups by using Sun Cluster commands.

▼ How to Modify a Hitachi TrueCopy Protection Group

Before You Begin Before modifying the configuration of your protection group, ensure that the protection group you want to modify exists locally.

1 Log in to a cluster node.

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see "Sun Cluster Geographic Edition Software and RBAC" in *Sun Cluster Geographic Edition System Administration Guide*.

2 Modify the configuration of the protection group.

This command modifies the properties of a protection group on all nodes of the local cluster. If the partner cluster contains a protection group of the same name, this command also propagates the new configuration information to the partner cluster.

```
# geopg set-prop -p property [-p...] \
protectiongroupname
```

`-p` *propertysetting*	Specifies the properties of the protection group. For more information about the properties you can set, see Appendix A, "Standard Sun Cluster Geographic Edition Properties," in *Sun Cluster Geographic Edition System Administration Guide*.
protectiongroupname	Specifies the name of the protection group.

For information about the names and values that are supported by Sun Cluster Geographic Edition software, see Appendix B, "Legal Names and Values of Sun Cluster Geographic Edition Entities," in *Sun Cluster Geographic Edition System Administration Guide*.

For more information about the `geopg` command, refer to the `geopg`(1M) man page.

Example 2–5 Modifying the Configuration of a Protection Group

This example modifies the `Timeout` property of the protection group that was created in Example 2–2.

```
# geopg set-prop -p Timeout=400 tcpg
```

Validating a Hitachi TrueCopy Protection Group

During protection group validation, the Hitachi TrueCopy data replication subsystem validates the following:

- The `horcmd` daemon is running on at least one node that is specified in the `Nodelist` property of the protection group. The data replication layer also confirms that a path to a Hitachi TrueCopy storage device exists from the node on which the `horcmd` daemon is running.

  For more information about the `horcmd` daemon, see the *Sun StorEdge SE 9900 V Series Command and Control Interface User and Reference Guide*
- The device group specified is a valid Sun Cluster device group or a CVM device group if the `Cluster_dgs` property is specified. The data replication layer also verifies that the device group is of a valid type.
- The properties are validated for each Hitachi TrueCopy device group that has been added to the protection group.

When the `geoadm status` output displays that the `Configuration` status of a protection group is `Error`, you can validate the configuration by using the `geopg validate` command. This command checks the current state of the protection group and its entities.

If the protection group and its entities are valid, then the `Configuration` status of the protection groups is set to `OK`. If the `geopg validate` command finds an error in the configuration files, then the command displays a message about the error and the configuration remains in the error state. In such a case, you can fix the error in the configuration, and run the `geopg validate` command again.

▼ How to Validate a Hitachi TrueCopy Protection Group

Before You Begin Ensure that the protection group you want to validate exists locally and that the Common Agent Container is online on all nodes of both clusters in the partnership.

1 Log in to a cluster node.

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see "Sun Cluster Geographic Edition Software and RBAC" in *Sun Cluster Geographic Edition System Administration Guide*.

2 Validate the configuration of the protection group.

This command validates the configuration of the protection group on the local cluster only. To validate the protection group configuration on the partner cluster, run the command again on the partner cluster.

```
# geopg validate protectiongroupname
```

protectiongroupname Specifies a unique name that identifies a single protection group

Example 2–6 Validating the Configuration of a Protection Group

This example validates a protection group.

```
# geopg validate tcpg
```

▼ How to Delete a Hitachi TrueCopy Protection Group

Before You Begin If you want to delete the protection group everywhere, you must run the `geopg delete` command on each cluster where the protection group exists.

Before deleting a protection group, ensure that the following conditions are met:

- The protection group you want to delete exists locally.
- The protection group is offline on the local cluster.

Note – You must remove the application resource groups from the protection group in order to keep the application resource groups online while deleting the protection group. See Example 2–8 and Example 2–10 for examples of this procedure.

1 Log in to a node on the primary cluster.

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see "Sun Cluster Geographic Edition Software and RBAC" in *Sun Cluster Geographic Edition System Administration Guide*.

2 Delete the protection group.

This command deletes the configuration of the protection group from the local cluster. The command also removes the replication resource group for each Hitachi TrueCopy device group in the protection group. This command does not alter the pair state of the Hitachi TrueCopy device group.

```
# geopg delete protectiongroupname
```

protectiongroupname Specifies the name of the protection group

3 To delete the protection group on the secondary cluster, repeat step 1 and step 2 on `cluster-newyork`.

Example 2–7 Deleting a Protection Group

This example deletes a protection group from both partner clusters.

`cluster-paris` is the primary cluster. For a reminder of the sample cluster configuration, see "Example Sun Cluster Geographic Edition Cluster Configuration" in *Sun Cluster Geographic Edition System Administration Guide*.

```
# rlogin phys-paris-1 -l root
phys-paris-1# geopg delete tcpg
# rlogin phys-newyork-1 -l root
phys-newyork-1# geopg delete tcpg
```

Example 2–8 Deleting a Hitachi TrueCopy Protection Group While Keeping Application Resource Groups Online

This example keeps online two application resource groups, `apprg1` and `apprg2`, while deleting their protection group, `tcpg`. Remove the application resource groups from the protection group, then delete the protection group.

```
# geopg remove-resource-group apprg1,apprg2 tcpg
# geopg stop -e global tcpg
# geopg delete tcpg
```

Administering Hitachi TrueCopy Application Resource Groups

To make an application highly available, the application must be managed as a resource in an application resource group.

All the entities you configure for the application resource group on the primary cluster, such as application resources, installation, application configuration files, and resource groups, must be replicated to the secondary cluster. The resource group names must be identical on both clusters. Also, the data that the application resource uses must be replicated to the secondary cluster.

This section contains information about the following tasks:

- "How to Add an Application Resource Group to a Hitachi TrueCopy Protection Group" on page 37
- "How to Delete an Application Resource Group From a Hitachi TrueCopy Protection Group" on page 39

▼ How to Add an Application Resource Group to a Hitachi TrueCopy Protection Group

Before You Begin

You can add an existing resource group to the list of application resource groups for a protection group. Before you add an application resource group to a protection group, ensure that the following conditions are met:

- The protection group is defined.
- The resource group exists on both clusters and is in an appropriate state.
- The `Auto_start_on_new_cluster` property of the resource group is set to `False`. You can view this property by using the `scrgadm` command.

  ```
  # scrgadm -pvv -g apprg | grep Auto_start_on_new_cluster
  ```

  Setting the `Auto_start_on_new_cluster` property to `False` prevents the Sun Cluster resource group manager from automatically starting the resource groups in the protection group. Therefore, after the Sun Cluster Geographic Edition software restarts and communicates with the remote cluster to ensure that the remote cluster is running and that the remote cluster is the secondary cluster for that resource group. The Sun Cluster Geographic Edition software does not automatically start the resource group on the primary cluster.

Application resource groups should be online only on primary cluster when the protection group is activated.

Set the `Auto_start_on_new_cluster` property to `False` as follows:

```
scrgadm -c -g apprg1 -y Auto_start_on_new_cluster=False
```

- The application resource group must not have dependencies on resource groups and resources outside of this protection group. To add several application resource groups that share dependencies, you must add the application resource groups to the protection group in a single operation. If you add the application resource groups separately, the operation fails.

The protection group can be activated or deactivated and the resource group can be either `Online` or `Unmanaged`.

If the resource group is `Unmanaged` and the protection group is `Active` after the configuration of the protection group has changed, the local state of the protection group becomes `Degraded`.

If the resource group to add is `Online` and the protection group is deactivated, the request is rejected. You must activate the protection group before adding an active resource group.

1 Log in to a cluster node.

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see "Sun Cluster Geographic Edition Software and RBAC" in *Sun Cluster Geographic Edition System Administration Guide*.

2 Add an application resource group to the protection group.

This command adds an application resource group to a protection group on the local cluster. Then the command propagates the new configuration information to the partner cluster if the partner cluster contains a protection group of the same name.

`# geopg add-resource-group` *resourcegrouplist protectiongroup*

resourcegrouplist Specifies the name of the application resource group.

You can specify more than one resource group in a comma-separated list.

protectiongroup Specifies the name of the protection group.

For information about the names and values that are supported by Sun Cluster Geographic Edition software, see Appendix B, "Legal Names and Values of Sun Cluster Geographic Edition Entities," in *Sun Cluster Geographic Edition System Administration Guide*.

If the add operation is unsuccessful on the local cluster, the configuration of the protection group is not modified. Otherwise, the `Configuration` status is set to `OK` on the local cluster.

If the `Configuration` status is `OK` on the local cluster, but the add operation is unsuccessful on the partner cluster, the `Configuration` status is set to `Error` on the partner cluster.

After the application resource group is added to the protection group, the application resource group is managed as an entity of the protection group. Then the application resource group is affected by protection group operations such as start, stop, switchover, and takeover.

Example 2–9 Adding an Application Resource Group to a Protection Group

This example adds two application resource groups, `apprg1` and `apprg2`, to `tcpg`.

```
# geopg add-resource-group apprg1,apprg2 tcpg
```

▼ How to Delete an Application Resource Group From a Hitachi TrueCopy Protection Group

You can remove an application resource group from a protection group without altering the state or contents of an application resource group.

Before You Begin Ensure that the following conditions are met:

- The protection group is defined on the local cluster.
- The resource group to be removed is part of the application resource groups of the protection group. For example, you cannot remove a resource group that belongs to the data replication management entity.

1 Log in to a cluster node.

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see "Sun Cluster Geographic Edition Software and RBAC" in *Sun Cluster Geographic Edition System Administration Guide*.

2 Remove the application resource group from the protection group.

This command removes an application resource group from the protection group on the local cluster. If the partner cluster contains a protection group of the same name, then the command removes the application resource group from the protection group on the partner cluster.

```
# geopg remove-resource-group resourcegrouplist protectiongroup
```

resourcegrouplist Specifies the name of the application resource group.

You can specify more than one resource group in a comma-separated list.

protectiongroup Specifies the name of the protection group.

If the remove operation is unsuccessful on the local cluster, the configuration of the protection group is not modified. Otherwise, the `Configuration` status is set to `OK` on the local cluster.

If the `Configuration` status is `OK` on the local cluster, but the remove operation is unsuccessful on the partner cluster, the `Configuration` status is set to `Error` on the partner cluster.

Example 2–10 Deleting an Application Resource Group From a Protection Group

This example removes two application resource groups, `apprg1` and `apprg2`, from `tcpg`.

```
# geopg remove-resource-group apprg1,apprg2 tcpg
```

Administering Hitachi TrueCopy Data Replication Device Groups

This section provides the following information about administering Hitachi TrueCopy data replication device groups:

For details about configuring a Hitachi TrueCopy data replication protection group, see "How to Create and Configure a Hitachi TrueCopy Protection Group That Does Not Use Oracle Real Application Clusters" on page 28.

▼ How to Add a Data Replication Device Group to a Hitachi TrueCopy Protection Group

1 Log in to a cluster node.

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see "Sun Cluster Geographic Edition Software and RBAC" in *Sun Cluster Geographic Edition System Administration Guide*.

2 Create a data replication device group in the protection group.

This command adds a device group to a protection group on the local cluster and propagates the new configuration to the partner cluster if the partner cluster contains a protection group of the same name.

```
# geopg add-device-group -p property [-p...] devicegroupname protectiongroupname
```

`-p` *property*

Specifies the properties of the data replication device group.

You can specify the `Fence_level` property which defines the fence level that is used by the device group. The fence level determines the level of consistency among the primary and secondary volumes for that device group.

You can set this property to `data`, `status`, `never`, or `async`. When you use a `Fence_level` of `never` or `async`, the application can continue to write to the primary cluster even after failure on the secondary cluster. However, when you set the `Fence_level` property to `data` or `status`, the application on the primary cluster might fail because the secondary cluster is not available for the following reasons:

- Data replication link failure
- Secondary cluster and storage is down
- Storage on the secondary cluster is down

Caution – To avoid application failure on the primary cluster, specify a `Fence_level` of `never` or `async`.

If you specify a `Fence_level` of `never`, the data replication roles do not change after you perform a takeover.

If you have special requirements to use a `Fence_level` of `data` or `status`, consult your Sun representative.

For more information about application errors associated with different fence levels, see the *Sun StorEdge SE 9900 V Series Command and Control Interface User and Reference Guide*.

The other properties you can specify depend on the type of data replication you are using. For details about these properties, see Appendix A, "Standard Sun Cluster Geographic Edition Properties," in *Sun Cluster Geographic Edition System Administration Guide*.

devicegroupname

Specifies the name of the new data replication device group.

protectiongroupname

Specifies the name of the protection group that will contain the new data replication device group.

For information about the names and values that are supported by Sun Cluster Geographic Edition software, see Appendix B, "Legal Names and Values of Sun Cluster Geographic Edition Entities," in *Sun Cluster Geographic Edition System Administration Guide*.

For more information about the `geopg` command, refer to the `geopg`(1M) man page.

Example 2–11 Adding a Data Replication Device Group to a Hitachi TrueCopy Protection Group

This example creates a Hitachi TrueCopy data replication device group in the `tcpg` protection group.

```
# geopg add-device-group -p Fence_level=data devgroup1 tcpg
```

Validations Made by the Data Replication Subsystem

When the Hitachi TrueCopy device group, configured as `dev_group` in the `/etc/horcm.conf` file, is added to a protection group, the data replication layer makes the following validations.

- Validates that the `horcmd` daemon is running on at least one node in the `Nodelist` property of the protection group.

  For more information about the `horcmd` daemon, see the *Sun StorEdge SE 9900 V Series Command and Control Interface User and Reference Guide*
- Checks that the path to the storage device exists from all the nodes that are specified in the `Nodelist` property. The storage device controls the new Hitachi TrueCopy device group.
- The Hitachi TrueCopy device group properties that are specified in the `geopg add-device-group` command are validated as described in the following table.

Hitachi TrueCopy Device Group Property	Validation
devicegroupname	Checks that the specified Hitachi TrueCopy device group is configured on all of the cluster nodes that are specified in the `Nodelist` property.
`Fence_level`	If a pair is already established for this Hitachi TrueCopy device group, the data replication layer checks that the specified `Fence_level` matches the already established fence level. If a pair is not yet established, for example, if a pair is in the `SMPL` state, any `Fence_level` is accepted.

When a Hitachi TrueCopy device group is added to a protection group, a Sun Cluster resource is automatically created by this command. This resource monitors data replication. The name of the resource is `r-tc-`*protectiongroupname*`-`*devicegroupname*. This resource is placed in the corresponding Sun Cluster resource group, which is named `rg-tc-`*protectiongroupname*.

Caution – You must use caution before you modify these replication resources with Sun Cluster commands. These resources are for internal implementation purposes only.

How the State of the Hitachi TrueCopy Device Group Is Validated

For validation purposes, Sun Cluster Geographic Edition gives each Hitachi TrueCopy device group a state according to the current state of its pair. This state is returned by the `pairvolchk -g <DG> -ss` command.

The remainder of this section describes the individual device group states and how these states are validated against the local role of the protection group.

Determining the State of an Individual Hitachi TrueCopy Device Group

An individual Hitachi TrueCopy device group can be in one of the following states:

- `SMPL`
- `Regular Primary`
- `Regular Secondary`
- `Takeover Primary`
- `Takeover Secondary`

The state of a particular device group is determined by using the value that is returned by the `pairvolchk -g <DG> -ss` command. The following table describes the device group state associated with the values returned by the `pairvolchk` command.

TABLE 2–1 Individual Hitachi TrueCopy Device Group States

Output of `pairvolchk`	**Individual Device Group State**
`11 = SMPL`	`SMPL`

`22 / 42 = PVOL_COPY`
`23 / 42 = PVOL_PAIR`
`26 / 46 = PVOL_PDUB`
`47 = PVOL_PFUL`
`48 = PVOL_PFUS`
	`Regular Primary`

`24 / 44 = PVOL_PSUS`
`25 / 45 = PVOL_PSUE`
For these return codes, determining the individual device group category requires that the `horcmd` process be active on the remote cluster so that the `remote-pair-state` for this device group can be obtained.
	`Regular Primary`, if `remote-cluster-state !=SSWS`
	or
	`Takeover Secondary`, if `remote-cluster-state == SSWS`
	`SSWS`, when you use the `pairdisplay -g <DG> -fc` command.

TABLE 2–1 Individual Hitachi TrueCopy Device Group States *(Continued)*

Output of `pairvolchk`	**Individual Device Group State**

`32 / 52 = SVOL_COPY`
`33 / 53 = SVOL_PAIR`
`35 / 55 = SVOL_PSUE`
`36 / 56 = SVOL_PDUB`
`57 = SVOL_PFUL`
`58 = SVOL_PFUS`
	`Regular Secondary`

`34 / 54 = SVOL_PSUS`	`Regular Secondary`, if `local-cluster-state !=SSWS`
	or
	`Takeover Primary`, if `local-cluster-state == SSWS`
	`SSWS`, when you use the `pairdisplay -g <DG> -fc` command.

Determining the Aggregate Hitachi TrueCopy Device Group State

If a protection group contains only one Hitachi TrueCopy device group, then the aggregate device group state is the same as the individual device group state.

When a protection group contains multiple Hitachi TrueCopy device groups, the aggregate device group state is obtained as described in the following table.

TABLE 2–2 Conditions That Determine the Aggregate Device Group State

Condition	**Aggregate Device Group State**
All individual device group states are `SMPL`	`SMPL`
All individual device group states are either `Regular Primary` or `SMPL`	`Regular Primary`
All individual device group states are either `Regular Secondary` or `SMPL`	`Regular Secondary`
All individual device group states are either `Takeover Primary` or `SMPL`	`Takeover Primary`
All individual device group states are either `Takeover Secondary` or `SMPL`	`Takeover Secondary`

The aggregate device group state cannot be obtained for any other combination of individual device group states. This is considered a pair-state validation failure.

Validating the Local Role of the Protection Group Against the Aggregate Device Group State

The local role of a Hitachi TrueCopy protection group is validated against the aggregate device group state as described in the following table.

TABLE 2–3 Validating the Aggregate Device Group State Against the Local Role of a Protection Group

Aggregate Device Group State	Valid Local Protection Group Role
SMPL	primary or secondary
Regular Primary	primary
Regular Secondary	secondary
Takeover Primary	primary
Takeover Secondary	secondary

EXAMPLE 2–12 Validating the Aggregate Device Group State

This example validates the state of a Hitachi TrueCopy device group against the role of the Hitachi TrueCopy protection group to which it belongs.

First, the protection group is created as follows:

```
phys-paris-1# geopg create -s paris-newyork-ps -o primary -d truecopy tcpg
```

A device group, `devgroup1`, is added to the protection group, `tcpg`, as follows:

```
phys-paris-1# geopg add-device-group -p fence_level=async devgroup1 tcpg
```

The current state of a Hitachi TrueCopy device group, `devgroup1`, is provided in the output of the `pairdisplay` command as follows:

```
phys-paris-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-A , 0, 1) 12345   1..P-VOL PAIR ASYNC,54321  609   -
devgroup1 pair1(R) (CL1-C , 0, 20)54321 609..S-VOL PAIR ASYNC,-----   1    -
devgroup1 pair2(L) (CL1-A , 0, 2) 12345   2..P-VOL PAIR ASYNC,54321  610   -
devgroup1 pair2(R) (CL1-C , 0,21) 54321 610..S-VOL PAIR ASYNC,-----   2    -
```

The `pairvolchk -g <DG> -ss` command is run and returns a value of 23.

```
phys-paris-1# pairvolchk -g devgroup1 -ss
parivolchk : Volstat is P-VOL.[status = PAIR fence = ASYNC]
phys-paris-1# echo $?
23
```

EXAMPLE 2–12 Validating the Aggregate Device Group State *(Continued)*

The output of the `pairvolchk` command is `23`, which corresponds in Table 2–1 to an individual device group state of `Regular Primary`. Because the protection group contains only one device group, the aggregate device group state is the same as the individual device group state. The device group state is valid because the local role of the protection group, specified by the `-o` option, is `primary`, as specified in Table 2–3.

▼ How to Modify a Hitachi TrueCopy Data Replication Device Group

1 Log in to a cluster node.

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see "Sun Cluster Geographic Edition Software and RBAC" in *Sun Cluster Geographic Edition System Administration Guide*.

2 Modify the device group.

This command modifies the properties of a device group in a protection group on the local cluster. Then the command propagates the new configuration to the partner cluster if the partner cluster contains a protection group of the same name.

```
# geopg modify-device-group -p property [-p...] \
TCdevicegroupname  protectiongroupname
```

`-p` *property*	Specifies the properties of the data replication device group. For more information about the properties you can set, see Appendix A, "Standard Sun Cluster Geographic Edition Properties," in *Sun Cluster Geographic Edition System Administration Guide*.
TCdevicegroupname	Specifies the name of the new data replication device group.
protectiongroupname	Specifies the name of the protection group that will contain the new data replication device group.

Example 2–13 Modifying the Properties of a Hitachi TrueCopy Data Replication Device Group

This example modifies the properties of a data replication device group that is part of a Hitachi TrueCopy protection group.

```
# geopg modify-device-group -p fence_level=async tcdg tcpg
```

▼ How to Delete a Data Replication Device Group From a Hitachi TrueCopy Protection Group

Before You Begin You might delete a data replication device group from a protection group if you added a data replication device group to a protection group. Normally, after an application is configured to write to a set of disks, you would not change the disks.

Deleting a data replication device group does not stop replication or change the replication status of the data replication device group.

For information about deleting protection groups, refer to "How to Delete a Hitachi TrueCopy Protection Group" on page 35. For information about deleting application resource groups from a protection group, refer to "How to Delete an Application Resource Group From a Hitachi TrueCopy Protection Group" on page 39.

1 Log in to a cluster node.

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see "Sun Cluster Geographic Edition Software and RBAC" in *Sun Cluster Geographic Edition System Administration Guide*.

2 Remove the device group.

This command removes a device group from a protection group on the local cluster. Then the command propagates the new configuration to the partner cluster if the partner cluster contains a protection group of the same name.

```
# geopg remove-device-group devicegroupname protectiongroupname
```

devicegroupname Specifies the name of the data replication device group

protectiongroupname Specifies the name of the protection group

When a device group is deleted from a Hitachi TrueCopy protection group, the corresponding Sun Cluster resource, `r-tc-`*protectiongroupname*`-`*devicegroupname*, is removed from the replication resource group. As a result, the deleted device group is no longer monitored. The resource group is removed when the protection group is deleted.

Example 2–14 Deleting a Replication Device Group From a Hitachi TrueCopy Protection Group

This example removes a Hitachi TrueCopy data replication device group.

```
# geopg remove-device-group tcdg tcpg
```

Replicating the Hitachi TrueCopy Protection Group Configuration to a Secondary Cluster

After you have configured data replication, resource groups, and resources on your primary and secondary clusters, you can replicate the configuration of the protection group to the secondary cluster.

▼ How to Replicate the Hitachi TrueCopy Protection Group Configuration to a Secondary Cluster

Before You Begin Before you replicate the configuration of a Hitachi TrueCopy protection group to a secondary cluster, ensure that the following conditions are met:

- The protection group is defined on the remote cluster, not on the local cluster.
- The device groups in the protection group on the remote cluster exist on the local cluster.
- The application resource groups in the protection group on the remote cluster exist on the local cluster.
- The `Auto_start_on_new_cluster` property of the resource group is set to `False`. You can view this property by using the `scrgadm` command.

  ```
  # scrgadm -pvv -g apprg1 | grep Auto_start_on_new_cluster
  ```

  Setting the `Auto_start_on_new_cluster` property to `False` prevents the Sun Cluster resource group manager from automatically starting the resource groups in the protection group. Therefore, after the Sun Cluster Geographic Edition software restarts and communicates with the remote cluster to ensure that the remote cluster is running and that the remote cluster is the secondary cluster for that resource group. The Sun Cluster Geographic Edition software does not automatically start the resource group on the primary cluster.

  Application resource groups should be online only on primary cluster when the protection group is activated.

  Set the `Auto_start_on_new_cluster` property to `False` as follows:

  ```
  scrgadm -c -g apprg1 -y Auto_start_on_new_cluster=False
  ```

1 Log in to `phys-newyork-1`.

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see "Sun Cluster Geographic Edition Software and RBAC" in *Sun Cluster Geographic Edition System Administration Guide*.

`phys-newyork-1` is the only node on the secondary cluster. For a reminder of which node is `phys-newyork-1` , see "Example Sun Cluster Geographic Edition Cluster Configuration" in *Sun Cluster Geographic Edition System Administration Guide*.

2 Replicate the protection group configuration to the partner cluster by using the `geopg get` command.

This command retrieves the configuration information of the protection group from the remote cluster and creates the protection group on the local cluster.

```
phys-newyork-1# geopg get -s partnershipname [protectiongroup]
```

`-s` *partnershipname*	Specifies the name of the partnership from which the protection group configuration information should be retrieved and the name of the partnership where the protection will be created locally.
protectiongroup	Specifies the name of the protection group. If no protection group is specified, then all protection groups that exist in the specified partnership on the remote partner are created on the local cluster.

Note – The `geopg get` command replicates Sun Cluster Geographic Edition related entities. For information about how to replicate Sun Cluster entities, see "Replicating and Upgrading Configuration Data for Resource Groups, Resource Types, and Resources" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

Example 2–15 Replicating the Hitachi TrueCopy Protection Group to a Partner Cluster

This example replicates the configuration of `tcpg` from `cluster-paris` to `cluster-newyork`.

```
# rlogin phys-newyork-1 -l root
phys-newyork-1# geopg get -s paris-newyork-ps tcpg
```

Activating a Hitachi TrueCopy Protection Group

When you activate a protection group, the protection group assumes the role that you assigned to it during configuration. For more information about configuring protection groups, see "How to Create and Configure a Hitachi TrueCopy Protection Group That Does Not Use Oracle Real Application Clusters" on page 28.

You can activate a protection group in the following ways:

- Globally – Activates a protection group on both clusters where the protection group is configured.
- On the primary cluster only – Secondary cluster remains inactive.

- On the secondary cluster only – Primary cluster remains inactive.

Activating a Hitachi TrueCopy protection group on a cluster has the following effect on the data replication layer:

- The data replication configuration of the protection group is validated. During validation, the current local role of a protection group is compared with the aggregate device group state as described in Table 2–3. If validation is successful, data replication is started.
- Data replication is started on the data replication device groups that are configured for the protection group, no matter whether the activation occurs on a primary or secondary cluster. Data is always replicated from the cluster on which the local role of the protection group is `primary` to the cluster on which the local role of the protection group is `secondary`.

Application handling proceeds only after data replication has been started successfully.

Activating a protection group has the following effect on the application layer:

- When a protection group is activated on the primary cluster, the application resource groups that are configured for the protection group are also started.
- When a protection group is activated on the secondary cluster, the application resource groups are *not* started.

The Hitachi TrueCopy command that is used to start data replication depends on the following factors:

- Aggregate device group state
- Local role of the protection group
- Current pair state

The following table describes the Hitachi TrueCopy command that is used to start data replication for each of the possible combinations of factors. In the commands, `dg` is the device group name and `fl` is the fence level that is configured for the device group.

TABLE 2–4 Commands Used to Start Hitachi TrueCopy Data Replication

Aggregate Device Group State	Valid Local Protection Group Role	Hitachi TrueCopy Start Command

`SMPL`	`primary` or `secondary`	`paircreate -vl -g dg -f fl`
		`paircreate -vr -g dg -f fl`
		Both commands require that the `horcmd` process is running on the remote cluster.

TABLE 2–4 Commands Used to Start Hitachi TrueCopy Data Replication *(Continued)*

Aggregate Device Group State	Valid Local Protection Group Role	Hitachi TrueCopy Start Command
`Regular Primary`	`primary`	If the local state code is 22, 23, 25, 26, 29, 42, 43, 45, 46, or 47, no command is run because data is already being replicated. If the local state code is 24, 44, or 48, then the following command is run: `pairresync -g dg [-l]`. If the local state code is 11, then the following command is run: `paircreate -vl -g dg -f fl`. Both commands require that the `horcmd` process is running on the remote cluster.
`Regular Secondary`	`secondary`	If the local state code is 32, 33, 35, 36, 39, 52, 53, 55, 56, or 57, no command is run because data is already being replicated. If the local state code is 34, 54, or 58, then the following command is run: `pairresync -g dg` If the local state code is 11, the following command is run: `paircreate -vr -g dg -f fl` Both commands require that the `horcmd` process is up on the remote cluster.
`Takeover Primary`	`primary`	If the local state code is 34 or 54, the following command is run: `pairresync -swaps -g`. If the local state code is 11, then the following command is run: `paircreate -vl -g dg -f fl`. The `paircreate` command requires that the `horcmd` process is running on the remote cluster.
`Takeover Secondary`	`secondary`	If the local state code is 24, 44, 25, or 45, the following command is run: `pairresync -swapp -g dg`. If the local state code is 11, the following command is run: `paircreate -vr -g dg -f fl`. Both commands require that the `horcmd` process is running on the remote cluster.

▼ How to Activate a Hitachi TrueCopy Protection Group

1 Log in to a cluster node.

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see "Sun Cluster Geographic Edition Software and RBAC" in *Sun Cluster Geographic Edition System Administration Guide*.

2 Activate the protection group.

When you activate a protection group, its application resource groups are also brought online.

```
# geopg start -e scope [-n] protectiongroupname
```

`-e` *scope* — Specifies the scope of the command.

If the scope is `Local`, then the command operates on the local cluster only. If the scope is `Global`, the command operates on both clusters that deploy the protection group.

Note – The property values, such as `Global` and `Local`, are *not* case sensitive.

`-n` — Prevents the start of data replication at protection group startup.

If you omit this option, the data replication subsystem starts at the same time as the protection group.

protectiongroupname — Specifies the name of the protection group.

The `geopg start` command uses the `scswitch -Z -g` *resourcegroups* command to bring resource groups and resources online. For more information about using this command, see the `scswitch`(1M) man page.

Example 2–16 How the Sun Cluster Geographic Edition Software Issues the Command to Start Replication

This example illustrates how the Sun Cluster Geographic Edition determines the Hitachi TrueCopy command that is used to start data replication.

First, the Hitachi TrueCopy protection group is created.

```
phys-paris-1# geopg create -s paris-newyork-ps -o primary -d truecopy tcpg
```

A device group, `devgroup1`, is added to the protection group.

```
phys-paris-1# geopg add-device-group -p fence_level=async devgroup1 tcpg
```

The current state of a Hitachi TrueCopy device group, `devgroup1`, is provided in the output of the `pairdisplay` command:

```
phys-paris-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-A , 0, 1) 12345   1..SMPL ----  ----, ----- ----   -
devgroup1 pair1(R) (CL1-C , 0, 20)54321 609..SMPL ----  ----, ----- ----   -
devgroup1 pair2(L) (CL1-A , 0, 2) 12345   2..SMPL ----  ----, ----- ----   -
devgroup1 pair2(R) (CL1-C , 0,21) 54321 610..SMPL ----  ----, ----- ----   -
```

The aggregate device group state is `SMPL`.

Next, the protection group, `tcpg`, is activated by using the `geopg start` command.

```
phys-paris-1# geopg start -e local tcpg
```

The Sun Cluster Geographic Edition software runs the `paircreate -g devgroup1 -vl -f async` command at the data replication level. If the command is successful, the state of `devgroup1` is provided in the output of the `pairdisplay` command:

```
phys-paris-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-A , 0, 1) 12345   1..P-VOL COPY ASYNC,54321  609   -
devgroup1 pair1(R) (CL1-C , 0, 20)54321 609..S-VOL COPY ASYNC,-----   1    -
devgroup1 pair2(L) (CL1-A , 0, 2) 12345   2..P-VOL COPY ASYNC,54321  610   -
devgroup1 pair2(R) (CL1-C , 0,21) 54321 610..S-VOL COPY ASYNC,-----   2    -
```

Example 2–17 Activating a Hitachi TrueCopy Protection Group Globally

This example activates a protection group globally.

```
# geopg start -e global tcpg
```

The protection group, `tcpg`, is activated on both clusters where the protection group is configured.

Example 2–18 Activating a Hitachi TrueCopy Protection Group Locally

This example activates a protection group on a local cluster only. This local cluster might be a primary cluster or a secondary cluster, depending on the role of the cluster.

```
# geopg start -e local tcpg
```

Deactivating a Hitachi TrueCopy Protection Group

You can deactivate a protection group on the following levels:

- Globally – Deactivates a protection group on both clusters where the protection group is configured
- On the primary cluster only – Secondary cluster remains active
- On the secondary cluster only – Primary cluster remains active

Deactivating a Hitachi TrueCopy protection group on a cluster has the following effect on the data replication layer:

- The data replication configuration of the protection group is validated. During validation, the current local role of the protection group is compared with the aggregate device group state as described in Table 2–3. If validation is successful, data replication is stopped.
- Data replication is stopped on the data replication device groups that are configured for the protection group, whether the deactivation occurs on a primary or secondary cluster.

Deactivating a protection group has the following effect on the application layer:

- When a protection group is deactivated on the primary cluster, all of the application resource groups that are configured for the protection group are stopped and unmanaged.
- When a protection group is deactivated on the secondary cluster, the resource groups on the secondary cluster are not affected. Application resource groups that are configured for the protection group might remain active on the primary cluster, depending on the activation state of the primary cluster.

The Hitachi TrueCopy command that is used to stop data replication depends on the following factors:

- Aggregate device group state
- Local role of the protection group
- Current pair state

The following table describes the Hitachi TrueCopy command used to stop data replication for each of the possible combinations of factors. In the commands, `dg` is the device group name.

TABLE 2–5 Commands Used to Stop Hitachi TrueCopy Data Replication

Aggregate Device Group State	Valid Local Protection Group Role	Hitachi TrueCopyStop Command
`SMPL`	`primary` or `secondary`	No command is run because no data is being replicated.

TABLE 2–5 Commands Used to Stop Hitachi TrueCopy Data Replication *(Continued)*

Aggregate Device Group State	Valid Local Protection Group Role	Hitachi TrueCopyStop Command
`Regular Primary`	`primary`	If the local state code is 22, 23, 26, 29, 42, 43, 46, or 47, then the following command is run: `pairsplit -g dg [-l]`. If the local state code is 11, 24, 25, 44, 45, or 48, then no command is run because no data is being replicated.
`Regular Secondary`	`secondary`	If the local state code is 32, 33, 35, 36, 39, 52, 53, 55, 56, or 57, the following command is run: `pairsplit -g dg`. If the local state code is 33 or 53 and the remote state is `PSUE`, no command is run to stop replication. If the local state code is 11, 34, 54, or 58, then no command is run because no data is being replicated.
`Takeover Primary`	`primary`	No command is run because no data is being replicated.
`Takeover Secondary`	`secondary`	No command is run because no data is being replicated.

▼ How to Deactivate a Hitachi TrueCopy Protection Group

1 Log in to a cluster node.

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see "Sun Cluster Geographic Edition Software and RBAC" in *Sun Cluster Geographic Edition System Administration Guide*.

2 Deactivate the protection group.

When you deactivate a protection group, its application resource groups are also unmanaged.

```
# geopg stop -e scope [-D] protectiongroupname
```

`-e` *scope* — Specifies the scope of the command.

If the scope is `Local`, then the command operates on the local cluster only. If the scope is `Global`, the command operates on both clusters where the protection group is deployed.

	Note – The property values, such as `Global` and `Local`, are *not* case sensitive.
`-D`	Specifies that only data replication should be stopped and the protection group should be online. If you omit this option, the data replication subsystem and the protection group are both stopped.
protectiongroupname	Specifies the name of the protection group.

Example 2–19 How the Sun Cluster Geographic Edition Software Issues the Command to Stop Replication

This example illustrates how the Sun Cluster Geographic Edition software determines the Hitachi TrueCopy command that is used to stop data replication.

The current state of the Hitachi TrueCopy device group, `devgroup1`, is provided in the output of the `pairdisplay` command:

```
phys-paris-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-A , 0, 1) 12345   1..P-VOL PAIR ASYNC,54321  609   -
devgroup1 pair1(R) (CL1-C , 0, 20)54321 609..S-VOL PAIR ASYNC,-----   1    -
devgroup1 pair2(L) (CL1-A , 0, 2) 12345   2..P-VOL PAIR ASYNC,54321  610   -
devgroup1 pair2(R) (CL1-C , 0,21) 54321 610..S-VOL PAIR ASYNC,-----   2    -
```

A device group, `devgroup1`, is added to the protection group as follows:

```
phys-paris-1# geopg add-device-group -p fence_level=async devgroup1 tcpg
```

The Sun Cluster Geographic Edition software runs the `pairvolchk -g <DG> -ss` command at the data replication level, which returns a value of `43`.

```
pairvolchk -g devgroup1 -ss
Volstat is P-VOL.[status = PAIR fence = ASYNC]
phys-paris-1# echo $?
43
```

Next, the protection group, `tcpg`, is deactivated by using the `geopg stop` command.

```
phys-paris-1# geopg stop -s local tcpg
```

The Sun Cluster Geographic Edition software runs the `pairsplit -g devgroup1` command at the data replication level.

If the command is successful, the state of `devgroup1` is provided in the output of the `pairdisplay` command:

```
phys-paris-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-A , 0, 1) 12345   1..P-VOL PSUS ASYNC,54321  609   -
devgroup1 pair1(R) (CL1-C , 0, 20)54321 609..S-VOL SSUS ASYNC,-----   1    -
devgroup1 pair2(L) (CL1-A , 0, 2) 12345   2..P-VOL PSUS ASYNC,54321  610   -
devgroup1 pair2(R) (CL1-C , 0,21) 54321 610..S-VOL SSUS ASYNC,-----   2    -
```

Example 2–20 Deactivating a Protection Group on All Clusters

This example deactivates a protection group on all clusters.

```
# geopg stop -e global tcpg
```

Example 2–21 Deactivating a Protection Group on a Local Cluster

This example deactivates a protection group on the local cluster.

```
# geopg stop -e local tcpg
```

Example 2–22 Stopping Data Replication While Leaving the Protection Group Online

This example stops only data replication on a local cluster.

```
# geopg stop -e local -D tcpg
```

If the administrator decides later to deactivate both the protection group and its underlying data replication subsystem, the administrator can rerun the command without the `-D` option:

```
# geopg stop -e local tcpg
```

Example 2–23 Deactivating a Hitachi TrueCopy Protection Group While Keeping Application Resource Groups Online

This example keeps two application resource groups, `apprg1` and `apprg2`, online while deactivating their protection group, `tcpg`, on both clusters.

1. Remove the application resource groups from the protection group.

   ```
   # geopg remove-resource-group apprg1,apprg2 tcpg
   ```

2. Deactivate the protection group.

```
# geopg stop -e global tcpg
```

Resynchronizing a Hitachi TrueCopy Protection Group

You can resynchronize the configuration information of the local protection group with the configuration information that is retrieved from the partner cluster. You need to resynchronize a protection group when its `Synchronization` status in the output of the `geoadm status` command is `Error`.

For example, you might need to resynchronize protection groups after booting the cluster. For more information, see "Booting a Cluster" in *Sun Cluster Geographic Edition System Administration Guide*.

Resynchronizing a protection group updates only entities that are related to Sun Cluster Geographic Edition software. For information about how to update Sun Cluster entities, see "Replicating and Upgrading Configuration Data for Resource Groups, Resource Types, and Resources" in *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

▼ How to Resynchronize a Protection Group

Before You Begin The protection group must be deactivated on the cluster where you are running the `geopg update` command. For information about deactivating a protection group, see "Deactivating a Hitachi TrueCopy Protection Group" on page 54.

1 Log in to a cluster node.

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see "Sun Cluster Geographic Edition Software and RBAC" in *Sun Cluster Geographic Edition System Administration Guide*.

2 Resynchronize the protection group.

`# geopg update` *protectiongroupname*

protectiongroupname Specifies the name of the protection group

Example 2–24 Resynchronizing a Protection Group

This example resynchronizes a protection group.

```
# geopg update tcpg
```

Checking the Runtime Status of Hitachi TrueCopy Data Replication

You can obtain an overall view of the status of replication, as well as a more detailed runtime status of the Hitachi TrueCopy replication resource groups. The following sections describe the procedures for checking each status.

Displaying a Hitachi TrueCopy Runtime Status Overview

The status of each Hitachi TrueCopy data replication resource indicates the status of replication on a particular device group. The status of all the resources under a protection group are aggregated in the replication status. This replication status is the second component of the protection group state. For more information about the states of protection groups, refer to "Monitoring the Runtime Status of the Sun Cluster Geographic Edition Software" in *Sun Cluster Geographic Edition System Administration Guide*.

To view the overall status of replication, look at the protection group state as described in the following procedure.

▼ How to Check the Overall Runtime Status of Replication

1 Access a node of the cluster where the protection group has been defined.

You must be assigned the Basic Solaris User RBAC rights profile to complete this procedure. For more information about RBAC, see "Sun Cluster Geographic Edition Software and RBAC" in *Sun Cluster Geographic Edition System Administration Guide*.

2 Check the runtime status of replication.

```
# geoadm status
```

Refer to the `Protection Group` section of the output for replication information. The information that is displayed by this command includes the following:

- Whether the local cluster is enabled for partnership participation
- Whether the local cluster is involved in a partnership
- Status of the heartbeat configuration
- Status of the defined protection groups
- Status of current transactions

3 Check the runtime status of data replication for each Hitachi TrueCopy device group.

```
# scstat -g
```

Refer to the `Status` and `Status Message` fields for the data replication device group you want to check.

See Also For more information about these fields, see Table 2–6.

Displaying a Detailed Hitachi TrueCopy Runtime Status

The Sun Cluster Geographic Edition software internally creates and maintains one replication resource group for each protection group. The name of the replication resource group has the following format:

`rg-tc_`*truecopyprotectiongroupname*

If you add a Hitachi TrueCopy device group to a protection group, Sun Cluster Geographic Edition software creates a resource for each device group. This resource monitors the status of replication for its device group. The name of each resource has the following format:

`r-tc-`*truecopyprotectiongroupname*`-`*truecopydevicegroupname*

You can monitor the status of replication of this device group by checking the `Status` and `Status Message` of this resource. Use the `scstat -g` command to display the resource status and the status message.

The following table describes the `Status` and `Status Message` values that are returned by the `scstat -g` command when the `State` of the Hitachi TrueCopy replication resource group is `Online`.

TABLE 2–6 State and Status Messages of an Online Hitachi TrueCopy Replication Resource Group

Status	Status Message
Online	P-Vol/S-Vol:PAIR
Online	P-Vol/S-Vol:PAIR:Remote horcmd not reachable
Online	P-Vol/S-Vol:PFUL
Online	P-Vol/S-Vol:PFUL:Remote horcmd not reachable
Degraded	SMPL:SMPL
Degraded	SMPL:SMPL:Remote horcmd not reachable
Degraded	P-Vol/S-Vol:COPY
Degraded	P-Vol/S-Vol:COPY:Remote horcmd not reachable

TABLE 2–6 State and Status Messages of an Online Hitachi TrueCopy Replication Resource Group *(Continued)*

Status	Status Message
Degraded	P-Vol/S-Vol:PSUS
Degraded	P-Vol/S-Vol:PSUS:Remote horcmd not reachable
Degraded	P-Vol/S-Vol:PFUS
Degraded	P-Vol/S-Vol:PFUS:Remote horcmd not reachable
Faulted	P-Vol/S-Vol:PDFUB
Faulted	P-Vol/S-Vol:PDUB:Remote horcmd not reachable
Faulted	P-Vol/S-Vol:PSUE
Faulted	P-Vol/S-Vol:PSUE:Remote horcmd not reachable
Degraded	S-Vol:SSWS:Takeover Volumes
Faulted	P-Vol/S-Vol:Suspicious role configuration. Actual Role=x, Config Role=y

For more information about these values, refer to the Hitachi TrueCopy documentation.

For more information about the `scstat` command, see the `scstat`(1M) man page.

◆ ◆ ◆ CHAPTER 3

Migrating Services That Use Hitachi TrueCopy Data Replication

This chapter provides information about migrating services for maintenance or as a result of cluster failure. This chapter contains the following sections:

Detecting Cluster Failure on a System That Uses Hitachi TrueCopy Data Replication

This section describes the internal processes that occur when failure is detected on a primary or a secondary cluster.

Detecting Primary Cluster Failure

When the primary cluster for a given protection group fails, the secondary cluster in the partnership detects the failure. The cluster that fails might be a member of more than one partnership, resulting in multiple failure detections.

The following actions take place when a primary cluster failure occurs. During a failure, the appropriate protection groups are in the `Unknown` state.

- Heartbeat failure is detected by a partner cluster.

- The heartbeat is activated in emergency mode to verify that the heartbeat loss is not transient and that the primary cluster has failed. The heartbeat remains in the `Online` state during this default timeout interval, while the heartbeat mechanism continues to retry the primary cluster.

  This query interval is set by using the `Query_interval` heartbeat property. If the heartbeat still fails after the interval you configured, a heartbeat-lost event is generated and logged in the system log. When you use the default interval, the emergency-mode retry behavior might delay heartbeat-loss notification for about nine minutes. Messages are displayed in the graphical user interface (GUI) and in the output of the `geoadm status` command.

  For more information about logging, see "Viewing the Sun Cluster Geographic Edition Log Messages" in *Sun Cluster Geographic Edition System Administration Guide*.

Detecting Secondary Cluster Failure

When a secondary cluster for a given protection group fails, a cluster in the same partnership detects the failure. The cluster that failed might be a member of more than one partnership, resulting in multiple failure detections.

During failure detection, the following actions occur:

- Heartbeat failure is detected by a partner cluster.
- The heartbeat is activated in emergency mode to verify that the secondary cluster is dead.
- The cluster notifies the administrator. The system detects all protection groups for which the cluster that failed was acting as secondary. The state of the appropriate protection groups is marked `Unknown`.

Migrating Services That Use Hitachi TrueCopy Data Replication With a Switchover

Perform a switchover of a Hitachi TrueCopy protection group when you want to migrate services to the partner cluster in an orderly fashion. A switchover consists of the following:

- Application services are offline on the former primary cluster, `cluster-paris`.

  For a reminder of which cluster is `cluster-paris`, see "Example Sun Cluster Geographic Edition Cluster Configuration" in *Sun Cluster Geographic Edition System Administration Guide*.
- The data replication role is reversed and now continues to run from the new primary, `cluster-newyork`, to the former primary, `cluster-paris`.
- Application services are brought online on the new primary cluster, `cluster-newyork`.

Validations That Occur Before a Switchover

When a switchover is initiated by using the `geopg switchover` command, the data replication subsystem runs several validations on both clusters. The switchover is performed only if the validation step succeeds on both clusters.

First, the replication subsystem checks that the Hitachi TrueCopy device group is in a valid aggregate device group state. Then, it checks that the local device group states on the target primary cluster, `cluster-newyork`, are 23, 33, 43, or 53. The local device group state is returned by the `pairvolchk -g device-group-name -ss` command. These values correspond to a `PVOL_PAIR` or `SVOL_PAIR` state. The Hitachi TrueCopy commands that are run on the new primary cluster, `cluster-newyork`, are described in the following table.

TABLE 3–1 Hitachi TrueCopy Switchover Validations on the New Primary Cluster

Aggregate Device Group State	**Valid Device Group State on Local Cluster**	**Hitachi TrueCopy Switchover Commands That Are Run on `cluster-newyork`**
SMPL	None	None
Regular primary	23, 43	No command is run, because the Hitachi TrueCopy device group is already in the `PVOL_PAIR` state.
Regular secondary	33, 53	`horctakeover -g dg [-t]` The `-t` option is specified when the `fence_level` of the Hitachi TrueCopy device group is `async`. The value is calculated as 80% of the `Timeout` property of the protection group. For example, if the protection group has a `Timeout` of 200 seconds, the value of `-t` used in this command is 80% of 200 seconds, or 160 seconds.
Takeover primary	None	None
Takeover secondary	None	None

Results of a Switchover From a Replication Perspective

After a successful switchover, at the data replication level the roles of the primary and secondary volumes have been switched. The `PVOL_PAIR` volumes that were in place before the switchover become the `SVOL_PAIR` volumes. The `SVOL_PAIR` volumes in place before the switchover become the `PVOL_PAIR` volumes. Data replication will continue from the new `PVOL_PAIR` volumes to the new `SVOL_PAIR` volumes.

The `Local-role` property of the protection group is also switched regardless of whether the application could be brought online on the new primary cluster as part of the switchover operation. On the cluster on which the protection group had a `Local role` of `Secondary`, the `Local-role` property of the protection group becomes `Primary`. On the cluster on which the protection group had a `Local-role` of `Primary`, the `Local-role` property of the protection group becomes `Secondary`.

▼ How to Switch Over a Hitachi TrueCopy Protection Group From Primary to Secondary

Before You Begin For a successful switchover, data replication must be active between the primary and the secondary clusters and data volumes on the two clusters must be synchronized.

Before you switch over a protection group from the primary cluster to the secondary cluster, ensure that the following conditions are met:

- The Sun Cluster Geographic Edition software is running on the both clusters.
- The secondary cluster is a member of a partnership.
- Both cluster partners can be reached.
- The protection group is in the `OK` state.

Caution – If you have configured the `Cluster_dgs` property, only applications that belong to the protection group can write to the device groups specified in the `Cluster_dgs` property.

1 Log in to a cluster node.

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see "Sun Cluster Geographic Edition Software and RBAC" in *Sun Cluster Geographic Edition System Administration Guide*.

2 Initiate the switchover.

The application resource groups that are a part of the protection group are stopped and started during the switchover.

```
# geopg switchover [-f] -m newprimarycluster protectiongroupname
```

`-f`	Forces the command to perform the operation without asking you for confirmation
`-m` *newprimarycluster*	Specifies the name of the cluster that is to be the new primary cluster for the protection group
protectiongroupname	Specifies the name of the protection group

Example 3–1 Forcing a Switchover From Primary to Secondary

This example performs a switchover to the secondary cluster.

```
# geopg switchover -f -m cluster-newyork tcpg
```

Forcing a Takeover on a System That Uses Hitachi TrueCopy Data Replication

Perform a takeover when applications need to be brought online on the secondary cluster regardless of whether the data is completely consistent between the primary volume and the secondary volume. The information in this section assumes that the protection group has been started.

The following steps occur after a takeover is initiated:

- If the former primary cluster, `cluster-paris`, can be reached and the protection group is not locked for notification handling or some other reason, the application services are taken offline on the former primary cluster.

  For a reminder of which cluster is `cluster-paris`, see "Example Sun Cluster Geographic Edition Cluster Configuration" in *Sun Cluster Geographic Edition System Administration Guide*.
- Data volumes of the former primary cluster, `cluster-paris`, are taken over by the new primary cluster, `cluster-newyork`.

  Note – This data might not be consistent with the original primary volumes. After the takeover, data replication from the new primary cluster, `cluster-newyork`, to the former primary cluster, `cluster-paris`, is stopped.
- Application services are brought online on the new primary cluster, `cluster-newyork`.

For details about the possible conditions of the primary and secondary cluster before and after takeover, see Appendix C, "Takeover Postconditions," in *Sun Cluster Geographic Edition System Administration Guide*.

The following sections describe the steps you must perform to force a takeover by a secondary cluster.

Validations That Occur Before a Takeover

When a takeover is initiated by using the `geopg takeover` command, the data replication subsystem runs several validations on both clusters. These steps are conducted on the original primary cluster only if the primary cluster can be reached. If validation on the original primary cluster fails, the takeover still occurs.

First, the replication subsystem checks that the Hitachi TrueCopy device group is in a valid aggregate device group state. Then, the replication subsystem checks that the local device group states on the target primary cluster, `cluster-newyork`, are not 32 or 52. These values correspond to a `SVOL_COPY` state, for which the `horctakeover` command fails. The Hitachi TrueCopy commands that are used for the takeover are described in the following table.

TABLE 3–2 Hitachi TrueCopy Takeover Validations on the New Primary Cluster

Aggregate Device Group State	Valid Local State Device Group State	Hitachi TrueCopy Takeover Commands That Are Run on `cluster-newyork`
SMPL	All	No command is run.
Regular primary	All	No command is run.
Regular secondary	All Regular secondary states except 32 or 52 For a list of Regular secondary states, refer to Table 2–1 and Table 2–2.	`horctakeover -S -g dg [-t]` The `-t` option is given when the `fence_level` of the Hitachi TrueCopy device group is `async`. The value is calculated as 80% of the `Timeout` property of the protection group. For example, if the protection group has a `Timeout` of 200 seconds, the value of `-t` used in this command will be 80% of 200 seconds, or 160 seconds.
Takeover primary	All	No command is run.
Takeover secondary	All	`pairsplit -R-g dg pairsplit -S-g dg`

Results of a Takeover From a Replication Perspective

From a replication perspective, after a successful takeover, the `Local-role` property of the protection group is changed to reflect the new role, it is immaterial whether the application could be brought online on the new primary cluster as part of the takeover operation. On `cluster-newyork`, where the protection group had a `Local-role` of `Secondary`, the `Local-role` property of the protection group becomes `Primary`. On `cluster-paris`, where the protection group had a `Local-role` of `Primary`, the following might occur:

- If the cluster can be reached, the `Local-role` property of the protection group becomes `Secondary`.
- If the cluster cannot be reached, the `Local-role` property of the protection group remains `Primary`.

If the takeover is successful, the applications are brought online. You do not need to run a separate `geopg start` command.

Caution – After a successful takeover, data replication between the new primary cluster, `cluster-newyork`, and the old primary cluster, `cluster-paris`, is stopped. If you want to run a `geopg start` command, you must use the `-n` option to prevent replication from resuming.

▼ How to Force Immediate Takeover of Hitachi TrueCopy Services by a Secondary Cluster

Before You Begin Before you force the secondary cluster to assume the activity of the primary cluster, ensure that the following conditions are met:

- Sun Cluster Geographic Edition software is running on the cluster.
- The cluster is a member of a partnership.
- The `Configuration` status of the protection group is `OK` on the secondary cluster.

1 Log in to a node in the secondary cluster.

You must be assigned the Geo Management RBAC rights profile to complete this procedure. For more information about RBAC, see "Sun Cluster Geographic Edition Software and RBAC" in *Sun Cluster Geographic Edition System Administration Guide*.

2 Initiate the takeover.

```
# geopg takeover [-f] protectiongroupname
```

`-f`	Forces the command to perform the operation without your confirmation
protectiongroupname	Specifies the name of the protection group

Example 3–2 Forcing a Takeover by a Secondary Cluster

This example forces the takeover of `tcpg` by the secondary cluster `cluster-newyork`.

The `phys-newyork-1` cluster is the first node of the secondary cluster. For a reminder of which node is `phys-newyork-1`, see "Example Sun Cluster Geographic Edition Cluster Configuration" in *Sun Cluster Geographic Edition System Administration Guide*.

```
phys-newyork-1# geopg takeover -f tcpg
```

Next Steps For information about the state of the primary and secondary clusters after a takeover, see Appendix C, "Takeover Postconditions," in *Sun Cluster Geographic Edition System Administration Guide*.

Recovering Services to a Cluster on a System That Uses Hitachi TrueCopy Replication

After a successful takeover operation, the secondary cluster, `cluster-newyork`, becomes the primary for the protection group and the services are online on the secondary cluster. After the recovery of the original primary cluster,`cluster-paris`, the services can be brought online again on the original primary by using a process called *failback*.

Sun Cluster Geographic Edition software supports the following two kinds of failback:

- **Failback-switchover.** During a failback-switchover, applications are brought online again on the original primary cluster, `cluster-paris`, after the data of the original primary cluster was resynchronized with the data on the secondary cluster, `cluster-newyork`.

  For a reminder of which clusters are `cluster-paris` and `cluster-newyork`, see "Example Sun Cluster Geographic Edition Cluster Configuration" in *Sun Cluster Geographic Edition System Administration Guide*.
- **Failback-takeover.** During a failback-takeover, applications are brought online again on the original primary cluster, `cluster-paris`, and use the current data on the original primary cluster. Any updates that occurred on the secondary cluster, `cluster-newyork`, while it was acting as primary are discarded.

You can perform a fail

To continue using the new primary, `cluster-newyork`, as the primary cluster and the original primary cluster, `cluster-paris`, as the secondary after the original primary is running again, resynchronize and revalidate the protection group configuration without performing a switchover or takeover.

▼ How to Resynchronize and Revalidate the Protection Group Configuration

Use this procedure to resynchronize and revalidate data on the original primary cluster, `cluster-paris`, with the data on the current primary cluster, `cluster-newyork`.

Before You Begin

Before you resynchronize and revalidate the protection group configuration, a takeover has occurred on `cluster-newyork`. The clusters now have the following roles:

- If the original primary cluster, `cluster-paris`, has been down, confirm that the cluster is booted and that the Sun Cluster Geographic Edition infrastructure is enabled on the cluster. For more information about booting a cluster, see "Booting a Cluster" in *Sun Cluster Geographic Edition System Administration Guide*.
- The protection group on `cluster-newyork` has the `primary` role.

- The protection group on `cluster-paris` has either the `primary` role or `secondary` role, depending on whether `cluster-paris` could be reached during the takeover from `cluster-newyork`.

1 Resynchronize the original primary cluster, `cluster-paris`, with the current primary cluster, `cluster-newyork`.

`cluster-paris` forfeits its own configuration and replicates the `cluster-newyork` configuration locally. Resynchronize both the partnership and protection group configurations.

a. On `cluster-paris`, resynchronize the partnership.

```
# geops update partnershipname
```

partnershipname Specifies the name of the partnership

Note – You need to perform this step only once, even if you are resynchronizing multiple protection groups.

For more information about synchronizing partnerships, see "Resynchronizing a Partnership" in *Sun Cluster Geographic Edition System Administration Guide*.

b. On `cluster-paris`, resynchronize each protection group.

Because the role of the protection group on `cluster-newyork` is `primary`, this step ensures that the role of the protection group on `cluster-paris` is `secondary`.

```
# geopg update protectiongroupname
```

protectiongroupname Specifies the name of the protection group

For more information about synchronizing protection groups, see "Resynchronizing a Hitachi TrueCopy Protection Group" on page 58.

2 On `cluster-paris`, validate the cluster configuration for each protection group.

```
# geopg validate protectiongroupname
```

protectiongroupname Specifies a unique name that identifies a single protection group

For more information, see "How to Validate a Hitachi TrueCopy Protection Group" on page 35.

3 On `cluster-paris`, activate each protection group.

Because the protection group on `cluster-paris` has a role of secondary, the `geopg start` command does not restart the application on `cluster-paris`.

```
# geopg start -e local protectiongroupname
```

`-e local` Specifies the scope of the command.

By specifying a `local` scope, the command operates on the local cluster only.

protectiongroupname Specifies the name of the protection group.

Caution – Do not use the `-n` option because the data needs to be synchronized from the current primary, `cluster-newyork`, to the current secondary, `cluster-paris`.

Because the protection group has a role of secondary, the data is synchronized from the current primary, `cluster-newyork`, to the current secondary, `cluster-paris`.

For more information about the `geopg start` command, see "How to Activate a Hitachi TrueCopy Protection Group" on page 52.

4 Confirm that the data is completely synchronized.

The state of the protection group on `cluster-newyork` must be `OK`.

```
phys-newyork-1# geoadm status
```

Refer to the `Protection Group` section of the output.

The protection group has a local state of `OK` when the Hitachi TrueCopy device groups on `cluster-newyork` have a state of `PVOL_PAIR` and the Hitachi TrueCopy device groups on `cluster-paris` have a state of `SVOL_PAIR`.

▼ How to Perform a Failback-Switchover on a System That Uses Hitachi TrueCopy Replication

Use this procedure to restart an application on the original primary cluster, `cluster-paris`, after the data on this cluster has been resynchronized with the data on the current primary cluster, `cluster-newyork`.

Note – The failback procedures apply only to clusters in a partnership. You need to perform the following procedure only once per partnership.

Before You Begin Before you perform a failback-switchover, a takeover has occurred on `cluster-newyork`. The clusters have the following roles:

- If the original primary cluster, `cluster-paris`, has been down, confirm that the cluster is booted and that the Sun Cluster Geographic Edition infrastructure is enabled on the cluster. For more information about booting a cluster, see "Booting a Cluster" in *Sun Cluster Geographic Edition System Administration Guide*.
- The protection group on `cluster-newyork` has the `primary` role.
- The protection group on `cluster-paris` has either the `primary` role or `secondary` role, depending on whether `cluster-paris` could be reached during the takeover from `cluster-newyork`.

1 Resynchronize the original primary cluster, `cluster-paris`, with the current primary cluster, `cluster-newyork`.

`cluster-paris` forfeits its own configuration and replicates the `cluster-newyork` configuration locally. Resynchronize both the partnership and protection group configurations.

a. On `cluster-paris`, resynchronize the partnership.

```
phys-paris-1# geops update partnershipname
```

partnershipname Specifies the name of the partnership

Note – You need to perform this step only once per partnership, even if you are performing a failback-switchover for multiple protection groups in the partnership.

For more information about synchronizing partnerships, see "Resynchronizing a Partnership" in *Sun Cluster Geographic Edition System Administration Guide*.

b. On `cluster-paris`, resynchronize each protection group.

Because the local role of the protection group on `cluster-newyork` is now `primary`, this steps ensures that the role of the protection group on `cluster-paris` becomes `secondary`.

```
phys-paris-1# geopg update protectiongroupname
```

protectiongroupname Specifies the name of the protection group

For more information about synchronizing protection groups, see "Resynchronizing a Hitachi TrueCopy Protection Group" on page 58.

2 On `cluster-paris`, validate the cluster configuration for each protection group.

Ensure that the protection group is not in an error state. A protection group cannot be started when it is in an error state.

```
phys-paris-1# geopg validate protectiongroupname
```

protectiongroupname Specifies a unique name that identifies a single protection group

For more information, see "How to Validate a Hitachi TrueCopy Protection Group" on page 35.

3 On `cluster-paris`, activate each protection group.

Because the protection group on `cluster-paris` has a role of `secondary`, the `geopg start` command does not restart the application on `cluster-paris`.

```
phys-paris-1# geopg start -e local protectiongroupname
```

`-e local` Specifies the scope of the command.

By specifying a `local` scope, the command operates on the local cluster only.

protectiongroupname Specifies the name of the protection group.

Caution – Do not use the `-n` option because the data needs to be synchronized from the current primary, `cluster-newyork`, to the current secondary, `cluster-paris`.

Because the protection group has a role of `secondary`, the data is synchronized from the current primary, `cluster-newyork`, to the current secondary, `cluster-paris`.

For more information about the `geopg start` command, see "How to Activate a Hitachi TrueCopy Protection Group" on page 52.

4 Confirm that the data is completely synchronized.

The state of the protection group on `cluster-newyork` must be `OK`.

```
phys-newyork-1# geoadm status
```

Refer to the `Protection Group` section of the output.

The protection group has a local state of `OK` when the Hitachi TrueCopy device groups on `cluster-newyork` have a state of `PVOL_PAIR` and the Hitachi TrueCopy device groups on `cluster-paris` have a state of `SVOL_PAIR`.

5 On either cluster, perform a switchover from `cluster-newyork` to `cluster-paris` for each protection group.

```
# geopg switchover [-f] -m clusterparis protectiongroupname
```

For more information, see "How to Switch Over a Hitachi TrueCopy Protection Group From Primary to Secondary" on page 66.

`cluster-paris` resumes its original role as primary cluster for the protection group.

6 Ensure that the switchover was performed successfully.

Verify that the protection group is now primary on `cluster-paris` and secondary on `cluster-newyork` and that the state for Data replication and Resource groups is `OK` on both clusters.

```
# geoadm status
```

Check the runtime status of the application resource group and data replication for each Hitachi TrueCopy protection group.

```
# scstat -g
```

Refer to the `Status` and `Status Message` fields that are presented for the data replication device group you want to check. For more information about these fields, see Table 2–1.

For more information about the runtime status of data replication see, "Checking the Runtime Status of Hitachi TrueCopy Data Replication" on page 59.

▼ How to Perform a Failback-Takeover on a System That Uses Hitachi TrueCopy Replication

Use this procedure to restart an application on the original primary cluster, `cluster-paris`, and use the current data on the original primary cluster. Any updates that occurred on the secondary cluster, `cluster-newyork`, while it was acting as primary are discarded.

The failback procedures apply only to clusters in a partnership. You need to perform the following procedure only once per partnership.

Note – Conditionally, you can resume using the data on the original primary, `cluster-paris`. You must not have replicated data from the new primary, `cluster-newyork`, to the original primary cluster, `cluster-paris`, at any point after the takeover operation on `cluster-newyork`. To prevent data replication between the new primary and the original primary, you must use the `-n` option when you run the `geopg start` command.

Before You Begin Ensure that the clusters have the following roles:

- The protection group on `cluster-newyork` has the `primary` role.
- The protection group on `cluster-paris` has either the `primary` role or `secondary` role, depending on whether the protection group could be reached during the takeover.

1 Resynchronize the original primary cluster, `cluster-paris`, with the original secondary cluster, `cluster-newyork`.

`cluster-paris` forfeits its own configuration and replicates the `cluster-newyork` configuration locally.

a. On `cluster-paris`, resynchronize the partnership.

```
phys-paris-1# geops update partnershipname
```

partnershipname Specifies the name of the partnership

Note – You need to perform this step only once per partnership, even if you are performing a failback-takeover for multiple protection groups in the partnership.

For more information about synchronizing partnerships, see "Resynchronizing a Partnership" in *Sun Cluster Geographic Edition System Administration Guide*.

b. Place the Hitachi TrueCopy device group, `devgroup1`, in the SMPL state.

Use the `pairsplit` commands to place the Hitachi TrueCopy device groups that are in the protection group on both `cluster-paris` and `cluster-newyork` in the SMPL state. The

`pairsplit` command you use depends on the pair state of the Hitachi TrueCopy device group. The following table gives some examples of the command you need to use on `cluster-paris` for some typical pair states.

Pair State on `cluster-paris`	**Pair State on** `cluster-newyork`	`pairsplit` **Command Used on** `cluster-paris`
PSUS or PSUE	SSWS	`pairsplit -R -g dgname` `pairsplit -S -g dgname`
SSUS	PSUS	`pairsplit -S -g dgname`

For more information about the `pairsplit` commands, see the *Sun StorEdge SE 9900 V Series Command and Control Interface User and Reference Guide*.

If the command is successful, the state of `devgroup1` is provided in the output of the `pairdisplay` command:

```
phys-paris-1# pairdisplay -g devgroup1
Group PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#,P/S,Status,Fence,Seq#,P-LDEV# M
devgroup1 pair1(L) (CL1-A , 0, 1) 12345   1..SMPL ----  ----,-----  ----   -
devgroup1 pair1(R) (CL1-C , 0, 20)54321 609..SMPL ----  ----,-----  ----   -
devgroup1 pair2(L) (CL1-A , 0, 2) 12345   2..SMPL ----  ----,-----  ----   -
devgroup1 pair2(R) (CL1-C , 0,21) 54321 610..SMPL ----  ----,-----  ----   -
```

.

c. On `cluster-paris`, resynchronize each protection group.

```
phys-paris-1# geopg update protectiongroupname
```

protectiongroupname Specifies the name of the protection group

For more information about resynchronizing protection groups, see "How to Resynchronize a Protection Group" on page 58.

2 On `cluster-paris`, validate the configuration for each protection group.

Ensure that the protection group is not in an error state. A protection group cannot be started when it is in a error state.

```
phys-paris-1# geopg validate protectiongroupname
```

protectiongroupname Specifies a unique name that identifies a single protection group

For more information, see "How to Validate a Hitachi TrueCopy Protection Group" on page 35.

3 On `cluster-paris`, activate each protection group in the secondary role *without* data replication.

Because the protection group on `cluster-paris` has a role of `secondary`, the `geopg start` command does not restart the application on `cluster-paris`.

```
phys-paris-1# geopg start -e local -n protectiongroupname
```

`-e local` Specifies the scope of the command

.

By specifying a `local` scope, the command operates on the local cluster only.

`-n` Prevents the start of data replication at protection group startup.

Note – You must use the `-n` option.

---

protectiongroupname Specifies the name of the protection group.

For more information, see "How to Activate a Hitachi TrueCopy Protection Group" on page 52.

Replication from `cluster-newyork` to `cluster-paris` is not started because the `-n` option is used on `cluster-paris`.

4 On `cluster-paris`, initiate a takeover for each protection group.

```
phys-paris-1# geopg takeover [-f] protectiongroupname
```

`-f` Forces the command to perform the operation without your confirmation

protectiongroupname Specifies the name of the protection group

For more information about the `geopg takeover` command, see "How to Force Immediate Takeover of Hitachi TrueCopy Services by a Secondary Cluster" on page 69.

The protection group on `cluster-paris` now has the `primary` role, and the protection group on `cluster-newyork` has the role of `secondary`. The application services are now online on `cluster-paris`.

5 On `cluster-newyork`, activate each protection group.

At the end of step 4, the local state of the protection group on `cluster-newyork` is `Offline`. To start monitoring the local state of the protection group, you must activate the protection group on `cluster-newyork`.

Because the protection group on `cluster-newyork` has a role of `secondary`, the `geopg start` command does not restart the application on `cluster-newyork`.

```
phys-newyork-1# geopg start -e local [-n] protectiongroupname
```

`-e local` Specifies the scope of the command.

By specifying a `local` scope, the command operates on the local cluster only.

`-n` Prevents the start of data replication at protection group startup.

If you omit this option, the data replication subsystem starts at the same time as the protection group.

protectiongroupname Specifies the name of the protection group.

For more information about the `geopg start` command, see "How to Activate a Hitachi TrueCopy Protection Group" on page 52.

6 Ensure that the takeover was performed successfully.

Verify that the protection group is now primary on `cluster-paris` and secondary on `cluster-newyork` and that the state for "Data replication" and "Resource groups" is `OK` on both clusters.

```
# geoadm status
```

Check the runtime status of the application resource group and data replication for each Hitachi TrueCopy protection group.

```
# scstat -g
```

Refer to the `Status` and `Status Message` fields that are presented for the data replication device group you want to check. For more information about these fields, see Table 2–1.

For more information about the runtime status of data replication, see "Checking the Runtime Status of Hitachi TrueCopy Data Replication" on page 59.

Recovering From a Switchover Failure on a System That Uses Hitachi TrueCopy Replication

When you run the `geopg switchover` command, the `horctakeover` command runs at the Hitachi TrueCopy data replication level. If the `horctakeover` command returns a value of `1`, the switchover is successful.

In Hitachi TrueCopy terminology, a switchover is called a *swap-takeover*. In some cases, the `horctakeover` command might not be able to perform a swap-takeover. In these cases, a return value other than 1 is returned, which is considered a switchover failure.

Note – In a failure, the `horctakeover` command usually returns a value of `5`, which indicates a SVOL-SSUS-takeover.

One reason the `horctakeover` command might fail to perform a swap-takeover is because the data replication link, ESCON/FC, is down.

Any result other than a swap-takeover implies that the secondary volumes might not be fully synchronized with the primary volumes. Sun Cluster Geographic Edition software does not start the applications on the new intended primary cluster in a switchover failure scenario.

The remainder of this section describes the initial conditions that lead to a switchover failure and how to recover from a switchover failure.

Switchover Failure Conditions

This section describes a switchover failure scenario. In this scenario, `cluster-paris` is the original primary cluster and `cluster-newyork` is the original secondary cluster.

A switchover switches the services from `cluster-paris` to `cluster-newyork` as follows:

```
phys-newyork-1# geopg switchover -f -m cluster-newyork tcpg
```

While processing the `geopg switchover` command, the `horctakeover` command performs an SVOL-SSUS-takeover and returns a value of 5 for the Hitachi TrueCopy device group, `devgroup1`. As a result, the `geopg switchover` command returns with the following failure message:

```
Processing operation.... this may take a while ....
"Switchover" failed for the following reason:
            Switchover failed for Truecopy DG devgroup1
```

After this failure message has been issued, the two clusters are in the following states:

```
cluster-paris:
        tcpg role: Secondary
cluster-newyork:
        tcpg role: Secondary

phys-newyork-1# pairdisplay -g devgroup1 -fc
Group  PairVol(L/R) (Port#,TID,LU),Seq#,LDEV#.P/S, Status,Fence,%, P-LDEV# M
devgroup1 pair1(L) (CL1-C , 0, 20)12345 609..S-VOL SSWS  ASYNC,100   1    -
devgroup1 pair1(R) (CL1-A , 0, 1) 54321   1..P-VOL PSUS  ASYNC,100  609   -
```

Recovering From Switchover Failure

This section describes procedures to recover from the failure scenario described in the previous section. These procedures bring the application online on the appropriate cluster.

1. Place the Hitachi TrueCopy device group, `devgroup1`, in the SMPL state.

   Use the `pairsplit` commands to place the device groups that are in the protection group on both `cluster-paris` and `cluster-newyork` in the SMPL state. For the pair states that are shown in the previous section, run the following `pairsplit` commands:

   ```
   phys-newyork-1# pairsplit -R -g devgroup1
   phys-newyork-1# pairsplit -S -g devgroup1
   ```

2. Designate one of the clusters `Primary` for the protection group.

   Designate the original primary cluster, `cluster-paris`, `Primary` for the protection group if you intend to start the application on the original primary cluster. The application uses the current data on the original primary cluster.

   Designate the original secondary cluster, `cluster-newyork`, `Primary` for the protection group if you intend to start the application on the original secondary cluster. The application uses the current data on the original secondary cluster.

Caution – Because the `horctakeover` command did not perform a swap-takeover, the data volumes on `cluster-newyork` might not be synchronized with the data volumes on `cluster-paris`. If you intend to start the application with the same data that appears on the original primary cluster, you must not make the original secondary cluster `Primary`.

▼ How to Make the Original Primary Cluster Primary for a Hitachi TrueCopy Protection Group

1 Deactivate the protection group on the original primary cluster.

```
phys-paris-1# geopg stop -e Local tcpg
```

2 Resynchronize the configuration of the protection group.

This command updates the configuration of the protection group on `cluster-paris` with the configuration information of the protection group on `cluster-newyork`.

```
phys-paris-1# geopg update tcpg
```

After the `geopg update` command completes successfully, `tcpg` has the following role on each cluster:

```
cluster-paris:
        tcpg role: Primary
cluster-newyork:
        tcpg role: secondary
```

3 Activate the protection group on both clusters in the partnership.

```
phys-paris-1# geopg start -e Global tcpg
```

This command starts the application on `cluster-paris`. Data replication starts from `cluster-paris` to `cluster-newyork`.

▼ How to Make the Original Secondary Cluster Primary for a Hitachi TrueCopy Protection Group

1 Resynchronize the configuration of the protection group.

This command updates the configuration of the protection group on `cluster-newyork` with the configuration information of the protection group on `cluster-paris`.

```
phys-newyork-1# geopg update tcpg
```

After the `geopg update` command completes successfully, `tcpg` has the following role on each cluster:

```
cluster-paris:
        tcpg role: Secondary
cluster-newyork:
        tcpg role: Primary
```

2 Activate the protection group on both clusters in the partnership.

```
phys-newyork-1# geopg start -e Global tcpg
```

This command starts the application on `cluster-newyork`. Data replication starts from `cluster-newyork` to `cluster-paris`.

Caution – This command overwrites the data on `cluster-paris`.

Recovering From a Hitachi TrueCopy Data Replication Error

When an error occurs at the data replication level, the error is reflected in the status of the resource in the replication resource group of the relevant device group.

How to Detect Data Replication Errors

For information about how different `Resource status` values map to actual replication pair states, see Table 2–6.

You can check the status of the replication resources by using the `scstat -g` command as follows:

```
phys-paris-1# scstat -g
```

Running the `scstat -g` command might return the following:

```
...

--Resources --
            Resource Name       Node Name           State         Status Message
            -------------       ---------           -----         --------------
Resource: r-tc-tcpg1-devgroup1  phys-paris-2        Offline       Offline
Resource: r-tc-tcpg1-devgroup1  phys-paris-1        Online        Faulted - P-VOL:PSUE

Resource: hasp4nfs              phys-paris-1        Offline       Offline
Resource: hasp4nfs              phys-paris-2        Offline       Offline

...
```

The aggregate resource status for all device groups in the protection group is provided by using the `geoadm status` command. For example, the output of the `scstat -g` command in the preceding example indicates that the Hitachi TrueCopy device group, `devgroup1`, is in the `PSUE` state on `cluster-paris`. Table 2–6 indicates that the `PSUE` state corresponds to a resource status of `FAULTED`. So, the data replication state of the protection group is also `FAULTED`. This state is reflected in the output of the `geoadm status` command, which displays the state of the protection group as `Error`.

```
phys-paris-1# geoadm status
Cluster: cluster-paris

Partnership "paris-newyork-ps"  : OK
   Partner clusters             : cluster-newyork
   Synchronization              : OK
   ICRM Connection              : OK

   Heartbeat "paris-to-newyork" monitoring "cluster-newyork": OK
      Heartbeat plug-in "ping_plugin"             : Inactive
      Heartbeat plug-in "tcp_udp_plugin"          : OK

Protection group "tcpg"   : Error
      Partnership         : paris-newyork-ps
      Synchronization     : OK

      Cluster cluster-paris    : Error
         Role                  : Primary
         PG activation state   : Activated
         Configuration         : OK
         Data replication      : Error
         Resource groups       : OK

      Cluster cluster-newyork  : Error
         Role                  : Secondary
         PG activation state   : Activated
         Configuration         : OK
```

```
        Data replication      : Error
        Resource groups       : OK

Pending Operations
      Protection Group        : "tcpg"
      Operations              : start
```

How to Recover From a Hitachi TrueCopy Data Replication Error

To recover from an error state, you might perform some or all of the steps in the following procedure.

1. **Use the procedures in the Hitachi TrueCopy documentation to determine the causes of the** `FAULTED` **state. This state is indicated as** `PSUE`**.**

2. **Recover from the faulted state by using the Hitachi TrueCopy procedures.**

   If the recovery procedures change the state of the device group, this state is automatically detected by the resource and is reported as a new protection group state.

3. **Revalidate the protection group configuration.**

   ```
   phys-paris-1# geopg validate protectiongroupname
   ```

   protectiongroupname Specifies the name of the Hitachi TrueCopy protection group

4. **Review the status of the protection group configuration.**

   ```
   phys-paris-1# geopg list protectiongroupname
   ```

   protectiongroupname Specifies the name of the Hitachi TrueCopy protection group

5. **Review the runtime status of the protection group.**

   ```
   phys-paris-1# geoadm status
   ```

◆ ◆ ◆ APPENDIX A

Sun Cluster Geographic Edition Properties for Hitachi TrueCopy

This appendix provides the properties of Sun Cluster Geographic Edition data replication device groups.

This appendix contains the following sections:

Note – The property values, such as `True` and `False`, are *not* case sensitive.

Hitachi TrueCopy Properties

The following table describes the Hitachi TrueCopy properties that the Sun Cluster Geographic Edition software defines.

TABLE A–1 Hitachi TrueCopy Properties

Property	Description
Data Replication Property: `Cluster_dgs` (string array)	Lists the device groups where the data is written. The list is comma delimited. Only applications that belong to the protection group should write to these device groups. Tuning recommendations: This property can only be tuned when the protection group is offline. Category: Optional Default: Empty

TABLE A–1 Hitachi TrueCopy Properties *(Continued)*

Property	Description
Data Replication Property: `Nodelist` (string array)	Lists the host names of the machines that can be primary for the replication mechanism. This list is comma delimited. Tuning recommendations: This property can be tuned at any time. Category: Optional Default: All nodes in the cluster
Device Group Property: `Fence_level` (enum)	Defines the fence level that is used by the device group. The fence level determines the level of consistency among the primary and secondary volumes for that device group. Possible values are `Never` and `Async`. To use the `data` or `status` fence levels, contact your Sun representative. **Note –** If you specify a `Fence_level` of `never`, the data replication roles do not change after you perform a takeover. For more information about setting this property, see "How to Add a Data Replication Device Group to a Hitachi TrueCopy Protection Group" on page 40. Tuning recommendations: This property can only be tuned when the protection group is offline. Category: Required Default: None

Hitachi TrueCopy Properties That Must Not Be Changed

The Sun Cluster Geographic Edition software internally changes some properties for the `SUNWscgreptc` resource type. Therefore, you must not edit these properties manually.

For Hitachi TrueCopy, do not edit the following properties:

- `Dev_group` – Specifies the Hitachi TrueCopy device group that contains the volumes that are being replicated.
- `Replication_role` – Defines the local data replication role.

Index

R

S

T

V