Sun Cluster Geographic Edition System Administration Guide

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

Part No: 817–7501–10
August 2005

050909@13215

Contents

Tables

Figures

Examples

Preface

Sun Cluster Geographic Edition System Administration Guide provides procedures for administering 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 Documentation

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*
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

Documentation, Support, and Training

Sun Function	URL	Description
Documentation	`http://www.sun.com/documentation/`	Download PDF and HTML documents, and order printed documents
Support and Training	`http://www.sun.com/supportraining/`	Obtain technical support, download patches, and learn about Sun courses

Obtaining Help

If you have problems installing or using Sun Cluster Geographic Edition software, contact your service provider and provide the following information:

- Your name and email address (if available)
- Your company name, address, and phone number

- The model and serial numbers of your systems
- The release number of the operating system (for example, Solaris 9)
- The release number of the Sun Cluster Geographic Edition software (for example, 3.1 8/05)
- The contents of the `/var/opt/SUNWcacao/logs/cacao.0/1/2` file

Use the following commands to gather information about each node on your system for your service provider.

Command	Function
`prtconf -v`	Displays the size of the system memory and reports information about peripheral devices
`psrinfo -v`	Displays information about processors
`showrev –p`	Reports which patches are installed
`prtdiag -v`	Displays system diagnostic information
`geoadm -V`	Displays the Sun Cluster Geographic Edition software release and package version information
`scstat`	Provides a snapshot of the cluster status
`scconf –p`	Lists cluster configuration information
`geoadm status`	Prints the Sun Cluster Geographic Edition runtime status of the local cluster

Also have available the contents of the `/var/adm/messages` file.

Typographic Conventions

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

TABLE P–1 Typographic Conventions

Typeface or Symbol	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.`

TABLE P–1 Typographic Conventions *(Continued)*

Typeface or Symbol	Meaning	Example
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*. Perform a *patch analysis*. Do *not* save the file. [Note that some emphasized items appear bold online.]

Shell Prompts in Command Examples

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

TABLE P–2 Shell Prompts

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

CHAPTER 1

Introduction to Administering the Sun Cluster Geographic Edition Software

Sun Cluster Geographic Edition software protects applications from unexpected disruptions by using multiple clusters that are geographically separated. These clusters contain identical copies of the Sun Cluster Geographic Edition infrastructure, which manage replicated data between the clusters. Sun Cluster Geographic Edition software is a layered extension of the Sun Cluster software.

This chapter contains the following topics:

Sun Cluster Geographic Edition Administration Tasks

Familiarize yourself with the planning information in the *Sun Cluster Geographic Edition Installation Guide* and the *Sun Cluster Geographic Edition Overview* before beginning administration tasks. This guide contains the standard tasks that are used to administer and maintain the Sun Cluster Geographic Edition configurations.

For general Sun Cluster, data service, and hardware administration tasks, refer to the Sun Cluster documentation.

You can perform all administration tasks on a cluster that is running the Sun Cluster Geographic Edition software without causing any nodes or the cluster to fail. The Sun Cluster Geographic Edition software can be installed, configured, started, used, stopped, and uninstalled on an operational cluster.

Note – You might be required to take nodes or the cluster offline for preparatory actions, such as installing data replication software and performing Sun Cluster administrative tasks. Refer to the appropriate product documentation for administration restrictions.

Sun Cluster Geographic Edition Administration Tools

You can perform administrative tasks on a cluster that is running Sun Cluster Geographic Edition software by using a graphical user interface (GUI) or the command-line interface (CLI).

The procedures in this guide describe how to perform administrative tasks using the CLI.

Graphical User Interface

Sun Cluster software supports the SunPlex™ Manager, a GUI tool that you can use to perform various administrative tasks on your cluster. For specific information about how to use SunPlex Manager, see the Sun Cluster online help.

Note – To administer Sun Cluster Geographic Edition software using the GUI, the root passwords must be the same on all nodes of both clusters in the partnership.

You can only use the GUI to administer Sun Cluster Geographic Edition software after the software infrastructure has been enabled using the `geoadm start` command. Use the CLI to issue the `geoadm start` and `geoadm stop` commands. For information on enabling and disabling the Sun Cluster Geographic Edition infrastructure, see Chapter 3.

The GUI does not support creating custom heartbeats outside of a partnership. If you want to specify a custom heartbeat in a partnership join operation, use the CLI to execute the `geops join-partnership` command.

Command-Line Interface

Table 1–1 lists the commands that you can use to administer the Sun Cluster Geographic Edition software. For more information about each command, refer to the *Sun Cluster Geographic Edition Reference Manual*.

TABLE 1–1 Sun Cluster Geographic Edition CLI

Command	Description
`geoadm`	Enables or disables the Sun Cluster Geographic Edition software on the local cluster and prints the runtime status of the local cluster
`geohb`	Configures and manages the heartbeat mechanism that is provided with the Sun Cluster Geographic Edition software
`geops`	Creates and manages the partnerships between clusters
`geopg`	Configures and manages protection groups

Overview of Disaster Recovery Administration

This section provides an example of a disaster recovery scenario and actions an administrator might perform.

Company X has two geographically separated clusters, `cluster-paris` in Paris, and `cluster-newyork` in New York. These clusters are configured as partner clusters. The cluster in Paris is configured as the primary cluster and the cluster in New York is the secondary.

The `cluster-paris` cluster fails temporarily as a result of power outages during a windstorm. For an administrator, the following events occur:

1. The heartbeat communication is lost between `cluster-paris` and `cluster-newyork`. Because heartbeat notification was configured during the creation of the partnership, a heartbeat-loss notification email is sent to the administrator.

   For information about the configuring partnerships and heartbeat notification, see "Creating and Modifying a Partnership" on page 49.
2. The administrator receives the notification email and follows the company procedure to verify the disconnect occurred because of a situation that requires a takeover by the secondary cluster. Because a takeover is expensive, Company X does not allow takeovers unless the primary cluster cannot be repaired within two hours.

For information about verifying a disconnect on a system that uses Sun StorEdge Availability Suite 3.2.1, see "Detecting Cluster Failure on a System That Uses Sun StorEdge Availability Suite 3.2.1 Data Replication" on page 111.

For information about verifying a disconnect on a system that uses Hitachi TrueCopy, see "Detecting Cluster Failure on a System That Uses Hitachi TrueCopy Data Replication" on page 179.

3. Because the `cluster-paris` cluster cannot be brought online again for at least another day, the administrator executes a `geopg takeover` command on a New York node, which starts the protection group on the secondary cluster `cluster-newyork` in New York.

   For information about performing a takeover on a system that uses Sun StorEdge Availability Suite 3.2.1 data replication, see "Forcing a Takeover on Systems That Use Sun StorEdge Availability Suite 3.2.1" on page 115. For information about performing a takeover on a system that uses Hitachi TrueCopy data replication, see "Forcing a Takeover on a System That Uses Hitachi TrueCopy Data Replication" on page 183.

4. After the takeover, the secondary cluster `cluster-newyork` becomes the new primary cluster. The failed cluster in Paris is still configured to be primary, so when `cluster-paris` restarts, the cluster detects that it was down and lost contact with the partner cluster. Then, `cluster-paris` enters an error state that requires administrative action to repair. The cluster might also need to recover and resynchronize data.

   For information about recovering data after a takeover on a system that uses Sun StorEdge Availability Suite 3.2.1 data replication, see "Recovering Sun StorEdge Availability Suite 3.2.1 Data After a Takeover" on page 118. For information about performing a takeover on a system that uses Hitachi TrueCopy data replication, see "Failback of Services to the Original Primary Cluster on a System That Uses Hitachi TrueCopy Replication" on page 187.

CHAPTER 2

Before You Begin

This chapter describes what you need to know before you begin administering the Sun Cluster Geographic Edition software. Here you also learn about the Sun Cluster infrastructure that is required by the Sun Cluster Geographic Edition software. You also can find here common Sun Cluster concepts and tasks you need to understand before administering the Sun Cluster Geographic Edition software. This chapter also provides an example configuration that is used throughout this guide to illustrate the common Sun Cluster Geographic Edition administration tasks.

This chapter discusses the following topics:

Overview of Sun Cluster Administration Concepts

You must be an experienced Sun Cluster administrator to administer Sun Cluster Geographic Edition software.

This section describes some specific Sun Cluster administration topics that you need to understand before you administer Sun Cluster Geographic Edition software.

Configuring Resources and Resource Groups

You use either the `scrgadm` command or the SunPlex Manager to create failover and scalable resource groups.

For more information about administering resources and resource groups in Sun Cluster software, see the *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*.

Configuring Logical Hostnames

The logical hostname is a special high availability (HA) resource. The `geoadm start` command configures the logical hostname that corresponds to the cluster name. The IP address and host maps for the logical hostname must be set up before you run this command. Before assigning hostnames, familiarize yourself with the legal names and values described in Appendix B.

For more information about using the `geoadm start` command , see "Enabling the Sun Cluster Geographic Edition Software" on page 36.

Note – If you are using Sun StorEdge™ Availability Suite 3.2 for data replication, one logical hostname is created for each device group to be replicated. For more information, see Chapter 6.

The following table lists the Sun Cluster and Sun Cluster Geographic Edition components that need to be assigned IP addresses. Add these IP addresses to the following locations:

- All naming services that are being used
- The local `/etc/inet/hosts` file on each cluster node, after you install Solaris software

TABLE 2–1 IP Addresses Required by Sun Cluster Geographic Edition Software

Component	Number of IP Addresses Needed
Sun Cluster administrative console	1 per subnet
IP Network Multipathing groups	■ Single-adapter groups – 1 primary IP address. For the Solaris 8 release, also 1 test IP address for each adapter in the group. ■ Multiple-adapter groups – 1 primary IP address plus 1 test IP address for each adapter in the group.
Cluster nodes	1 per node, per subnet
Domain console network interface (Sun Fire™ 15000)	1 per domain

TABLE 2–1 IP Addresses Required by Sun Cluster Geographic Edition Software *(Continued)*

Console-access device	1
Logical addresses	1 per logical host resource, per subnet
Sun Cluster Geographic Edition infrastructure hostname	1 logical IP address per cluster infrastructure For example, if you have two clusters in your Sun Cluster Geographic Edition infrastructure, you need two IP addresses.
Replication with Sun StorEdge Availability Suite 3.2.1 software	1 dedicated logical IP address on the local cluster for each device group to be replicated For example, if you have two clusters in your Sun Cluster Geographic Edition infrastructure, you need two IP addresses.

For more information about configuring the IP address and host maps during the installation of Sun Cluster software, refer to Chapter 2, "Installing and Configuring Sun Cluster Software," in *Sun Cluster Software Installation Guide for Solaris OS*.

Managing Device Groups

A device group is a hardware resource that is managed by the Sun Cluster software. A device group is a type of global device that is used by the Sun Cluster software to register device resources, such as disks. A device group can include the device resources of disks, Solaris Volume Manager disksets, and VERITAS Volume Manager disk groups.

For information about configuring device groups in Sun Cluster software, refer to Chapter 4, "Administering Global Devices, Disk-Path Monitoring, and Cluster File Systems," in *Sun Cluster System Administration Guide for Solaris OS*.

The Sun Cluster Geographic Edition software configures Sun Cluster device groups to include replication.

For more information about configuring data replication in Sun Cluster Geographic Edition software, see Chapter 6 and Chapter 9.

Overview of Sun Cluster Geographic Edition Administration Tasks

This section provides a starting point for administering the Sun Cluster Geographic Edition software. It contains the following topics:

Prerequisite Administration Tasks

Before you begin administering the Sun Cluster Geographic Edition software, you must identify the Sun Cluster installations you need to host protection groups. Then, you need to adjust the Sun Cluster configuration and environment to support the formation of partnerships and protection groups with the Sun Cluster Geographic Edition software. The following table describes these prerequisite tasks.

TABLE 2–2 Sun Cluster Geographic Edition Prerequisite Tasks

Task	Description
Set the `SC-clustername` to the cluster name you want to use with the Sun Cluster Geographic Edition software.	Use the `scconf`(1M) command. For more information, see "How to Enable Sun Cluster Geographic Edition Software" on page 36.
Set up the IP address and host maps for the cluster that is enabled to run Sun Cluster Geographic Edition software.	See Chapter 2, "Installing and Configuring Sun Cluster Software," in *Sun Cluster Software Installation Guide for Solaris OS*.
Install and configure your data replication product.	See the Sun StorEdge Availability Suite 3.2.1 or Hitachi TrueCopy documentation. This step is required before you can create protection groups with the `geopg create` command.
Port and configure application configuration and corresponding resource groups on clusters that are candidates for partnership.	You can use the Sun Cluster `scsnapshot` tool to facilitate porting of application resource groups. See "Creating and Modifying a Partnership" on page 49 for more information.
Enable the common agent container on all the nodes of both clusters.	See "Enabling the Sun Cluster Geographic Edition Software" on page 36.

Sun Cluster Geographic Edition Administration Tasks

After you have completed the prerequisite administration tasks, you can install, configure, and administer the Sun Cluster Geographic Edition software as described in the following table.

TABLE 2–3 Sun Cluster Geographic Edition Administration Tasks

Task	Description and Documentation
Install Sun Cluster Geographic Edition software.	See the *Sun Cluster Geographic Edition Installation Guide*.
Set up security between the candidate partner clusters.	■ Exchange certificates, as described in "Configuring Secure Cluster Communication Using Security Certificates" on page 45. ■ (Optional) Configure a secure logical hostname that uses IP Security Architecture (IPsec), as described in "Configuring Secure Cluster Communication Using IPsec" on page 46.
Enable the Sun Cluster Geographic Edition software.	Use the `geoadm start` command. For more information, see "Enabling the Sun Cluster Geographic Edition Software" on page 36.
Create partnerships.	See "How to Create a Partnership" on page 50. This procedure includes the following: ■ Modifying the default heartbeat. For more information, see Chapter 12. ■ Configuring loss of heartbeat notification. For more information, see "Configuring Loss of Heartbeat Notification" on page 214.
Configure data replication.	For information about replicating data using Sun StorEdge Availability Suite 3.2.1, see Chapter 6 For information about replicating data using Hitachi TrueCopy, see Chapter 9.

TABLE 2–3 Sun Cluster Geographic Edition Administration Tasks *(Continued)*

Task	Description and Documentation
Create protection groups.	■ Create a protection group. See "How to Create and Configure a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 81 or "How to Create and Configure a Hitachi TrueCopy Protection Group" on page 144. ■ Add data replication device groups. See "How to Add a Data Replication Device Group to a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 94 or "How to Add a Data Replication Device Group to a Hitachi TrueCopy Protection Group" on page 155. ■ Add application resource groups to the protection group. See "How to Add an Application Resource Group to a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 90 or"How to Add an Application Resource Group to a Hitachi TrueCopy Protection Group" on page 152. ■ Create a protection group that does not require data replication. See "Creating a Protection Group That Does Not Require Data Replication" on page 88
Bring the protection groups online.	See "How to Activate a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 101 or "How to Activate a Hitachi TrueCopy Protection Group" on page 167.
Test the configured partnership and protection groups to validate the setup.	Perform a trial a switchover or takeover and test some simple failure scenarios. See Chapter 8 or Chapter 11.
Migrate services to the partner cluster.	See "How to Switch Over a Sun StorEdge Availability Suite 3.2.1 Protection Group From Primary to Secondary" on page 113 or "How to Switch Over a Hitachi TrueCopy Protection Group From Primary to Secondary" on page 182.
Take over services from primary to secondary during a disaster.	See "How to Force Immediate Takeover of Sun StorEdge Availability Suite 3.2.1 Services by a Secondary Cluster" on page 116 or"How to Force Immediate Takeover of Hitachi TrueCopy Services by a Secondary Cluster" on page 186.

TABLE 2–3 Sun Cluster Geographic Edition Administration Tasks *(Continued)*

Task	Description and Documentation
Recover from a takeover	■ Data recovery and error repair outside of the Sun Cluster Geographic Edition infrastructure. See the Sun StorEdge Availability Suite 3.2.1 or Hitachi TrueCopy documentation. ■ Resynchronize the partner clusters. See "Recovering Sun StorEdge Availability Suite 3.2.1 Data After a Takeover" on page 118 or "Failback of Services to the Original Primary Cluster on a System That Uses Hitachi TrueCopy Replication" on page 187.
Take a protection group offline.	See "How to Deactivate a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 104 or "How to Deactivate a Hitachi TrueCopy Protection Group" on page 171.
Delete a protection group.	See "How to Delete a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 86 or "How to Delete a Hitachi TrueCopy Protection Group" on page 148.
Delete a partnership.	See "Leaving or Deleting a Partnership" on page 56.
Disable the Sun Cluster Geographic Edition software.	See "How to Disable the Sun Cluster Geographic Edition Software" on page 39.
Uninstall the Sun Cluster Geographic Edition software.	See the *Sun Cluster Geographic Edition Installation Guide*.

Example Sun Cluster Geographic Edition Cluster Configuration

The following figure describes an example Sun Cluster Geographic Edition cluster configuration that is used throughout this guide to illustrate the Sun Cluster Geographic Edition administration tasks. The primary cluster, `cluster-paris` contains two nodes, `phys-paris-1` and `phys-paris-2`. The secondary cluster, `cluster-newyork`, also contains two nodes, `phys-newyork-1` and `phys-newyork-2`.

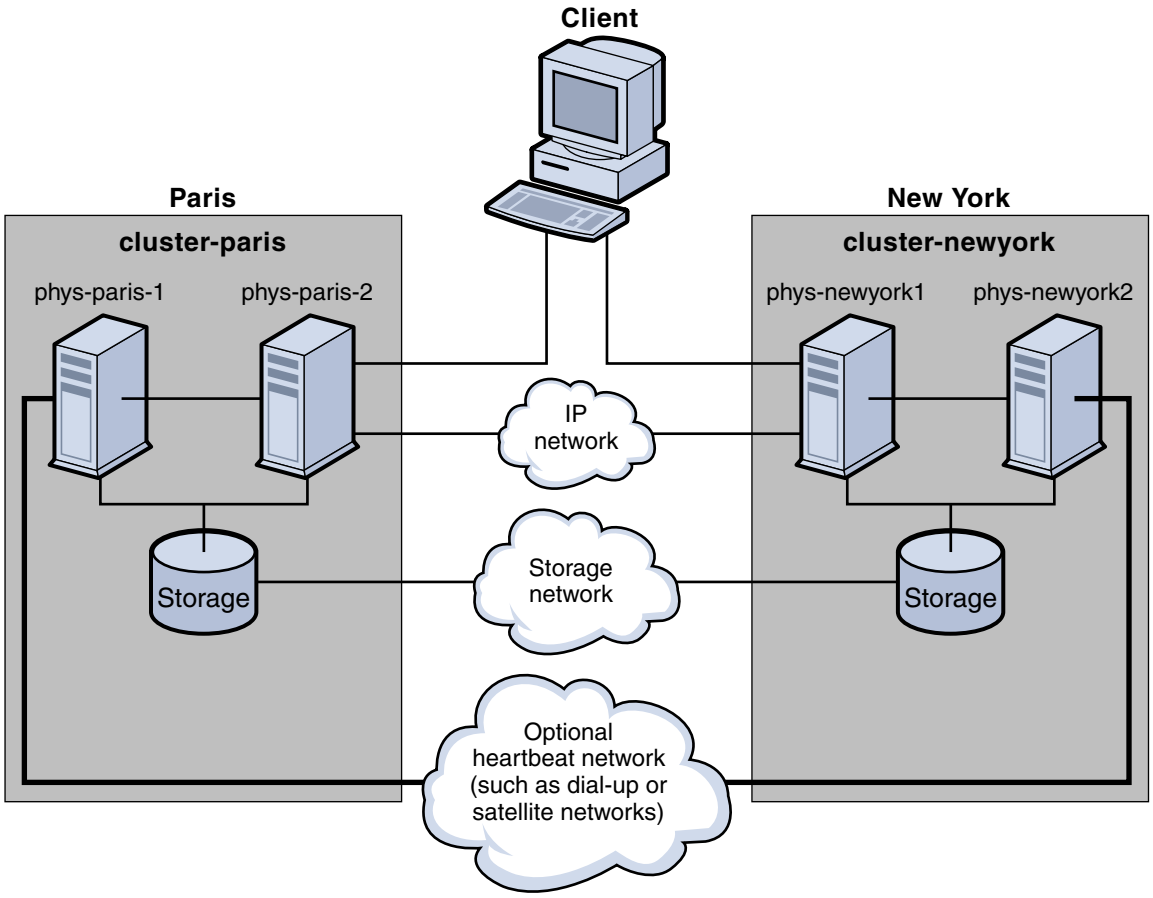

FIGURE 2–1 Example Cluster Configuration

CHAPTER 3

Administering the Sun Cluster Geographic Edition Infrastructure

This chapter contains information about enabling your cluster for participation in one or more partnerships. It also contains information for disabling the Sun Cluster Geographic Edition software so that your cluster no longer can participate in partnerships.

This chapter contains the following topics:

About the Sun Cluster Geographic Edition Infrastructure Resource Groups

When you enable the Sun Cluster Geographic Edition infrastructure, the following Sun Cluster resource groups are created:

- `geo-clusterstate` – A scalable resource group that the Sun Cluster Geographic Edition software uses to distinguish between node failover and cluster reboot scenarios. This resource group does not contain any resources.
- `geo-infrastructure` – A failover resource group that encapsulates the Sun Cluster Geographic Edition infrastructure. The resource group contains the following resources:

- `geo-clustername` – Encapsulates the logical hostname for the Sun Cluster Geographic Edition software. The Sun Cluster Geographic Edition software uses the logical hostname of a cluster for inter-cluster management communication and heartbeat communication. A logical hostname must be the same as the name of the cluster and be available on the namespace of each cluster.
- `geo-hbmonitor` – Encapsulates the heartbeat processes for the Sun Cluster Geographic Edition software.
- `geo-failovercontrol` – Encapsulates the Sun Cluster Geographic Edition software itself. The Sun Cluster Geographic Edition module uses this resource to load into the common agent container.

These resources are for internal purposes only, so you must not change them.

These internal resources are removed when you disable the Sun Cluster Geographic Edition infrastructure.

You can monitor the status of these resources by using the `scstat -g` command. For more information about this command, see the `scstat`(1M) man page.

Enabling the Sun Cluster Geographic Edition Software

When you enable the Sun Cluster Geographic Edition software, the cluster is ready to enter a partnership with another enabled cluster. You can use the CLI commands or the GUI to create a cluster partnership.

For more information about setting up and installing the Sun Cluster Geographic Edition software, see the *Sun Cluster Geographic Edition Installation Guide*.

▼ How to Enable Sun Cluster Geographic Edition Software

Before You Begin

Before you enable Sun Cluster Geographic Edition software on a cluster, ensure that the following conditions are met:

- The cluster is running the Solaris Operating System and the Sun Cluster software.
- The Sun Cluster management-agent container for SunPlex Manager is running.
- The Sun Cluster Geographic Edition software is installed.

Steps

1. **Log in to a cluster node.**

   You must be assigned the Geo Operation RBAC rights profile to complete this procedure. For more information about RBAC, see "Sun Cluster Geographic Edition Software and RBAC" on page 43.

2. **Ensure that the logical hostname, which is the same as the cluster name, is available and defined.**

   ```
   # scconf -p | grep -i "cluster name"
   ```

   If the cluster name is not the name you want to use, you can change the cluster name with the following command:

   ```
   # scconf -c -C cluster=cluster-name
   ```

   For more information, see the `scconf`(1M) man page.

   Note – After you have enabled the Sun Cluster Geographic Edition infrastructure, you must not change the cluster name while the infrastructure is enabled.

   ---

3. **Confirm that a logical hostname that matches the cluster name is available and defined in the local host files.**

   The local host file, `hosts`, is located in the `/etc/inet` directory.

   Confirm that the logical hostname is also defined in the network namespace database, for example, NIS.

4. **On one node of the cluster, create the Sun Cluster Geographic Edition infrastructure resource groups and enable the Sun Cluster Geographic Edition control module.**

   ```
   # geoadm start
   ```

   The `geoadm start` command enables the Sun Cluster Geographic Edition control module on the local cluster only. For more information, see the `geoadm`(1M) man page.

5. **Verify that you have enabled the infrastructure and that the Sun Cluster Geographic Edition resource groups are online.**

   For a list of the Sun Cluster Geographic Edition resource groups, see "About the Sun Cluster Geographic Edition Infrastructure Resource Groups" on page 35.

   ```
   # geoadm show
   # scstat -g
   ```

   The output for the `geoadm show` command should state that the Sun Cluster Geographic Edition infrastructure is active from a particular node in the cluster.

The output for the `scstat -g` command should state that the `geo-failovercontrol`, `geo-hbmonitor`, and `geo-clustername` resources and the `geo-infrastructure` resource groups are online on one node of the cluster.

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

Example 3–1 Enabling a Cluster

The following example illustrates how to enable Sun Cluster Geographic Edition software on `cluster-paris`.

1. Start the Sun Cluster Geographic Edition software on `cluster-paris`.

```
phys-paris-1# geoadm start
```

2. Determine whether the Sun Cluster Geographic Edition infrastructure was successfully enabled.

```
phys-paris-1# geoadm show

--- CLUSTER LEVEL INFORMATION ---
Sun Cluster Geographic Edition is active on cluster-paris from node phys-paris-1
Command execution successful
phys-paris-1#
```

3. Verify the status of the Sun Cluster Geographic Edition resource groups and resources.

```
phys-paris-1# scstat -g
-- Resource Groups and Resources --
      Group Name              Resources
      ----------           ---------
Resources: geo-clusterstate    -
Resources: geo-infrastructure geo-clustername geo-hbmonitor geo-failovercontrol

-- Resource Groups --
     Group Name            Node Name       State
     ----------            ---------       -----
Group: geo-clusterstate   phys-paris-1    Online
Group: geo-clusterstate   phys-paris-2    Online
Group: geo-infrastructure phys-paris-1    Online
Group:geo-infrastructure  phys-paris-2    Offline

-- Resources --
Resource Name                 Resources    State   Status Message
-------------                 ---------    -----   --------------
Resource: geo-clustername     phys-paris-1 Online  Online - LogicalHostname online
Resource: geo-clustername     phys-paris-2 Offline Offline
Resource: geo-hbmonitor       phys-paris-1 Online  Online- Daemon OK
Resource: geo-hbmonitor       phys-paris-2 Offline Offline
Resource: geo-failovercontrol phys-paris-1 Online  Online
Resource: geo-failovercontrol phys-paris-2 Offline Offline
```

Next Steps See Chapter 6 or Chapter 10 for information on creating protection groups.

Disabling the Sun Cluster Geographic Edition Software

You can disable the Sun Cluster Geographic Edition infrastructure by using the following procedure.

▼ How to Disable the Sun Cluster Geographic Edition Software

Before You Begin

Before you can disable the Sun Cluster Geographic Edition infrastructure, all protection groups on the local cluster must be offline.

Steps

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" on page 43.

2. **Confirm that all of the protection groups are offline on the local cluster.**

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

   For more information about the `geoadm status` command and its output, see "Monitoring the Runtime Status of the Sun Cluster Geographic Edition Software" on page 217.

   If you want to keep the application resource groups online while deactivating a Sun StorEdge Availability Suite 3.2.1 protection group, follow the procedure described in Example 7–18.

   If you want to keep the application resource groups online while deactivating a Hitachi TrueCopy protection group, follow the procedure described in Example 10–23.

3. **Disable the Sun Cluster Geographic Edition software.**

   ```
   phys-paris-1# geoadm stop
   ```

   This command removes the infrastructure resource groups that were created when you enabled the Sun Cluster Geographic Edition infrastructure.

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

Note – If you did not deleted the protection group entities by using the Sun Cluster Geographic Edition software, then the data replication resource groups are not removed when you disable the Sun Cluster Geographic Edition infrastructure. The data replication resource groups are available again when you enable the Sun Cluster Geographic Edition infrastructure.

4. **Verify that the software was disabled and that the Sun Cluster Geographic Edition resource groups are no longer displayed.**

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

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

Example 3–2 Disabling a Cluster

The following example illustrates how to disable `cluster-paris`.

1. Confirm that all protection groups are offline as follows.

```
phys-paris-1# geoadm status

Cluster: cluster-paris

Partnership "paris-newyork-ps" : OK

   Partner clusters: cluster-newyork
   Synchronization:  OK

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

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

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


    Cluster cluster-newyork  : OK
        Role                 : Secondary
        PG activation state  : Deactivated
        Configuration        : OK
        Data replication     : OK
        Resource groups      : OK
```

2. Disable the Sun Cluster Geographic Edition infrastructure.

```
phys-paris-1# geoadm stop
... verifying pre conditions and performing pre remove operations ... done
...removing product infrastructure ... please wait ...
```

3. Use the `geoadm show` command to determine if the Sun Cluster Geographic Edition infrastructure was successfully disabled.

```
phys-paris-1# geoadm show

--- CLUSTER LEVEL INFORMATION ---
Sun Cluster Geographic Edition is not active on cluster-paris

--- LOCAL NODE INFORMATION ---
Node phys-paris-1 does not host active product module.

Command execution successful
phys-paris-1#
```

4. Verify that Sun Cluster Geographic Edition resource groups and resources have been removed by using the `scstat -g` command.

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

Checking the Status of the Sun Cluster Geographic Edition Infrastructure

You can use the `geoadm show` command to determine whether the Sun Cluster Geographic Edition infrastructure is enabled on the local cluster and, if so, on which node the infrastructure is active. The Sun Cluster Geographic Edition infrastructure is considered active on the node on which the `geo-infrastructure` resource group has a state of `Online`.

For example, the `geoadm show` command is run on the `phys-paris-1` node of `cluster-paris` as follows:

```
phys-paris-1# geoadm show

--- CLUSTER LEVEL INFORMATION ---
Sun Cluster Geographic Edition is active on:
node phys-paris-2, cluster cluster-paris

Command execution successful
phys-paris-1#
```

Booting a Cluster

During a boot, the following steps occur:

1. After the Sun Cluster infrastructure is enabled, the Sun Cluster Geographic Edition software starts automatically. Verify that the software started by using the `geoadm show` command.
2. The heartbeat framework checks which partners it can reach.
3. Check the current status of the cluster by using the `geoadm status` command. For more information about this command and its output, see "Monitoring the Runtime Status of the Sun Cluster Geographic Edition Software" on page 217.

Applying Patches to a Sun Cluster Geographic Edition System

Complete the following steps when you apply patches to your Sun Cluster Geographic Edition system.

1. Shut down the Sun Cluster Geographic Edition infrastructure by using the `geoadm stop` command.

   Shutting down the infrastructure ensures that a patch installation on one cluster does not affect the other cluster in the partnership.
2. Apply the patches using the `patchadd` command.

   If you are applying Sun Cluster patches, use the Sun Cluster methods on both clusters.

   Note – Apply patches to the secondary cluster first to ensure that errors do not affect the services on the primary cluster. After you have verified the installation, apply the patches to the primary cluster.

   ---

3. Restart the Sun Cluster Geographic Edition infrastructure by using the `geoadm start` command.

CHAPTER 4

Administering Access and Security

This section describes how to administer access and security. It contains the following topics:

Sun Cluster Geographic Edition Software and RBAC

This section describes role-based access control (RBAC) in Sun Cluster Geographic Edition software. Topics that are covered include the following:

- Setting Up and Using RBAC
- RBAC Rights Profiles
- Modifying a User's RBAC Properties

Setting Up and Using RBAC

Sun Cluster Geographic Edition software bases its RBAC profiles on the RBAC rights profiles that are used in the Sun Cluster software. For general information about setting up and using RBAC with Sun Cluster, refer to Chapter 2, "Sun Cluster and RBAC," in *Sun Cluster System Administration Guide for Solaris OS*.

Sun Cluster Geographic Edition software adds the following new RBAC entities to the appropriate file in the `/etc/security` directory:

- RBAC authentication names to `auth_attr`
- RBAC execution profiles to `prof_attr`
- RBAC execution attributes to `exec_attr`

Note – The default search order for the `auth_attr` and `prof_attr` databases is `files nis`, which is defined in the `/etc/nsswitch.conf` file. If you have customized the search order in your environment, confirm that `files` is present in the search list. `files` enables your system to find the RBAC entries that Sun Cluster Geographic Edition defined.

RBAC Rights Profiles

The Sun Cluster Geographic Edition CLI and GUI use RBAC rights to control end-user access to operations. The general conventions for these rights are described in the following table.

TABLE 4–1 Sun Cluster Geographic Edition RBAC Rights Profiles

Rights Profile	Included Authorizations	Role Identity Permission
Geo Management	`solaris.cluster.geo.read`	Read information about the Sun Cluster Geographic Edition entities
	`solaris.cluster.geo.admin`	Perform administrative tasks with the Sun Cluster Geographic Edition software
	`solaris.cluster.geo.modify`	Modify the configuration of the Sun Cluster Geographic Edition software
Basic Solaris User	Solaris authorizations	Perform the same operations that the Basic Solaris User role identity can perform
	`solaris.cluster.geo.read`	Read information about the Sun Cluster Geographic Edition entities

Modifying a User's RBAC Properties

To modify the RBAC rights for a user, you must be logged in as the `root` user or assume a role that is assigned the Primary Administrator rights profile.

For example, you can assign the Geo Management RBAC profile to the user `admin` as follows:

```
# usermod -P "Geo Management" admin
# profiles admin
Geo Management
Basic Solaris User
#
```

For more information about how to modify the RBAC properties for a user, refer to Chapter 2, "Sun Cluster and RBAC," in *Sun Cluster System Administration Guide for Solaris OS*.

Configuring Secure Cluster Communication Using Security Certificates

You must configure the Sun Cluster Geographic Edition software for secure communication between partner clusters. The configuration must be reciprocal, which means that partner cluster `cluster-paris` must be configured to trust `cluster-newyork`, and that partner cluster `cluster-newyork` must be configured to trust `cluster-paris`.

If you are using the GUI to administer the Sun Cluster Geographic Edition software, the root password must be the same on all nodes of both partner clusters.

For information about setting up security certificates for partner clusters, see "Configuring Security" in *Sun Cluster Geographic Edition Installation Guide*.

For information about the example cluster configuration, see Figure 2–1.

Configuring Secure Cluster Communication Using IPsec

You can use IP Security Architecture (IPsec) to configure secure communication between partner clusters. IPsec enables you to set policies that permit or require either secure datagram authentication, or actual data encryption, or both, between machines communicating by using IP. Consider using IPsec for the following cluster communications:

- Secure Sun StorEdge Availability Suite 3.2.1 communications, if you use Sun StorEdge Availability Suite 3.2.1 for data replication
- Secure TCP/UDP heartbeat communications

Sun Cluster software and Sun Cluster Geographic Edition software support IPsec by using only manual keys. Keys must be stored manually on the cluster nodes for each combination of server and client IP address. The keys must also be stored manually on each client.

Refer to the *System Administration Guide: IP Services* for a full description of IPsec configuration parameters.

▼ How to Configure IPsec for Secure Cluster Communication

In the Sun Cluster Geographic Edition infrastructure, the hostname of a logical host is identical to the cluster name. The logical hostname is a special high availability (HA) resource. You must set up a number of IP addresses for various Sun Cluster Geographic Edition components, depending on your cluster configuration.

On each partner cluster, you must configure encryption and authorization for exchanging inbound and outbound packets from a physical node to the logical-hostname addresses. The values for the IPsec configuration parameters on these addresses must be consistent between partner clusters.

IPsec uses two configuration files:

- **IPsec policy file,** `/etc/inet/ipsecinit.conf`. Contains directional rules to support an authenticated, encrypted heartbeat. The contents of this file will be different on the two clusters in your partnership.
- **IPsec keys file,** `/etc/init/secret/ipseckeys`. Contains keys files for specific authentication and encryption algorithms. The contents of this file will be identical on both clusters in your partnership.

The following procedure configures an example cluster, `cluster-paris`, for IPsec secure communication with another example cluster, `cluster-newyork`. Both clusters are running the Solaris OS 9 release. The procedure assumes that the local logical hostname on `cluster-paris` is `lh-paris-1` and that the remote logical hostname is `lh-newyork-1`. Inbound messages are sent to `lh-paris-1` and outbound messages are sent to `lh-newyork-1`.

Use the following procedure on each node of `cluster-paris`.

Steps

1. **Log in to the first node of the primary cluster, `phys-paris-1`, as superuser.**

   For a reminder of which node is `phys-paris-1`, see "Example Sun Cluster Geographic Edition Cluster Configuration" on page 33.

2. **Set up an entry for the local address and remote address in the IPsec policy file.**

   The policy file is located at `/etc/inet/ipsecinit.conf`. Permissions on this file should be `644`. For more information about this file, see the `ipsecconf`(1M) man page.

   For information about the names and values that are supported by Sun Cluster Geographic Edition software, see Appendix B.

   a. **Configure the communication policy.**

      The default port for the or `tcp_udp` plug-in is 2084. This is specified in the `/etc/opt/SUNWcacao/modules/com.sun.cluster.agent.geocontrol.xml` file.

      The following command configures a policy with no preference for authorization or encryption algorithms:

      ```
      # {raddr lh-newyork-1 rport 2084} ipsec {auth_algs any encr_algs any \
      sa shared} {laddr lh-paris-1 lport 2084} ipsec {auth_algs any encr_algs \
      any sa shared}
      ```

      When you configure the communication policy on the secondary cluster, `cluster-newyork`, the policies need to be reversed:

      ```
      # {laddr lh-newyork-1 lport 2084} ipsec {auth_algs any encr_algs \
      any sa shared} {raddr lh-paris-1 rport 2084} ipsec {auth_algs any encr_algs \
      any sa shared}
      ```

   b. **Add the policy by rebooting the node or by running the following command.**

      ```
      # ipsecconf -a /etc/inet/ipsecinit.conf
      ```

3. **Set up encryption and authentication keys for inbound and outbound communication.**

   The communication file is located at `/etc/init/secret/ipseckeys`. Permissions on the file should be `600`.

   Add keys by running the following command:

   ```
   # ipseckey -f /etc/init/secret/ipseckeys
   ```

Key entries have the following general format:

```
# inbound to cluster-paris
add esp spi <paris-encr-spi> dst lh-paris-1 encr_alg <paris-encr-algorithm> \
encrkey <paris-encrkey-value>
add ah spi <newyork-auth-spi> dst lh-paris-1 auth_alg <paris-auth-algorith> \
authkey <paris-authkey-value>

# outbound to cluster-newyork
add esp spi <newyork-encr-spi> dst lh-newyork-1 encr_alg \
<newyork-encr-algorithm> encrkey <newyork-encrkey-value>
add ah spi <newyork-auth-spi> dst lh-newyork-1 auth_alg \
<newyork-auth-algorithm> authkey <newyork-authkey-value>
```

For more information about the communication files, see the `ipsecconf`(1M) man page.

CHAPTER 5

Administering Cluster Partnerships

This chapter provides the procedures for administering partnerships between two Sun Cluster Geographic Edition software-enabled clusters.

This chapter discusses the following topics:

Creating and Modifying a Partnership

A partnership establishes heartbeat monitoring between two clusters that are running Sun Cluster Geographic Edition software. Clusters in a partnership exchange heartbeats to monitor each other's presence.

You create a partnership with the `geops`(1m) command. After you have created a partnership, you can use this command to modify the properties of this partnership.

The names of the application resource groups that are managed by the Sun Cluster Geographic Edition software must be the same on both partner clusters. You can configure the names of these resource groups manually or by using the `scsnapshot` command.

The `scsnapshot` command replicates configuration data on a cluster that does not have configured resource groups, resource types, and resources. The `scsnapshot` command retrieves the configuration data from the cluster on which it is launched and generates a script called `scriptfile`. Edit the script to adapt it to the specific

features of the cluster where you want to replicate the configuration data. For example, if you change the IP address and host names in the script, you can launch the script from any node in the cluster where you want to replicate the configuration data. For more information about using this command, see the `scsnapshot`(1M) man page.

You can configure a partnership between only two clusters, and only one partnership can be defined between a given pair of clusters. A single cluster can participate in multiple partnerships.

▼ How to Create a Partnership

Before You Begin

Before you can create a partnership between two clusters, ensure that the following conditions are met:

- The cluster on which you want to create the partnership is up and running.
- The `geoadm start` command must have already been run on the this cluster and the partner cluster. For more information about using the `geoadm start` command, see "Enabling the Sun Cluster Geographic Edition Software" on page 36.
- The cluster name of the partner cluster is known.
- The host information of the partner cluster must defined in the local host file. The local cluster needs to know how to reach the partner cluster.
- Security has been configured on the two clusters by installing the appropriate certificates.

  See "Configuring Secure Cluster Communication Using Security Certificates" on page 45 for more information.

Note – When you create or join multiple partnerships, do not use port numbers that are being used by other partnerships. All `tcp_udp` requests go through the RPC server and the `tcp_udp_resp` is created on the remote cluster with the port number defined for the local cluster. Use an unused port number and ensure that the port number is not registered with the Internet Assigned Numbers Authority (IANA).

The Sun Cluster Geographic Edition software uses the port in the CCR table rather than reading from the XML file, so when you change the port number, you must reset the configuration by deleting the partnership and manually update the XML information.

Steps

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" on page 43.

2. **If the default port is being used by another application or if you are creating multiple partnerships on a cluster, change the port number as follows:**

a. **If a partnership exists, leave the partnership.**

```
# geops leave-partnership paris-newyork-ps
```

b. **On all the nodes of the cluster, change the value of the `tcp_udp.port` in the `/etc/opt/SUNWcacao/modules/com.sun.cluster.agent.geocontrol.xml` file to the port you want the partnership to use.**

Only one heartbeat can use the default port. If a partnership's heartbeat is using the default port, the other partnership's heartbeat must be created with a specific port. The port you specify in the file must be a port that is not used by any application other than the Sun Cluster Geographic Edition software.

If the cluster participates in more than one partnership, ensure that the port number is different from the other partners.

c. **Restart the common agent container on all the nodes of the cluster.**

```
# /opt/SUNWcacao/bin/cacaoadm restart
```

3. **Create the partnership.**

```
# geops create -c remote-partner-cluster-name [-h heartbeat-name] \
[-p property-setting [-p...]] partnership-name
```

`-c` *remote-cluster-name*	Specifies the name of the remote cluster that will participate in the partnership This name matches the logical hostname used by the Sun Cluster Geographic Edition infrastructure on the remote cluster.
`-h` *heartbeat-name*	Specifies a custom heartbeat to be used in the partnership to monitor the partner cluster's availability If you omit this option, the default Sun Cluster Geographic Edition heartbeat is used. Custom heartbeats are provided for special circumstances and require careful configuration. Consult your Sun specialist for assistance if your system requires the use of custom heartbeats. For more information about configuring custom heartbeats, see Chapter 12. The custom heartbeat that is specified by this option must already be configured before running the `geops` command.

Note – The presence of a custom heartbeat prevents the default heartbeat from being used during partnership creation. If you want to use the default heartbeat for your partnership, you must delete the custom heartbeat before running the `geops create` command.

`-p` *property-setting*	Sets the value of partnership properties with a string of *property=value* pair statements Specify a description of the partnership with the `Description` property. You can configure heartbeat-loss notification with the `Notification_emailaddrs` and `Notification_actioncmd` properties. For more information about configuring heartbeat-loss notification, see "Configuring Loss of Heartbeat Notification" on page 214. For more information about the properties you can set, see Appendix A.
partnership-name	Specifies the name of the partnership

For information about the names and values that are supported by Sun Cluster Geographic Edition software, see Appendix B.

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

4. **Verify that the partnership was created and the status of the partnership.**

```
# geoadm status
```

Example 5–1 Creating a Partnership

The following example illustrates the creation of the partnership, `paris-newyork-ps`, on `cluster-paris`:

```
# geops create -c cluster-newyork -p Description=Transatlantic \
-p Notification_emailaddrs=sysadmin@companyX.com paris-newyork-ps
# geoadm status
```

▼ How to Modify Partnership Properties

Steps

1. **Log in to one of the cluster nodes.**

   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" on page 43.

2. **Modify partnership properties.**

   ```
   # geops set-prop -p property-setting [-p...] partnership-name
   ```

   `-p` *property-setting* — Sets the value of partnership properties with a string of *property=value* pair statements

   Specify a description of the partnership with the `Description` property.

   You can configure heartbeat-loss notification with the `Notification_emailaddrs` and `Notification_actioncmd` properties. For more information about configuring heartbeat-loss notification, see "Configuring Loss of Heartbeat Notification" on page 214.

   For more information about the properties you can set, see Appendix A.

   partnership-name — Specifies the name of the partnership

   For information about the names and values that are supported by Sun Cluster Geographic Edition software, see Appendix B.

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

3. **Verify that your modification was made correctly.**

   ```
   # geops list
   ```

Example 5–2 Modifying the Properties of a Partnership

The following example illustrates how to modify the notification email address for `cluster-paris`:

```
# geops set-prop -p Notification_emailaddrs=operations@companyX.com \
paris-newyork-ps
# geops list
```

Joining an Existing Partnership

When you define and configure a partnership, the partnership specifies a second cluster to be a member of that partnership. Then, you must configure this second cluster to join the partnership.

▼ How to Join a Partnership

Before You Begin

Before you can configure a cluster to join a partnership, ensure that the following conditions are met:

- The local cluster is enabled to run the Sun Cluster Geographic Edition software.
- The partnership you want the cluster to join is defined and configured on another cluster (`cluster-paris`) and the local cluster (`cluster-newyork`) is specified as a member of this partnership.
- Security has been configured on the clusters by installing the appropriate certificates.

  See "Configuring Secure Cluster Communication Using Security Certificates" on page 45 for more information.

Note – When you create or join multiple partnerships, do not use port numbers that are being used by other partnerships. All `tcp_udp` requests go through the RPC server and the `tcp_udp_resp` is created on the remote cluster with the port number defined for the local cluster. Use an unused port number and ensure that the port number is not registered with the Internet Assigned Numbers Authority (IANA).

The Sun Cluster Geographic Edition software uses the port in the CCR table rather than reading from the XML file, so when you change the port number, you must reset the configuration by deleting the partnership and manually update the XML information.

Steps

1. **Log in to one of the nodes of the cluster that is to join the partnership.**

   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" on page 43.

2. **Confirm that the cluster to be joined, `cluster-paris`, can be reached at its logical hostname.**

   ```
   # ping lh-paris-1
   ```

For information about the logical hostname of the cluster, see "How to Enable Sun Cluster Geographic Edition Software" on page 36.

3. **Join an existing partnership.**

```
# geops join-partnership [-h heartbeat-name] remote-cluster-name partnership-name
```

`-h` *heartbeat-name*	Specifies an identifier to a custom heartbeat plug-in on a partner cluster that the local cluster can use to monitor the partner cluster's availability Custom heartbeats are provided for special circumstances and require careful configuration. Consult your Sun specialist for assistance if your system requires the use of custom heartbeats. For more information about configuring custom heartbeats, see Chapter 12. If you omit this option, the default Sun Cluster Geographic Edition heartbeat is used.
remote-cluster-name	Specifies the name of a cluster that is currently a member of the partnership being joined, from which the partnership configuration information will be retrieved
partnership-name	Specifies the name of the partnership

For information about the names and values that are supported by Sun Cluster Geographic Edition software, see Appendix B.

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

4. **Verify that the cluster was added to the partnership and that the partnership properties were defined correctly.**

```
# geops list
# geoadm status
```

Example 5–3 Joining a Partnership

The following example illustrates how `cluster-newyork` joins the partnership that was created on `cluster-paris` in Example 5–1.

```
# geops join-partnership cluster-paris paris-newyork-ps
# geops list
# geoadm status
```

Leaving or Deleting a Partnership

You can also use the `geops` command to remove a cluster from a partnership and release all of the partnership's associated resources.

Because this command destroys the local partnership configuration information, when the last member leaves a partnership, the partnership no longer exists.

▼ How to Leave a Partnership

Before You Begin

Before you leave a partnership, ensure that the following conditions are met:

- The local cluster is a member of the partnership you want to leave.
- This partnership does not contain any protection groups.

Steps

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" on page 43.

2. **Verify that the partnership does not have any protection groups.**

   ```
   # geopg list
   ```

   If you find that the partnership contains protection groups, you can delete them with the `geopg delete` command. For information about deleting Sun StorEdge Availability Suite 3.2.1 protection groups, see "How to Delete a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 86. For information about deleting a Hitachi TrueCopy protection group, see "How to Delete a Hitachi TrueCopy Protection Group" on page 148.

3. **Remove the partnership on one node of the cluster that is a member of the partnership.**

   ```
   # geops leave-partnership partnership-name
   ```

   partnership-name Specifies the name of the partnership

   Note – The `geops leave-partnership` command deletes the heartbeats configured for the partnership, including custom heartbeats.

   For more information, refer to the `geops`(1M) man page.

Example 5–4 Leaving a Partnership

The following example illustrates how `cluster-paris` leaves the `paris-newyork-ps` partnership.

```
phys-paris-1# geops leave-partnership paris-newyork-ps
```

Example 5–5 Deleting a Partnership

After `cluster-paris` leaves the `paris-newyork-ps` partnership, as described in the previous example, the only remaining member of the partnership is `cluster-newyork`. The `paris-newyork-ps` partnership can now be deleted by making its last remaining member, `cluster-newyork`, leave the partnership as follows:

```
phys-newyork-1# geops leave-partnership paris-newyork-ps
```

Next Steps Repeat this procedure on the other cluster in the partnership.

Resynchronizing a Partnership

Partners become disconnected during a disaster situation, forcing the administrator to perform a takeover for a protection group that the partners share. When both clusters are brought online again, both partner clusters report as the primary of the protection group. You must resynchronize the configuration information of the local protection group with the configuration information that is retrieved from the partner.

If a cluster that is a member of a partnership fails, when the cluster restarts, it detects whether the partnership parameters have been modified while it was down. You decide which partnership configuration information you want to keep: the information on the cluster that failed or the information on the failover cluster. Then, resynchronize the configuration of the partnership accordingly.

You can detect that a partnership needs to be resynchronized by looking at the output of the `geoadm status` command. If the `Configuration` status is `Synchronization Status Error`, the partnership needs to be synchronized. If the `Local status` is `Partnership Error`, you should not resynchronize the partnership. Instead, wait for the heartbeat exchange to occur.

▼ How to Resynchronize a Partnership

Before You Begin

Before you resynchronize a partnership, ensure that the following conditions are met:

- The local cluster is Sun Cluster Geographic Edition enabled.
- The local cluster was an active member of the partnership before failing.

Caution – Resynchronizing a partnership overwrites the partnership configuration on the cluster where the command is run with the information from the partner cluster.

Steps

1. **Log in to a node on the cluster that needs to be synchronized with the information retrieved from the partner 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" on page 43.

2. **Resynchronize the partnership.**

   ```
   # geops update partnership-name
   ```

   partnership-name Specifies the name of the partnership

Example 5–6 Resynchronizing a Partnership

The following example illustrates how to resynchronize a partnership:

```
# geops update paris-newyork-ps
```

CHAPTER 6

Replicating Data With Sun StorEdge Availability Suite 3.2.1 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.

Sun Cluster Geographic Edition software supports the use of Sun StorEdge Availability Suite 3.2.1 remote mirror software for data replication. Before you can replicate data with Sun StorEdge Availability Suite 3.2.1 software, you must be familiar with the Sun StorEdge Availability Suite 3.2.1 documentation, have the Sun StorEdge Availability Suite 3.2.1 product, and the latest Sun StorEdge Availability Suite 3.2.1 patches installed on your system. For information about installing Sun StorEdge Availability Suite 3.2.1 software and its latest patches, see *Sun StorEdge Availability Suite 3.2 Software Installation Guide*.

This chapter describes the procedures for configuring data replication with Sun StorEdge Availability Suite 3.2.1 software. The following topics are covered:

Task Summary of Replicating Data in a Sun StorEdge Availability Suite 3.2.1 Protection Group

This section summarizes the steps for configuring Sun StorEdge Availability Suite 3.2.1 data replication in a protection group.

TABLE 6–1 Administration Tasks for Sun StorEdge Availability Suite 3.2.1 Data Replication

Task	Description
Perform an initial configuration of the Sun StorEdge Availability Suite 3.2.1 software.	See "Initial Configuration of Sun StorEdge Availability Suite 3.2.1 Software" on page 62.
Create a protection group that is configured for Sun StorEdge Availability Suite 3.2.1 data replication.	See "How to Create and Configure a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 81.
Add a device group that is controlled by Sun StorEdge Availability Suite 3.2.1.	See "How to Add a Data Replication Device Group to a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 94.
Add application resource groups to the protection group.	See "How to Add an Application Resource Group to a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 90.
Replicate the protection group configuration to a secondary cluster.	See "How to Replicate the Sun StorEdge Availability Suite 3.2.1 Protection Group Configuration to a Partner Cluster" on page 100.
Activate the protection group.	See "How to Activate a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 101.
Verify the protection group configuration.	Perform a trial a switchover or takeover and test some simple failure scenarios before bringing your system online. See Chapter 8.
Check the runtime status of replication.	See "Checking the Runtime Status of Sun StorEdge Availability Suite 3.2.1 Data Replication" on page 107.
Detect failure.	See "Detecting Cluster Failure on a System That Uses Sun StorEdge Availability Suite 3.2.1 Data Replication" on page 111.
Migrate services by using a switchover.	See "Migrating Services That Use Sun StorEdge Availability Suite 3.2.1 With a Switchover" on page 112.
Migrate services by using a takeover.	See "Forcing a Takeover on Systems That Use Sun StorEdge Availability Suite 3.2.1" on page 115.
Recover data after forcing a takeover.	See "Recovering Sun StorEdge Availability Suite 3.2.1 Data After a Takeover" on page 118.

Overview of Sun StorEdge Availability Suite 3.2.1 Data Replication

This section provides an overview of Sun StorEdge Availability Suite 3.2.1 resource groups and outlines some limitations of Sun StorEdge Availability Suite 3.2.1 replication on clusters of more than two nodes.

Sun StorEdge Availability Suite 3.2.1 Lightweight Resource Groups

A device group that is controlled by Sun StorEdge Availability Suite 3.2.1 might be added to a protection group. Sun Cluster Geographic Edition software then creates a lightweight resource group for each device group. The name of the lightweight resource group has the following format:

AVS-device-group-name`-stor-rg`

For example, a device group named `avsdg` that is controlled by the Sun StorEdge Availability Suite 3.2.1 software has a lightweight resource group named `avsdg-stor-rg`.

The lightweight resource group collocates the logical host and the device group, a requirement of data replication with the Sun StorEdge Availability Suite 3.2.1 remote mirror software.

Each lightweight resource group contains two resources:

- A logical hostname resource for the local logical host that is used for replication of the device group. The name of this resource is the same as the hostname of the local logical host.
- A `HAStoragePlus` resource for controlling the collocation of the device group with the lightweight resource group. The name of this resource has the format *AVS-device-group-name*`-stor`.

For more information about lightweight resource groups, see the Sun StorEdge Availability Suite 3.2.1 documentation.

Sun StorEdge Availability Suite 3.2.1 Replication Resource Groups

In addition to the light weight resource group, when a device group controlled by the Sun StorEdge Availability Suite 3.2.1 software is added to a protection group, Sun Cluster Geographic Edition software creates a special replication resource for that

device group in the replication resource group. By monitoring these replication resource groups, the overall status of replication can be monitored. One replication resource group with one replication resource is created for each protection group.

The name of the replication resource group has the following format:

AVS-protection-group-name`-rep-rg`

The replication resource in the replication resource group monitors the replication status of the device group on the local cluster, which is reported by the Sun StorEdge Availability Suite 3.2.1 remote mirror software.

The name of the replication resource has the following format:

AVS-device-group-name`-rep-rs`

Initial Configuration of Sun StorEdge Availability Suite 3.2.1 Software

This section describes the initial steps you must perform before you can configure Sun StorEdge Availability Suite 3.2.1 replication in the Sun Cluster Geographic Edition product.

This section uses an example of a protection group, `avspg`, that is configured on a partnership that consists of two clusters, `cluster-paris` and `cluster-newyork`. An application, which is encapsulated in the `apprg1` resource group, is to be protected by the `avspg` protection group. The application data is held by some volume in the `avsdg` device group. These volumes can be Solaris Volume Manager volumes, VERITAS Volume Manager volumes, or raw device volumes.

The resource group, `apprg1`, and the device group, `avsdg`, are present on both `cluster-paris` and `cluster-newyork`. The application data is protected by `avspg` by data replication between `cluster-paris` and `cluster-newyork`.

Note – Replication of each device group requires one logical host on the local cluster and one logical host on the partner cluster.

You cannot use the slash character (`/`) in a cluster tag Sun Cluster Geographic Edition software. If you are using raw DID devices, you cannot use predefined DID device group names such as `dsk/s3`.

To use DIDs with raw device groups, complete the following procedure.

Sun StorEdge Availability Suite Volume Set

Before you can define the Sun StorEdge Availability Suite 3.2.1 volume set, you must determine the following:

- **The data volumes to be replicated**. Examples are `vol-data-paris` in `avsdg` on `cluster-paris` and `vol-data-newyork` in `avsdg` on `cluster-newyork`.
- **The bitmap volume that is needed for replication**. Examples are `vol-bitmap-paris` in `avsdg` on `cluster-paris` and `vol-bitmap-newyork` in `avsdg` on `cluster-newyork`.
- **The logical host to be used exclusively for replication of the device group** `avsdg`. Examples are the logical host `logicalhost-paris-1` on `cluster-paris` and the logical host `logicalhost-newyork-1` on `cluster-newyork`.

Note – The logical host that is used for Sun StorEdge Availability Suite 3.2.1 replication cannot be the same as the Sun Cluster Geographic Edition infrastructure logical host. For more information about configuring logical hostnames, see "Configuring Logical Hostnames" on page 28.

The `volset` file is located at `/var/cluster/geo/avs/`*device-group-name*`-volset.ini` on all the nodes of the protection group's primary cluster and secondary cluster. For example, the `volset` file for the device group `avsdg` would be located at `/var/cluster/geo/avs/avsdg-volset.ini`.

The fields in the volume set file that are handled by the Sun Cluster Geographic Edition software are described in the following table. The Sun Cluster Geographic Edition software does not handle other parameters of the volume set, including disk queue, size of memory queue, and number of asynchronous threads. You must adjust these parameters manually by using Sun StorEdge Availability Suite 3.2.1 commands.

Field	Meaning	Description
`phost`	Primary host	The logical host of the server on which the primary volume resides.
`pdev`	Primary device	Primary volume partition. Specify full path names only.
`pbitmap`	Primary bitmap	Volume partition in which the bitmap of the primary partition is stored. Specify full path names only.

Field	Meaning	Description
`shost`	Secondary host	The logical host of the server on which the secondary volume resides.
`sdev`	Secondary device	Secondary volume partition. Specify full path names only.
`sbitmap`	Secondary bitmap	Volume partition in which the bitmap of the secondary partition is stored. Specify full path names only.
`ip`	Network transfer protocol	Specify IP.
`sync` | `async`	Operating mode	`sync` is the mode in which the I/O operation is confirmed as complete only when the volume on the secondary cluster has been updated. `async` is the mode in which the primary host I/O operation is confirmed as complete before updating the volumes on the secondary cluster.
`g` *io-groupname*	I/O group name	An I/O group name can be specified by the `g` character. The set must be configured in the same I/O group on both the primary and the secondary cluster.
C	C tag	Specifies the device group name or resource tag of the local data and bitmap volumes in cases where this information is not implied by the name of the volume. For example, `/dev/md/avsset/rdsk/vol` indicates a device group named `avsset`. As another example, `/dev/vx/rdsk/avsdg/vol` indicates a device group named `avsdg`.

Sun Cluster Geographic Edition software does not modify the value of the Sun StorEdge Availability Suite 3.2.1 parameters. The software only controls the role of the volume set during switchover and takeover operations.

For more information about the format of the volume set files, refer to the Sun StorEdge Availability Suite 3.2.1 documentation.

▼ How to Use DIDs With Raw Device Groups

Steps

1. **Remove all the DIDs you want to use from its predefined DID device group.**
2. **Add the DIDs to a raw device group with a name that does not contain any slashes.**
3. **Create the same group name on each cluster of the partnership. You can use the same DIDs on each cluster.**
4. **Use this new name where a device group name is required.**

▼ How to Configure the Sun StorEdge Availability Suite 3.2.1 Volume in Sun Cluster

This procedure provides an example of how to configure Sun StorEdge Availability Suite 3.2.1 volumes in Sun Cluster. These volumes can be Solaris Volume Manager volumes, VERITAS Volume Manager volumes, or raw device volumes.

The volumes are encapsulated at the Sun Cluster device-group level. The Sun StorEdge Availability Suite 3.2.1 software interacts with the Solaris Volume Manager disksets, or VERITAS Volume Manager disk group, or raw device through this device group interface. The path to the volumes depends on the volume type as described in the following table.

Volume Type	Path
Solaris Volume Manager	`/dev/md/`*diskset-name*`/rdsk/d#`, where # represents a number
VERITAS Volume Manager	`/dev/vx/rdsk/`*disk-group-name*`/`*volume-name*
Raw device	`/dev/did/rdsk/d#s#`

Steps

1. **Create a diskset, `avsset`, by using Solaris Volume Manager or a disk group, `avsdg`, by using VERITAS Volume Manager or a raw device on `cluster-paris` and `cluster-newyork`.**

   For example, if you configure the volume by using a raw device, choose a raw device group, `dsk/d3`, on `cluster-paris` and `cluster-newyork`.

2. **Create two volumes in the diskset or disk group on `cluster-paris`.**

   The Sun StorEdge Availability Suite software requires a dedicated bitmap volume for each data volume to track which modifications to the data volume when the system is in logging mode.

   If you use a raw device to configure the volumes, create two partitions, `/dev/did/rdsk/d3s3` and `/dev/did/rdsk/d3s4`, on the `/dev/did/rdsk/d3` device on `cluster-paris`.

3. **Create two volumes in the diskset or disk group on `cluster-newyork`.**

   If you use a raw device to configure the volumes, create two partitions, `/dev/did/rdsk/d3s5` and `/dev/did/rdsk/d3s6`, on the `/dev/did/rdsk/d3` device on `cluster-paris`.

Enabling a Sun StorEdge Availability Suite 3.2.1 Volume Set

The Sun StorEdge Availability Suite 3.2.1 volume sets can be enabled in two ways:

- Automatically, when the device group is added to the protection group, `avspg`

  Use the automatic procedures to prepare the *device-group-name*`-volset.ini` file when you are setting up Sun StorEdge Availability Suite 3.2.1 software for the first time. After you have prepared the file, when you add the device group to the protection group, set the device group's `Enable_volume_set` property to `True`. The information in the *device-group-name*`-volset.ini` file will be read by the Sun StorEdge Availability Suite command to automatically enable the device group.

- Manually, after the device group is added to the protection group, `avspg`

  Use the manual procedures to enable the volume sets when you are creating volumes on a system that has already been configured.

Automatically Enabling a Solaris Volume Manager Volume Set

In this example, the `cluster-paris` cluster is the primary and `avsset` is a device group that contains a Solaris Volume Manager diskset.

EXAMPLE 6–1 Automatically Enabling a Solaris Volume Manager Volume Set

This example has the following entries in `/var/cluster/geo/avs/avsset-volset.ini`:

```
logicalhost-paris-1 /dev/md/avsset/rdsk/d100 /dev/md/avsset/rdsk/d101
logicalhost-newyork-1 /dev/md/avsset/rdsk/d100 /dev/md/avsset/rdsk/d101
ip async g - C avsset
```

EXAMPLE 6–1 Automatically Enabling a Solaris Volume Manager Volume Set *(Continued)*

The `avsset-volset.ini` file contains the following entries:

- `lh-paris-1` is the primary host.
- `/dev/md/avsset/rdsk/d100` is the primary data.
- `/dev/md/avsset/rdsk/d101` is the primary bitmap.
- `lh-newyork-1` is the secondary host.
- `/dev/md/avsset/rdsk/d100` is the secondary data.
- `/dev/md/avsset/rdsk/d101` is the secondary bitmap.
- `ip` is the protocol.
- `async` is the mode.
- `g` is the G flag.
- `-` is the IO group.
- `C` is the C tag.
- `avsset` is the diskset.

The sample configuration file defines a volume set that replicates `d100` from `cluster-paris` to `d100` on `cluster-newyork` by using the bitmap volumes and logical hostnames that are specified in the file.

Automatically Enabling a VERITAS Volume Manager Volume Set

In this example, the `cluster-paris` cluster is the primary and `avsdg` is a device group that contains a VERITAS Volume Manager disk group.

EXAMPLE 6–2 Automatically Enabling a VERITAS Volume Manager Volume Set

This example has the following entries in the `/var/cluster/geo/avs/avsdg-volset.ini` file:

```
logicalhost-paris-1 /dev/vx/rdsk/avsdg/vol-data-paris \
/dev/vx/rdsk/avsdg/vol-bitmap-paris
logicalhost-newyork-1 /dev/vx/rdsk/avsdg/vol-data-newyork \
/dev/vx/rdsk/avsdg/vol-bitmap-ny
ip async g - C avsdg
```

The `avsdg-volset.ini` file contains the following entries:

- `lh-paris-1` is the primary host.
- `/dev/vx/rdsk/avsdg/vol-data-paris` is the primary data.
- `/dev/vx/rdsk/avsdg/vol-bitmap-paris` is the primary bitmap.
- `lh-newyork-1` is the secondary host.
- `/dev/vx/rdsk/avsdg/vol-data-newyork` is the secondary data.
- `/dev/vx/rdsk/avsdg/vol-bitmap-ny` is the secondary bitmap.
- `ip` is the protocol.
- `async` is the mode.
- `g` is the G flag.

EXAMPLE 6–2 Automatically Enabling a VERITAS Volume Manager Volume Set *(Continued)*

- `-` is the IO group.
- `C` is the C flag.
- `avsdg` is the device group.

The sample configuration file defines a volume set that replicates `vol-data-paris` from `cluster-paris` to `vol-data-newyork` on `cluster-newyork`. The volume set uses the bitmap volumes and logical hostnames that are specified in the file.

Automatically Enabling a Raw Device Volume Set

In this example, the `cluster-paris` cluster is the primary and `rawdg` is the name of the device group that contains a raw device disk group, `/dev/did/rdsk/d3`.

EXAMPLE 6–3 Automatically Enabling a Raw Device Volume Set

This example has the following entries in `/var/cluster/geo/avs/avsdg-volset.ini` file:

```
logicalhost-paris-1 /dev/did/rdsk/d3s3 /dev/did/rdsk/d3s4
logicalhost-newyork-1 /dev/did/rdsk/d3s5 /dev/did/rdsk/d3s6
ip async g - C rawdg
```

The `rawdg-volset.ini` file contains the following entries:

- `logicalhost-paris-1` is the primary host.
- `/dev/did/rdsk/d3s3` is the primary data.
- `/dev/did/rdsk/d3s4` is the primary bitmap.
- `logicalhost-newyork-1` is the secondary host.
- `/dev/did/rdsk/d3s5` is the secondary data.
- `/dev/did/rdsk/d3s6` is the secondary bitmap.
- `ip` is the protocol.
- `async` is the mode.
- `g` is the G flag.
- `-` is the IO group.
- `C` is the C flag.
- `rawdg` is the device group.

The sample configuration file defines a volume set that replicates `d3s3` from `cluster-paris` to `d3s5` on `cluster-newyork`. The volume set uses the bitmap volumes and logical hostnames that are specified in the file.

Manually Enabling Volume Sets

After you have added the device group to the protection group, `avspg`, you can manually enable the Sun StorEdge Availability Suite 3.2.1 volume sets.

EXAMPLE 6–4 Manually Enabling the Sun StorEdge Availability Suite 3.2.1 Volume Set

The following example illustrates how to manually enable a Solaris Volume Manager volume set:

```
phys-paris-1# /usr/opt/SUNWesm/sbin/sndradm -e logicalhost-paris-1 \
/dev/md/avsset/rdsk/d100 /dev/md/avsset/rdsk/d101 \
logicalhost-newyork-1 /dev/md/avsset/rdsk/d100 \
/dev/md/avsset/rdsk/d101 ip async C avsset
```

EXAMPLE 6–5 Manually Enabling a VERITAS Volume Manager Volume Set

The following example illustrates how to manually enable a VERITAS Volume Manager volume set:

```
phys-paris-1# /usr/opt/SUNWesm/sbin/sndradm -e logicalhost-paris-1
/dev/vx/rdsk/avsdg/vol-data-paris /dev/vx/rdsk/avsdg/vol-bitmap-paris
logicalhost-newyork-1 /dev/vx/rdsk/avsdg/vol-data-newyork
/dev/vx/rdsk/avsdg/vol-bitmap-newyork ip async C avsdg
```

EXAMPLE 6–6 Manually Enabling a Raw Device Volume Set

The following example illustrates how to manually enable a raw device volume set:

```
phys-paris-1# /usr/opt/SUNWesm/sbin/sndradm -e logicalhost-paris-1
/dev/did/rdsk/d3s3 /dev/did/rdsk/d3s4 logicalhost-newyork-1 /dev/did/rdsk/d3s5
/dev/did/rdsk/d3s6 ip async C dsk/d3
```

Information about `sndradm` command execution is printed in the Sun StorEdge Availability Suite 3.2.1 log file, `/var/opt/SUNWesm/ds.log`. Refer to this file if you encounter errors while manually enabling the volume set.

▼ How to Configure the Sun Cluster Device Group That Is Controlled by Sun StorEdge Availability Suite 3.2.1

Sun StorEdge Availability Suite 3.2.1 software supports Solaris Volume Manager, VERITAS Volume Manager, and raw device volumes.

Steps

1. **Ensure that the device group that contains the volume set that you want to replicate is register with Sun Cluster.**

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

2. **If you are using a VERITAS Volume Manager device group, synchronize the VERITAS Volume Manager configuration by using one of the Sun Cluster commands, `scsetup` or `scconf`.**

3. **After you have finished configuring the device group, it should be displayed in the output of the `scstat -D` command.**

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

4. **Repeat steps 1–3 on both `cluster-paris` and `cluster-newyork`.**

▼ How to Configure a Highly Available Cluster Global File System for Use With Sun StorEdge Availability Suite 3.2.1

Steps

1. **Create the required file system on the volume set that you created in the previous step, `vol-data-paris`.**

   The application writes to this file system.

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

   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. **To handle the new file system, add the `HAStoragePlus` resource to the application resource group, `apprg1`.**

   Adding this resource ensures that the necessary file systems are remounted before the application is started.

   For more information about the `HAStoragePlus` resource type, refer to the *Sun Cluster 3.1 Data Service Planning and Administration Guide*.

4. **Repeat steps 1–3 on both `cluster-paris` and `cluster-newyork`.**

Example 6–7 Configuring a Highly Available Cluster Global File System for Solaris Volume Manager Volumes

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

1. Create a UNIX file system (UFS).

```
phys-paris-1# newfs /dev/md/avsset/rdsk/d100
```

2. An entry is created in the /etc/vfstab file as follows.

```
/dev/md/avsset/dsk/d100 /dev/md/avsset/rdsk/d100
/global/sample ufs 2 no logging
```

3. Add the HAStoragePlus resource.

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

Example 6–8 Configuring a Highly Available Cluster Global File System for VERITAS Volume Manager Volumes

This example assumes that the apprg1 resource group already exists.

1. Create a UNIX file system (UFS).

```
phys-paris-1# newfs /dev/vx/rdsk/avsdg/vol-data-paris
```

2. An entry is created in the /etc/vfstab file as follows:

```
/dev/vx/dsk/avsdg/vol-data-paris /dev/vx/rdsk/avsdg/vol-data-paris
/global/sample ufs 2 no logging
```

3. Add the HAStoragePlus resource.

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

Example 6–9 Configuring a Highly Available Cluster Global File System for Raw Device Volumes

This example assumes that the apprg1 resource group already exists.

1. Create a UNIX file system (UFS).

```
phys-paris-1# newfs /dev/did/rdsk/d3s3
```

2. An entry is created in the /etc/vfstab file as follows:

```
/dev/did/dsk/d3s3 /dev/did/rdsk/d3s3
/global/sample ufs 2 no logging
```

3. Add the HAStoragePlus resource.

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

CHAPTER 7

Administering Sun StorEdge Availability Suite 3.2.1 Protection Groups

This chapter describes the procedures for administering data replication with Sun StorEdge Availability Suite 3.2.1 software. The following topics are covered:

Strategies for Creating Sun StorEdge Availability Suite 3.2.1 Protection Groups

Before you begin creating protection groups, consider which of the following strategies is best for you:

- Stopping the application 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.

Before you create a protection group by using the steps in the following sections, ensure that the following prerequisites are met.

- The application has been configured by Sun Cluster software on both cluster.
- Corresponding device groups are configured for data replication.

Creating a Protection Group While the Application Is Offline

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

- Create the protection group from a node on one cluster.

  For more information, see "How to Create and Configure a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 81.

- Add the data replication device group to the protection group.

  For more information, see "How to Add a Data Replication Device Group to a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 94.

- Take the application resource group to the unmanaged state.
- Add the application resource group to the protection group on one cluster.

  For more information, see "How to Add an Application Resource Group to a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 90.

- On the other cluster, retrieve the protection group configuration.

  For more information, see "How to Replicate the Sun StorEdge Availability Suite 3.2.1 Protection Group Configuration to a Partner Cluster" on page 100.

- From either cluster, activate the protection group globally.

  For more information, see "How to Activate a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 101.

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 node on one cluster.

  For more information, see "How to Create and Configure a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 81.
- Add the data replication device group to the protection group

  For more information, see "How to Add a Data Replication Device Group to a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 94.
- Activate the protection group locally.

  For more information, see "How to Activate a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 101.
- Add the application resource group to the protection group.

  For more information, see "How to Add an Application Resource Group to a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 90.

Complete the following steps on the other cluster.

- Retrieve the protection group configuration.

  For more information, see "How to Replicate the Sun StorEdge Availability Suite 3.2.1 Protection Group Configuration to a Partner Cluster" on page 100.
- Activate the protection group locally.

  For more information, see "How to Activate a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 101.

EXAMPLE 7–1 Creating a Sun StorEdge Availability Suite 3.2.1 Protection Group While the Application Remains Online

This example describes how to create 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`.

   First, on a node of the cluster where an application resource group is already running, check and fix any `Nodelist` inconsistencies that may exist between the resource group and the device group with which the resource group has affinities. Then, create the protection group with a matching `Nodelist`.

   ```
   phys-paris-1# scrgadm -pvv -g apprg1 | grep Nodelist

   (apprg1) Res Group Nodelist:   phys-paris-1 phys-paris-2
   phys-paris-1# scconf -pvv | grep "avsdg.*group node list"
   (avsdg) Device group node list: phys-paris-2, phys-paris-1
   ```

EXAMPLE 7–1 Creating a Sun StorEdge Availability Suite 3.2.1 Protection Group While the Application Remains Online *(Continued)*

The node list of the device group is in a different order from the resource group's `Nodelist`. The order of the device group's node list is changed as follows:

```
phys-paris-1# scconf -c -D name=avsdg,nodelist=phys-paris-1:phys-paris-2
```

You can also change the node list of a resource group to meet this requirement.

The protection group is created with a `Nodelist` identical to the `Nodelist` of the resource group and device group:

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

2. Add the Sun StorEdge Availability Suite 3.2.1 device group, `avsdg`, to the protection group.

```
phys-paris-1# geopg add-device-group -p Local_logical_host=lh-paris-1 \
-p Remote_logical_host=lh-newyork-1 -p Enable_volume_set=True avsdg avspg
```

3. Verify that the data replication resource groups and the light weight resource groups have been created and brought online. Also, verify that the Sun StorEdge Availability Suite 3.2.1 volume set has been enabled, because the `Enable-volume-set` property has been set to `True`.

```
phys-paris-1# dsstat
name              t      s       pct     role    ckps    dkps    tps    svt
/avsdg/rdsk/d100 P      L      100.00   net      -       0       0      0
/avsdg/rdsk/d101                        bmp      0       0       0      0

phys-paris-1# scstat -g

-- Resource Groups and Resources --
       Group Name              Resources
       ----------              ---------
Resources: geo-clusterstate     -
Resources: geo-infrastructure  geo-clustername geo-hbmonitor geo-failovercontrol
Resources: avsdg-stor-rg       lh-paris-1 avsdg-stor
Resources: avspg-rep-rg        avsdg-rep-rs
Resources: apprg1              avs-lh avs-stor avs-server-res avs-listener-res


-- Resource Groups --

            Group Name             Node Name      State
            ----------             ---------      -----
Group: geo-clusterstate          phys-paris-1     Online
Group: geo-clusterstate          phys-paris-2     Online

Group: geo-infrastructure        phys-paris-1     Online
Group: geo-infrastructure        phys-paris-2     Offline

Group: avsdg-stor-rg             phys-paris-1     Online
```

EXAMPLE 7–1 Creating a Sun StorEdge Availability Suite 3.2.1 Protection Group While the Application Remains Online *(Continued)*

```
Group: avsdg-stor-rg            phys-paris-2     Offline

Group: avspg-rep-rg             phys-paris-1     Online
Group: avspg-rep-rg             phys-paris-2     Offline

Group: apprg1                   phys-paris-1     Online
Group: apprg1                   phys-paris-2     Offline

-- Resources--

Resource Name                   Node Name      State    Status Message
-------------                 ---------      ------   --------------
Resource: geo-clustername     phys-paris-1  Online    Online - \
LogicalHostname online
Resource: geo-clustername     phys-paris-2  Offline   Offline - \
LogicalHostname offline
Resource: geo-hbmonitor       phys-paris-1  Online    Online - \
Daemon OK
Resource: geo-hbmonitor       phys-paris-2  Offline   Offline

Resource: geo-failovercontrol phys-paris-1  Online    Online
Resource: geo-failovercontrol phys-paris-2  Offline   Offline

Resource: lh-paris-1          phys-paris-1  Online    Online - \
LogicalHostname online
Resource: lh-paris-1          phys-paris-2  Offline   Offline

Resource: avsdg-stor          phys-paris-1  Online    Online
Resource: avsdg-stor          phys-paris-2  Offline   Offline

Resource: avsdg-rep-rs        phys-paris-1  Online    Degraded - \
Logging
Resource: avsdg-rep-rs        phys-paris-2  Offline   Offline

Resource: avs-lh              phys-paris-1  Online    Online - \
LogicalHostname online
Resource: avs-lh              phys-paris-2  Offline   Offline

Resource: avs-server-res      phys-paris-1  Online    Online
Resource: avs-server-res      phys-paris-2  Offline   Offline

Resource: avs-listener-res    phys-paris-1  Online    Online
Resource: avs-listener-res    phys-paris-2  Offline   Offline
```

4. Activate the protection group locally.

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

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

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

EXAMPLE 7–1 Creating a Sun StorEdge Availability Suite 3.2.1 Protection Group While the Application Remains Online *(Continued)*

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

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

   Protection group "avspg"               : Unknown
      Partnership                         : paris-newyork-ps
      Synchronization                     : Error

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

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

6. On one node of the partner cluster, retrieve the protection group as follows:

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

7. Verify that the data replication resource groups and the light weight resource groups have been created and brought online as follows:

```
phys-newyork-1# dsstat
name                 t  s   pct    role   ckps   dkps   tps   svt
/avsdg/rdsk/d100     S  L  100.00  net     -      0      0     0
/avsdg/rdsk/d101                    bmp     0      0      0     0

phys-newyork-1# scstat -g

-- Resource Groups and Resources --

    Group Name                 Resources
    ----------                 ---------
Resources: geo-clusterstate        -
```

EXAMPLE 7–1 Creating a Sun StorEdge Availability Suite 3.2.1 Protection Group While the Application Remains Online *(Continued)*

```
Resources: geo-infrastructure  geo-clustername geo-hbmonitor \
geo-failovercontrol
Resources: avsdg-stor-rg       lh-newyork-1 avsdg-stor
Resources: avspg-rep-rg        avsdg-rep-rs
Resources: apprg1              avs-lh avs-stor avs-server-res avs-listener-res

-- Resource Groups --

     Group Name                  Node Name        State
     ----------                  ---------        -----
Group: geo-clusterstate             phys-newyork-1   Online
Group: geo-clusterstate          phys-newyork-2   Online

Group: geo-infrastructure        phys-newyork-1   Online
Group: geo-infrastructure        phys-newyork-2   Offline

Group: avsdg-stor-rg             phys-newyork-1   Online
Group: avsdg-stor-rg             phys-newyork-2   Offline

Group: avspg-rep-rg              phys-newyork-1   Online
Group: avspg-rep-rg              phys-newyork-2   Offline

Group: apprg1                    phys-newyork-1   Unmanaged
Group: apprg1                    phys-newyork-2   Unmanaged

-- Resources --

Resource Name                Node Name       State     Status Message
-------------                ---------       ------    --------------
Resource: geo-clustername    phys-newyork-1  Online    Online - \
LogicalHostname online
Resource: geo-clustername    phys-newyork-2  Offline   Offline - \
LogicalHostname offline

Resource: geo-hbmonitor      phys-newyork-1  Online    Online - Daemon OK
Resource: geo-hbmonitor      phys-newyork-2  Offline   Offline

Resource: geo-failovercontrol phys-newyork-1  Online    Online
Resource: geo-failovercontrol phys-newyork-2  Offline   Offline

Resource: lh-newyork-1       phys-newyork-1  Online    Online - \
LogicalHostname online
Resource: lh-newyork-1       phys-newyork-2  Offline   Offline

Resource: avsdg-stor         phys-newyork-1  Offline   Offline
Resource: avsdg-stor         phys-newyork-2  Offline   Offline

Resource: avsdg-rep-rs       phys-newyork-1  Online    Degraded - Logging
Resource: avsdg-rep-rs       phys-newyork-2  Offline   Offline

Resource: avs-lh             phys-newyork-1  Offline   Offline
Resource: avs-lh             phys-newyork-2  Offline   Offline
```

EXAMPLE 7–1 Creating a Sun StorEdge Availability Suite 3.2.1 Protection Group While the Application Remains Online *(Continued)*

```
Resource: avs-server-res      phys-newyork-1  Offline   Offline
Resource: avs-server-res      phys-newyork-2  Offline   Offline

Resource: avs-listener-res    phys-newyork-1  Offline   Offline
Resource: avs-listener-res    phys-newyork-2  Offline   Offline
```

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

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

9. Verify that the protection group was successfully created an 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    : cluster-newyork
  Synchronization     : OK

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

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

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

  Cluster cluster-newyork : OK
     Role                 : Secondary
     PG activation state  : Activated
     Configuration        : OK
     Data replication     : OK
     Resource groups      : OK
```

Creating, Modifying, Validating, and Deleting a Sun StorEdge Availability Suite 3.2.1 Protection Group

This section contains procedures for the following tasks:

- "How to Create and Configure a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 81
- "How to Modify a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 83
- "How to Validate a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 84
- "How to Delete a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 86

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` *data-replication-type* option when you use the `geopg` command. The `geoadm status` command shows a state for data replication of `NONE`.

▼ How to Create and Configure a Sun StorEdge Availability Suite 3.2.1 Protection Group

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 Sun StorEdge Availability Suite 3.2.1 Protection Group Configuration to a Partner Cluster" on page 99

Steps

1. **Log in to one of the cluster nodes.**

   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" on page 43.

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 partnership-name -d avs \
   -o local-role [-p property-settings [-p...]] \
   protection-group-name
   ```

`-s` *partnership-name*	Specifies the name of the partnership
`-d avs`	Specifies that the protection group data is replicated by Sun StorEdge Availability Suite 3.2.1
`-o` *local-role*	Specifies the role of this protection group on the local cluster as either `primary` or `secondary`
`-p` *property-setting*	Sets the properties of the protection group The properties you can set are the following: ■ `Description` – describes the protection group ■ `Timeout` – specifies the timeout period for the protection group in seconds ■ `Enable_volume_set` – when set to `true`, specifies that the volume set will be automatically enabled at protection group creation time For more information about automatically enabling a volume set, see "Enabling a Sun StorEdge Availability Suite 3.2.1 Volume Set" on page 66. ■ `Nodelist` – Lists the host names of the machines that can be primary for the device group in the protection group. For more information about the properties you can set, see Appendix A.
protection-group-name	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.

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

   Before creating the protection group, the data replication layer validates that the configuration is correct.

If the validation is successful, the local `Configuration` status is set to `OK` and the `Synchronization` status is set to `Error`.

If the validation is unsuccessful, the protection group is not created.

Example 7–2 Creating and Configuring a Protection Group

The following example illustrates how to create a Sun StorEdge Availability Suite 3.2.1 protection group on `cluster-paris`, which is set as the primary cluster:

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

▼ How to Modify a Sun StorEdge Availability Suite 3.2.1 Protection Group

Before You Begin

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

Steps

1. **Log in to one of the cluster nodes.**

   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" on page 43.

2. **Modify the configuration of the protection group.**

   This command modifies the properties of a protection group on 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-settings[-p...] protection-group-name
   ```

`-p` *property-setting*	Sets the properties of the protection group For more information about the properties you can set, see Appendix A.
protection-group-name	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.

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

Example 7–3 Modifying the Configuration of a Protection Group

The following example illustrates how to modify the `timeout` property of the protection group that was created in Example 7–2:

```
# geopg set-prop -p Timeout=300 avspg
```

More Information How the Data Replication Subsystem Validates the Modified Protection Group

When you launch the `geopg set-prop` command, the data replication subsystem revalidates the protection group with the new configuration information. If the validation 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 validation is in unsuccessful on the partner cluster, the `Configuration` status is set to `Error` on the partner cluster.

▼ How to Validate a Sun StorEdge Availability Suite 3.2.1 Protection Group

Before You Begin When the `Configuration` status of a protection group is displayed as `Error` in the `geoadm status` output, 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 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 issue the`geopg validate` command again.

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

Before validating the configuration of a protection group, 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.

Steps

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" on page 43.

2. **Validate the configuration of the protection group.**

   This command validates the configuration of the protection group on the local cluster only.

   ```
   # geopg validate protection-group-name
   ```

   protection-group-name Specifies a unique name that identifies a single protection group.

Example 7–4 Validating the Configuration of a Protection Group

The following example illustrates how to validate a protection group:

```
# geopg validate avspg
```

How the Data Replication Layer Validates the Application Resource Groups and Data Replication Entities

During protection group validation, the Sun StorEdge Availability Suite 3.2.1 data replication layer validates the application resource groups and the data replication entities as follows:

- Verifies that a application resource group in the protection group has its `Auto_start_on_new_cluster` property set to `False`.
- Verifies that the `Nodelist` property of an application resource group that has affinities with a device group defined by the `HAStoragePlus` resource contains the same entries in the same order as the `Nodelist` property of the protection group.
- Verifies that the `Nodelist` property of a device group in the protection group contains the same entries in the same order as the `Nodelist` property of the protection group.
- Verifies that a lightweight resource group is created for each device group in the protection group. Each lightweight resource group contains two resources, a logical hostname resource and a `HAStoragePlus` resource. For more information about lightweight resource groups and their resources, see "Sun StorEdge Availability Suite 3.2.1 Lightweight Resource Groups" on page 61.

- Verifies that a replication resource of the type `GeoCtlAVS` is created in the replication resource group of each device group in the protection group. For information about the format of the replication resource group, see "Sun StorEdge Availability Suite 3.2.1 Replication Resource Groups" on page 61.
- Verifies that the `Nodelist` property of the lightweight resource group and replication resource group contains the same entries in the same order as the `Nodelist` property of the protection group.

If the `Enable_volume_set` property of a successfully validated device group is set to `True`, then volume sets defined in the `/var/cluster/geo/avs/avsdg-volset.ini` file are enabled. Other volume sets for the device group are disabled. If you want to enable the other volume sets, you can add the volume sets to the `/var/cluster/geo/avs/avsdg-volset.ini` file or set the `Enable_volume_set` property to `False`.

When validation is complete, the Sun Cluster Geographic Edition software creates the lightweight resource group, the replication resource group, and the resources for this replication resource group, if nonexistent, and brings them online. If a resource group or resource of the same name already exists, the Sun Cluster Geographic Edition operations may modify their properties. Sun Cluster Geographic Edition software cannot create a new resource group or resource of the same name if one already exists.

The `Configuration` status is set to `OK` after successful validation. If validation is not successful, the `Configuration` status is set to `Error`.

▼ How to Delete a Sun StorEdge Availability Suite 3.2.1 Protection Group

Before You Begin

To delete the protection group on every cluster, 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 exists locally.
- The protection group is offline on the local cluster.

Note – To keep the application resource groups online while deleting a protection group, you must remove the application resource groups from the protection group.

Steps

1. **Log in to one of the nodes on the cluster where you want to delete the protection group, for example `cluster-paris`.**

   The `cluster-paris` is the primary cluster. See Figure 2–1 for For a a sample cluster configuration.

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" on page 43.

2. **Delete the protection group.**

   This command deletes the configuration of the protection group from the local cluster. The command also removes the lightweight resource group and the replication resource group for each device group in the protection group.

   ```
   # geopg delete protection-group-name
   ```

   protection-group-name Specifies the name of the protection group

   If the deletion is unsuccessful, the `Configuration` status is set to `Error`. Fix the cause of the error and reissue the `geopg delete` command.

Example 7–5 Deleting a Protection Group

The following example illustrates how to delete a protection group from both partner clusters.

```
# rlogin cluster-paris -l root
cluster-paris# geopg delete avspg
# rlogin cluster-newyork -l root
cluster-newyork# geopg delete avspg
```

Example 7–6 Deleting a Protection Group While Keeping Application Resource Groups Online

The following example illustrates how to keep two application resource groups (`apprg1` and `apprg2`) online while deleting their protection group, `avspg`. Remove the application resource groups from the protection group, then delete the protection group.

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

Creating a Protection Group That Does Not Require Data Replication

Some of the protection groups do not require data replication. If you are using the Sun Cluster Geographic Edition software to manage only resource groups and to handle data replication differently, you can create protection groups that do not replicate data. The `geoadm status` command displays that these protection groups are in the `Degraded` state. This section describes how to configure your protection group to not use data replication.

For information about how to create a Sun StorEdge Availability Suite 3.2.1 protection group that requires data replication, see"How to Create and Configure a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 81.

Note – You cannot add device groups to a protection group that does not use data replication.

▼ How to Create a Protection Group That Is Configured to Not Use Data Replication

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 that 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 more than one partnership on the same system.

Steps

1. **Log in to one of the cluster nodes.**

   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" on page 43.

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

   This command creates a protection group on the local cluster.

   ```
   # geopg create -s partnership-name -o local-role \
   [-p property-settings [-p...]] \
   protection-group-name
   ```

`-s` *partnership-name*	Specifies the name of the partnership
`-o` *local-role*	Specifies the role of this protection group on the local cluster as either `Primary` or `Secondary`
`-p` *property-setting*	Sets the properties of the protection group The properties you can set are the following: ■ `Description` – Describes the protection group. ■ `Timeout` – Specifies the timeout period for the protection group in seconds. ■ `RoleChange_ActionArgs` – Specifies a string that follows system-defined arguments at the end of the command line when the role-change callback command runs. ■ `RoleChange_ActionCmd` – Specifies the absolute path to the executable command to run when the primary cluster of the protection group changes. This path should be valid on all partner clusters that host the protection group. For more information about the properties you can set, see Appendix A.
protection-group-name	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.

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

Example 7–7 Creating and Configuring a Protection Group That Is Configured to Not Use Data Replication

The following example illustrates how to create an protection group that is configured to not use data replication.

```
# geopg create -s paris-newyork-ps -o primary example-pg
```

Next Steps See "Administering Sun StorEdge Availability Suite 3.2.1 Application Resource Groups" on page 90 for information on adding resource groups to a protection group.

Administering Sun StorEdge Availability Suite 3.2.1 Application Resource Groups

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

All of the entities that you configure for the application resource group on the primary cluster, such as application data resources, 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 Sun StorEdge Availability Suite 3.2.1 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 to add already 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 apprg1 | grep Auto_start_on_new_cluster
  ```

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

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

  - The `Nodelist` property of the failover application resource group that has affinities with a device group defined by the resource must contain the same entries in the same order as the `Nodelist` property of the protection group.

- 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 all the application resource groups that share dependencies to the protection group in a single operation. If you add the application resource groups separately, the operation will fail.

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

If the resource group is `Offline` and the protection group is activated after the configuration of the protection group has changed, then the protection group's local state becomes `Error`.

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 activate resource group.

Steps

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" on page 43.

2. **Add an application resource group to the protection group.**

   `# geopg add-resource-group` *resource-group-list protection-group*

resource-group-list	Specifies the name of the application resource group You can specify more than one resource group in a comma-separated list.
protection-group	Specifies the name of 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.

   For information about the names and values that are supported by Sun Cluster Geographic Edition software, see Appendix B.

   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.

If the application resource group is a failover type resource group that shares affinities with one of the device groups in the same protection group, the Sun Cluster Geographic Edition software alters its `RG_affinities` property to include a strong, positive affinity to an internal resource group, called a lightweight resource group. This affinity includes failover delegation.

The application resource group must not have strong, positive affinities that have failover delegation with other resource groups. Otherwise, trying to include a strong positive affinity with failover delegation on the lightweight resource group fails.

The Sun Cluster Geographic Edition software also creates strong dependencies between the `HAStoragePlus` resource in the application resource group and the `HAStoragePlus` resource in the lightweight resource group for this device group. This redirection occurs when the protection group is brought online or when an online application resource group is added to an online protection group.

Do not modify dependencies and resource group affinities between application resource groups and light weight resource groups.

Example 7–8 Adding an Application Resource Group to a Sun StorEdge Availability Suite 3.2.1 Protection Group

The following example illustrates how to add two application resource groups, `apprg1` and `apprg2`, to `avspg`:

```
# geopg add-resource-group apprg1,apprg2 \
avspg
```

▼ How to Delete an Application Resource Group From a Sun StorEdge Availability Suite 3.2.1 Protection Group

Before You Begin

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

Before you remove an application resource group from a protection group, 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 protection group's application resource groups. For example, you cannot remove a resource group that belongs to the data replication management entity.

Steps

1. **Log in to one of the cluster nodes.**

   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" on page 43.

2. **Remove the application resource group from the protection group.**

   This command removes an application resource group from a protection group on the local cluster. If the partner cluster contains a protection group of the same name, the application resource group is also removed from the partner cluster's protection group.

   ```
   # geopg remove-resource-group resource-group-list \
    protection-group
   ```

   resource-group-list — Specifies the name of the application resource group

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

   protection-group — Specifies the name of the protection group

   If the resource group being removed shares dependencies with other resource groups in the protection group, then you must remove all the resource groups that share dependencies in the same command.

   If the remove operation failed 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.

   Sun Cluster Geographic Edition software removes the affinity and resource dependencies between the application resource group and the lightweight resource group.

Example 7–9 Deleting an Application Resource Group From a Protection Group

The following example illustrates how to remove two application resource groups, `apprg1` and `apprg2`, from `avspg`:

```
# geopg remove-resource-group apprg1,apprg2 \
avspg
```

Administering Sun StorEdge Availability Suite 3.2.1 Data Replication Device Groups

This section describes the following tasks for administering data replication device groups in a Sun StorEdge Availability Suite 3.2.1 protection group:

- "How to Add a Data Replication Device Group to a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 94
- "How to Modify a Sun StorEdge Availability Suite 3.2.1 Data Replication Device Group" on page 98
- "How to Delete a Data Replication Device Group From a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 98

For details about configuring a Sun StorEdge Availability Suite 3.2.1 protection group, see "How to Create and Configure a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 81.

▼ How to Add a Data Replication Device Group to a Sun StorEdge Availability Suite 3.2.1 Protection Group

Before You Begin

A protection group is the container for the application resource groups, which contain data for services that are protected from disaster. Sun Cluster Geographic Edition software protects the data by replicating it from the primary cluster to the secondary cluster. By adding a Sun Cluster device group to a protection group, Sun Cluster Geographic Edition software monitors the replication status of all volumes in the device group that belong to a Sun StorEdge Availability Suite 3.2.1 volume set. Sun Cluster Geographic Edition software also controls the role and state of the volume set during protection group operations like start, stop, switchover, and takeover.

Before you add a device group to a protection group, ensure that the following conditions are met:

- The protection group is defined on the local cluster.
- The protection group is offline on the local cluster and the partner cluster, if the partner cluster can be reached.
- The device group exists on both the local cluster and the partner cluster.
- The `Nodelist` property of the device group contains the same entries in the same order as the `Nodelist` property of the protection group.

- The `Local_logical_host` property specifies a valid hostname that can be hosted by the local cluster and that is reserved for this device group.
- The `Remote_logical_host` property specifies a valid hostname that can be hosted by the remote cluster and that has been reserved for this device group.
- If the `Enable_volume_set` property is set to `true`, then the `/var/cluster/geo/avs/avsdg-volset.ini` file must exist and contain valid entries on all nodes of both partner clusters. For information about configuring this file, see "Enabling a Sun StorEdge Availability Suite 3.2.1 Volume Set" on page 66.

Steps

1. **Log in to one of the cluster nodes.**

   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" on page 43.

2. **Add a data replication device group to 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-settings [-p...] \
   AVS-device-group-name  protection-group-name
   ```

`-p` *property-settings*	Sets the properties of the data replication device group The Sun StorEdge Availability Suite 3.2.1 specific properties you can set are the following: ■ `Local_logical_host` – Specifies the name of the local logical host that is used to replicate the device group ■ `Remote_logical_host` – Specifies the name of the remote logical host that is used to replicate the device group ■ `Enable_volume_set` – Specifies whether the volume sets that are given in the file will be automatically enabled. Set to either `True` or `False` For more information about the properties you can set, see Appendix A.
AVS-device-group-name	Specifies the name of the new data replication device group
protection-group-name	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.

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

Example 7–10 Adding a Data Replication Device Group to a Sun StorEdge Availability Suite 3.2.1 Protection Group

The following example illustrates how to create a Sun StorEdge Availability Suite 3.2.1 data replication device group in the `avspg` protection group:

```
# geopg add-device-group -p Local_logical_host=lh-paris-1 \
-p Remote_logical_host=lh-newyork-1 avsdg avspg
```

How the Data Replication Subsystem Verifies the Device Group

When the device group controlled by Sun StorEdge Availability Suite 3.2.1 is added to a protection group, the data replication layer verifies that the device group exists and that the value of its `Nodelist` property contains the same entries in the same order as the `Nodelist` property of the protection group.

When the `geopg add-device-group` command runs, a lightweight resource group for the device group is created and brought online. The lightweight resource group contains the following resources:

- A logical hostname resource that is used for data replication, as specified in the `Local_logical_host` property
- A `HAStoragePlus` resource that controls the collocation of the device group with the lightweight resource group

For more information about lightweight resource groups and their resources, see "Sun StorEdge Availability Suite 3.2.1 Lightweight Resource Groups" on page 61.

When you run the `geopg add-device-group` command, a replication resource of the type `GeoCtlAVS` is created in the replication resource group of each device group in the protection group and brought online. For information about the format of the replication resource group, see "Sun StorEdge Availability Suite 3.2.1 Replication Resource Groups" on page 61.

The `Nodelist` property of the lightweight resource group and replication resource group contains the same entries in the same order as the `Nodelist` property of the protection group.

If a resource or resource group of the same name is already configured on the local cluster, Sun Cluster Geographic Edition verifies the configuration and sets the `Configuration` to `Error` if the configuration is not correct.

If the `Enable_volume_set` property of this device group is set to `True`, then volume sets that are defined in the `/var/cluster/geo/avs/AVS-devicegroup-volset.ini` file are enabled. Otherwise, Sun Cluster Geographic Edition software controls and monitors all volume sets that you enable manually by using the Sun StorEdge Availability Suite 3.2.1 commands.

If the `geopg add-device-group` command is unsuccessful, the configuration of the protection group is not modified.

If the `geopg add-device-group` command is successful and the `Configuration` status on the local cluster is set to `OK`, then the new configuration is propagated to the partner cluster. This propagation causes the whole protection group configuration to be revalidated on the partner cluster. During revalidation, the same entities are created on the partner cluster, including the lightweight resource group and the replication resource group. Volume sets are also enabled on the partner clusters if the `/var/cluster/geo/avs/AVS-devicegroup-volset.ini` file exists on the partner cluster and contains correctly defined volume sets. If the validation is unsuccessful, then the `Configuration` status on the partner cluster is set to `Error` on the partner cluster.

Caution – Do not change, remove, or bring offline these resources or resource groups. Lightweight resource groups, replication resource groups, and their resources are internal entities managed by Sun Cluster Geographic Edition software and should be administered by using only the Sun Cluster Geographic Edition commands. Altering the configuration or state of these entities directly with Sun Cluster commands could result in unrecoverable failure.

Conditionally, if the validation of the device group on the partner cluster is successful, the volume sets defined in the `/var/cluster/geo/avs/AVS-devicegroup-volset.ini` file are enabled on the partner cluster. The `Enable_volume_set` property of this device group must be set to `true`. Other volume sets of the device group are disabled.

After the device group has been added to the protection group, you can enable or disable the volume sets of the device group directly by using the Sun StorEdge Availability Suite 3.2.1 commands. The `/var/cluster/geo/avs/AVS-devicegroup-volset.ini` file is only consulted when the protection group that contains the device group is successfully validated for the first time.

▼ How to Modify a Sun StorEdge Availability Suite 3.2.1 Data Replication Device Group

Steps

1. **Log in to one of the cluster nodes.**

   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" on page 43.

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-settings [-p...] \
   AVS-device-group-name  protection-group-name
   ```

`-p` *property-settings*	Sets the properties of the data replication device group For more information about the properties you can set, see Appendix A.
AVS-device-group-name	Specifies the name of the new data replication device group
protection-group-name	Specifies the name of the protection group that will contain the new data replication device group

▼ How to Delete a Data Replication Device Group From a Sun StorEdge Availability Suite 3.2.1 Protection Group

Before You Begin

You might need to 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.

Before you remove a data replication device group, ensure that the following conditions are met:

- The protection group is defined on the local cluster.
- The protection group is offline on the local cluster and the partner cluster, if the partner cluster can be reached.
- The device group is managed by the protection group.

For information about deleting protection groups, refer to "How to Delete a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 86.

Steps

1. **Log in to one of the cluster nodes.**

   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" on page 43.

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.

   This command removes the device group from the protection group. This command also disables all volume sets associated with the device group and deletes the lightweight resource group and replication resource group for this device group.

   ```
   # geopg remove-device-group AVS-device-group-name protection-group-name
   ```

AVS-device-group-name	Specifies the name of the data replication device group
protection-group-name	Specifies the name of the protection group

Example 7–11 Deleting a Replication Device Group From a Sun StorEdge Availability Suite 3.2.1 Protection Group

The following example illustrates how to delete a data replication device group from a Sun StorEdge Availability Suite 3.2.1 protection group:

```
# geopg remove avsdg avspg
```

Replicating the Sun StorEdge Availability Suite 3.2.1 Protection Group Configuration to a Partner Cluster

You can replicate the configuration of a protection group to the partner cluster either before or after you configure data replication, resource groups, and resources on both clusters.

▼ How to Replicate the Sun StorEdge Availability Suite 3.2.1 Protection Group Configuration to a Partner Cluster

Before You Begin

Before you replicate the configuration of a Sun StorEdge Availability Suite 3.2.1 protection group to a partner 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 groups is set to `False`. You can view this property by using the `scrgadm` command.

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

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

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

Steps

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" on page 43.

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.

   ```
   # geopg get -s partnership-name \
   AVS-protection-group
   ```

`-s` *partnership-name*	Specifies the name of the partnership from which the protection group configuration information should be gathered.
AVS-protection-group	Specifies the name of the protection group.

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 7–12 Replicating the Sun StorEdge Availability Suite 3.2.1 Protection Group to a Partner Cluster

This example illustrates how to replicate the configuration of `avspg` to `cluster-newyork`.

The configuration of the protection group is retrieved from the remote cluster, in this example `cluster-paris` and then validated by the data replication subsystem on the local cluster, `cluster-newyork`.

If the validation is successful, the `Configuration` status is set to `OK` and the protection group is created on the local cluster. This protection group contains a device group and an application group that are configured identically to the device group and application group on the remote cluster.

If the validation fails, the `Configuration` status is set to `Error`. Fix the cause of the error and revalidate the protection group or delete the invalid protection group from all nodes on the local cluster.

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

Activating and Deactivating a Protection Group

When you activate a protection group, it assumes the role that you assigned to it during configuration.

For more information about configuring protection groups, see "How to Create and Configure a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 81.

▼ How to Activate a Sun StorEdge Availability Suite 3.2.1 Protection Group

Before You Begin

You can activate a protection group in the following ways:

- Globally, meaning you activate a protection group on both clusters where the protection group is configured
- On the primary cluster only
- On a secondary cluster only

When a protection group is activated on a primary or secondary cluster, the outcome depends on the type of data replication you are using. If you are using Sun StorEdge Availability Suite 3.2.1 software, data replication can only be started from the primary cluster. So, when you activate a protection group on the secondary cluster, the activation will not start data replication.

Steps

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" on page 43.

2. **Activate the protection group.**

   This command activates the protection group on the local cluster.

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

   ```
   # geopg start -e scope [-n] AVS-protection-group
   ```

   `-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 and the command performs the following operations on each device group in the protection group:

     - Verifies that the role configured for the replication resource is the same as the role of the protection group on the local cluster.
     - Verifies that the role of the volume sets associated with the device group is the same as the role of the protection group on the local cluster.
     - If the role of the protection group on the local cluster is `secondary`, unmounts the local volumes defined in all volume sets associated with the device group.

- If the role of the protection group on the local cluster is `primary`, enables the autosynchronization feature of the Sun StorEdge Availability Suite 3.2.1 remote mirror software. Also, resynchronizes the volume sets associated with the device group.

AVS-protection-group Specifies the name of the protection group

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

The `geopg start` command performs the following actions if the role of the protection group is `primary` on the local cluster:

- The command executes a script defined in the `RoleChange_ActionCmd`.
- The command brings the application resource groups in the protection group online on the local cluster.
- If the application resource group is a failover type resource group that shares affinities with one of the device groups in the same protection group, the command adds strong, positive affinities and failover delegation between the application resource group and the lightweight resource group.

  The application resource group must not have strong, positive affinities with failover delegation. Otherwise, the attempt to add strong, positive affinities with failover delegation with the lightweight resource group will fail.
- The command creates strong dependencies between the `HAStoragePlus` resource in the application resource group and the `HAStoragePlus` resource in the lightweight resource group for this device group.

If the command fails, the `Configuration` status might be set to `Error`, depending on the cause of the failure. The protection group remains deactivated, but data replication may be started and some resource groups may be brought online. You should run `geoadm status` to get the status of your system.

If the `Configuration` status is set to `Error`, revalidate the protection group by using the procedures described in "How to Validate a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 84.

Example 7–13 Activating a Sun StorEdge Availability Suite 3.2.1 Protection Group Globally

The following example illustrates how to activate a protection group globally:

```
# geopg start -e global avspg
```

Example 7–14 Activating a Sun StorEdge Availability Suite 3.2.1 Protection Group Locally

The following example illustrates how to activate a protection group on a local cluster only. This local cluster might be a primary cluster or a secondary cluster, depending on the cluster's role.

```
# geopg start -e local avspg
```

▼ How to Deactivate a Sun StorEdge Availability Suite 3.2.1 Protection Group

Before You Begin

You can deactivate a protection group in the following ways:

- Globally, meaning you deactivate a protection group on both the primary and the secondary cluster where the protection group is configured
- On the primary cluster only
- On the secondary cluster only

The result of deactivating a protection group on primary or secondary cluster depends on the type of data replication you are using. If you are using Sun StorEdge Availability Suite 3.2.1 software, data replication can only be stopped from the primary cluster. So, when you deactivate a protection group on the secondary cluster, this deactivate command will not stop data replication.

Steps

1. **Log in to one of the cluster nodes.**

   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" on page 43.

2. **Deactivate the protection group.**

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

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

   ```
   # geopg stop -e scope [-D] protection-group-name
   ```

   `-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 while leaving the protection group online

If you omit this option, the data replication subsystem and the protection group are both stopped. If the role of the protection group on the local cluster is `primary`, omitting the `-d` option also results in the following actions:

- Removal of resource group affinities and resource dependencies between the application resource groups in the protection group and the internal resource group
- Taking the application resource groups offline and putting them in an `unmanaged` state

protection-group-name — Specifies the name of the protection group

If the role of the protection group is `primary` on the local cluster, the `geopg stop` command disables the autosynchronization of each device group and places the volume sets into logging mode.

If the `geopg stop` command fails, execute the `geoadm status` command to see the status of each component. For example, the `Configuration` status might be set to `Error` depending upon the cause of the failure. The protection group might remain activated even though some resource groups might be taken offline. The protection group might be deactivated with data replication running.

If the `Configuration` status is set to `Error`, revalidate the protection group by using the procedures described in "How to Validate a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 84.

Example 7–15 Deactivating a Sun StorEdge Availability Suite 3.2.1 Protection Group on All Clusters

The following example illustrates how to deactivate a protection group on all clusters:

```
# geopg stop -e global avspg
```

Example 7–16 Deactivating a Sun StorEdge Availability Suite 3.2.1 Protection Group on a Local Cluster

The following example illustrates how to deactivate a protection group on the local cluster:

```
# geopg stop -e local avspg
```

Example 7–17 Stopping Sun StorEdge Availability Suite 3.2.1 Data Replication While Leaving the Protection Group Online

The following example illustrates how to stop only data replication on a local cluster:

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

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

```
# geopg stop -e local avspg
```

Example 7–18 Deactivating a Sun StorEdge Availability Suite 3.2.1 Protection Group While Keeping Application Resource Groups Online

The following example illustrates how to keep two application resource groups, `apprg1` and `apprg2`, online while deactivating their protection group, `avspg`.

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

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

2. Deactivate the protection group.

   ```
   # geopg stop -e global avspg
   ```

Resynchronizing a Sun StorEdge Availability Suite 3.2.1 Protection Group

You can resynchronize the configuration information of the local protection group with the configuration information 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" on page 42.

Resynchronizing a protection group updates only entities that are related to Sun Cluster Geographic Edition. 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 Sun StorEdge Availability Suite 3.2.1 Protection Group

Before You Begin The protection group must be deactivated on the cluster where you run the `geopg update` command.

Steps

1. **Log in to one of the cluster nodes.**

   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" on page 43.

2. **Resynchronize the protection group.**

   ```
   # geopg update protection-group-name
   ```

   protection-group-name Specifies the name of the protection group

Example 7–19 Resynchronizing a Sun StorEdge Availability Suite 3.2.1 Protection Group

The following example illustrates how to resynchronize a protection group:

```
# geopg update avspg
```

Checking the Runtime Status of Sun StorEdge Availability Suite 3.2.1 Data Replication

You can obtain an overall view of the status of replication, as well as a more detailed runtime status of the Sun StorEdge Availability Suite 3.2.1 software from the status of the replication resource groups. The following sections describe the procedures for checking each status.

Printing a Sun StorEdge Availability Suite 3.2.1 Runtime Status Overview

The status of each Sun StorEdge Availability Suite 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.

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

Steps

1. **Access a node of a cluster where the protection group is 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" on page 43.

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 printed 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 Sun StorEdge Availability Suite 3.2.1 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 7–1.

Printing a Detailed Sun StorEdge Availability Suite 3.2.1 Runtime Status

You must create one replication resource group for each protection group. The name of the replication resource group has the following format:

avs-protection-group-name`-rep-rg`

If you add a Sun StorEdge Availability Suite 3.2.1 device group to a protection group, the 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:

avs-devicegroup-name`-rep-rs`

You can monitor the state of the replication resource group to give you the overall status of replication. Use the `scstat -g` command to obtain the `State` and `Status Message` values for the replication resource group. The `State` is `Online` while the resource is online.

The following table describes the `Status` and `State Message` values that are returned by the `scstat` command when the `State` of the Sun StorEdge Availability Suite 3.2.1 replication resource group is `Online`.

TABLE 7–1 State and Status Messages of an Online Sun StorEdge Availability Suite 3.2.1 Replication Resource Group

Status	Status Message
Faulted	Replication service disabled
Faulted	Incorrect role
Faulted	Volume failed
Faulted	Bitmap failed
Faulted	Queue failed
Faulted	Need sync
Faulted	Need reverse sync
Faulted	Reverse synching
Degraded	Synching
Degraded	Queuing
Degraded	Logging
Online	Replicating

For more details about these values, refer to the *Sun StorEdge Availability Suite 3.2.1 Remote Mirror Software Administration and Operations Guide*.

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

CHAPTER 8

Migrating Services That Use Sun StorEdge Availability Suite 3.2.1 Data Replication

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

Detecting Cluster Failure on a System That Uses Sun StorEdge Availability Suite 3.2.1 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 occur when a protection group's overall state changes to 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 `OK` state during this default timeout interval, while the heartbeat mechanism continues to retry the primary cluster. Only the heartbeat plug-ins appear in the `Error` state.

  This query interval is set by using the `Query_interval` property of the heartbeat. If the heartbeat still fails after four times the `Query_interval` you configured (three retries and one emergency mode probing), a heartbeat-lost event is generated and logged in the system log. When using 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" on page 224.

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 these protection groups becomes `Unknown`.

Migrating Services That Use Sun StorEdge Availability Suite 3.2.1 With a Switchover

You perform a switchover of a Sun StorEdge Availability Suite 3.2.1 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 taken offline on the former primary cluster, `cluster-paris`.

  For a reminder of which cluster is `cluster-paris`, see Figure 2–1.
- 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`.

▼ How to Switch Over a Sun StorEdge Availability Suite 3.2.1 Protection Group From Primary to Secondary

Before You Begin

For a switchover to occur, data replication must be active between the primary cluster and the secondary cluster. Additionally, the data volumes on the two clusters must be in a synchronized state.

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

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

Steps

1. **Log in to one of the cluster nodes.**

   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" on page 43.

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 new-primary-cluster  protection-group-name
   ```

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

Example 8–1 Forcing a Switchover From Primary to Secondary

The following example illustrates how to perform a switchover to the secondary cluster:

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

More Information

Actions Performed by the Sun Cluster Geographic Edition Software During a Switchover

When the `geopg switchover` command is executed, the software confirms that the volume sets associated with the device groups are in the `replicating` state. Then, the software performs the following actions on the original primary cluster:

- Removes affinities and resource dependencies between all of the application resource groups in the protection group and the internal resource group, such as the lightweight resource groups
- Takes the application resource groups offline and puts them into the `unmanaged` state
- Waits for writes to complete
- Unmounts the primary volumes that correspond to the device groups in the protection group
- Stops data replication by placing all volume sets into logging mode
- Reverses the role of all volume sets

On the original secondary cluster, the command takes the following actions:

- Places all volume sets into logging mode
- Reverses the role of all volume sets
- Starts data replication by issuing update synchronization with the autosynchronization feature active
- Runs the script defined in the `RoleChange_ActionCmd` property
- Brings all application resource groups online and adds the affinities between the application resource groups and the internal resource groups, such as the lightweight resource group

If the command executes successfully, the secondary cluster, `cluster-newyork`, becomes the new primary cluster for the protection group. The original primary cluster, `cluster-paris`, becomes the new secondary cluster. Volume sets associated with a device group of the protection group have their role reversed according to the role of the protection group on the local cluster. The application resource group is online on the new primary cluster. Data replication from the new primary cluster to the new secondary cluster begins.

This command returns an error if any of the previous operations fails. Execute the `geoadm status` command to view the status of each component. For example, the `Configuration` status of the protection group might be set to `Error`, depending on the cause of the failure. The protection group might be activated or deactivated.

If the `Configuration` status of the protection group is set to `Error`, revalidate the protection group by using the procedures described in "How to Validate a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 84.

If the configuration of the protection group is not the same on each partner cluster, you need to resynchronize the configuration by using the procedures described in "How to Resynchronize a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 107.

Forcing a Takeover on Systems That Use Sun StorEdge Availability Suite 3.2.1

You 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 following steps occur after takeover is initiated:

- If the former primary cluster, `cluster-paris`, can be reached, the protection group is deactivated.

  For a reminder of which cluster is `cluster-paris`, see Figure 2–1.
- 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. Data replication from the new primary cluster, `cluster-newyork`, to the former primary cluster, `cluster-paris`, is stopped.

- Protection group is activated without data replication.

For details about the possible conditions of the primary and secondary cluster before and after takeover, see Appendix C.

The following procedures describe the steps you must perform to force takeover by a secondary cluster, and how to recover data afterward.

▼ How to Force Immediate Takeover of Sun StorEdge Availability Suite 3.2.1 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 up and 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.

Steps

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" on page 43.

2. **Initiate the takeover.**

   ```
   # geopg takeover  [-f] protection-group-name
   ```

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

Example 8–2 Forcing a Takeover by a Secondary Cluster

The following example illustrates how to force the takeover of `avspg` by the secondary cluster `cluster-newyork`.

`phys-newyork-1` 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" on page 33.

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

More Information

Actions Performed by the Sun Cluster Geographic Edition Software During a Takeover

When the `geopg takeover` command executes, the software confirms that the volume sets are in a `Replicating` or `Logging` state on the secondary cluster.

If the original primary cluster, `cluster-paris`, can be reached, the software performs the following actions:

- Removes affinities and resource dependencies between all of the application resource groups in the protection group and the internal resource group if the protection group was active
- Takes the application resource groups offline and put into an `unmanaged` state
- Unmounts the primary volumes that correspond to the device groups in the protection group
- Stops data replication by placing all volume sets into logging mode
- Reverses the role of all volume sets

On the original secondary cluster, `cluster-newyork`, the software performs the following actions:

- Places all volume sets into logging mode
- Reverses the role of all volume sets
- Runs the script specified in the `RoleChange_ActionCmd` property
- If the protection group was active on the original secondary cluster before the takeover, brings all application resource groups online and adds affinities and resource dependencies between the application resource group and the internal resource group

If the command executes successfully, the secondary cluster, `cluster-newyork`, becomes the new primary cluster for the protection group. Volume sets associated with a device group in the protection group have their role reversed according to the role of the protection group on the local cluster. If the protection group was active on the original secondary cluster before the takeover, the application resource groups are brought online on the new primary cluster. If the original primary cluster can be reached, it becomes the new secondary cluster of the protection group. Replication of all volume sets associated with the device groups of the protection group is stopped.

Caution – After a successful takeover, data replication is stopped. If you want to continue to suspend replication, specify the `-n` option any time you use the `geopg start` command. This option prevents the start of data replication from the new primary cluster to the new secondary cluster.

This command returns an error if any of the previous operations fails. Execute the `geoadm status` command to view the status of each component. For example, the `Configuration` status of the protection group might be set to `Error`, depending on the cause of the failure. The protection group might be activated or deactivated.

If the `Configuration` status of the protection group is set to `Error`, revalidate the protection group by using the procedures described in "How to Validate a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 84.

If the configuration of the protection group is not the same on each partner cluster, you need to resynchronize the configuration by using the procedures described in "How to Resynchronize a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 107.

Recovering Sun StorEdge Availability Suite 3.2.1 Data After a Takeover

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, 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 primary cluster's data was resynchronized with the data on the secondary cluster, `cluster-newyork`.

  For a reminder of which clusters are `cluster-paris` and `cluster-newyork`, see Figure 2–1.
- Failback-takeover. During a failback-takeover, applications are brought online again on the original primary cluster and use the current data on the primary cluster. Any updates that occurred on the secondary cluster are discarded.

▼ How to Perform a Failback-Switchover on a System That Uses Sun StorEdge Availability Suite 3.2.1 Replication

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

Before You Begin

Before you perform a failback-switchover, a takeover has occurred on `cluster-newyork`. The clusters now 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.

Steps 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`, deactivate the protection group on the local cluster.

```
# geopg stop -e Local protection-group-name
```

`-e Local`	Specifies the scope of the command By specifying a `local` scope, the command operates on the local cluster only.
protection-group-name	Specifies the name of the protection group

If the protection group is already deactivated, the state of the resource group in the protection group is probably `Error`. The state is `Error` because the application resource groups are managed and offline.

Deactivating the protection group will result in the application resource groups no longer being managed, clearing the `Error` state.

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

```
# geops update partnership-name
```

partnership-name	Specifies the name of the partnership

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

For more information about synchronizing partnerships, see "Resynchronizing a Partnership" on page 57.

c. 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 protection-group-name
```

protection-group-name	Specifies the name of the protection group

For more information about synchronizing protection groups, see "Resynchronizing a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 106.

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

   ```
   # geopg validate protection-group-name
   ```

protection-group-name	Specifies a unique name that identifies a single protection group

   For more information, see "How to Validate a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 84.

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

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

   ```
   # geopg start -e Global protection-group-name
   ```

`-e Global`	Specifies the scope of the command By specifying a `Global` scope, the command operates on both clusters where the protection group is deployed.
protection-group-name	Specifies the name of the protection group

   ---

   Note – The `-n` option must *not* be given when doing a failback-switchover 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 Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 101.

4. **Confirm that the data is completely synchronized.**

   First, confirm that the state of the protection group on `cluster-newyork` is `OK`.

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

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

   Next, confirm that all resources in the replication resource group, *AVS-protection-group-name*`-rep-rg`, report a status of `OK`.

   ```
   phys-newyork-1# scstat -g
   ```

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

   ```
   # geopg  switchover [-f] -m cluster-paris protection-group-name
   ```

   For more information, see "How to Switch Over a Sun StorEdge Availability Suite 3.2.1 Protection Group From Primary to Secondary" on page 113.

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

6. **Ensure that the switchover was performed successfully by using the `geoadm status` on either cluster to verify that the replication resource and the application resource groups and resources are online.**

   Also, you must verify that the protection group is now primary on `cluster-paris` and secondary on `cluster-newyork` and that "Data replication" and "Resource groups" are listed in `OK` states for both clusters.

   ```
   # geoadm status
   ```

▼ How to Perform a Failback-Takeover on a System That Uses Sun StorEdge Availability Suite 3.2.1 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.

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`.

Before You Begin

Before you begin the failover-takeover operation, 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.

Steps

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.**

      `# geops update` *partnership-name*

      partnership-name Specifies the name of the partnership

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

      For more information about synchronizing partnerships, see "Resynchronizing a Partnership" on page 57.

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

      If the protection group has been activated, deactivate the protection group by using the `geopg stop` command. For more information about deactivating a protection group, see "How to Deactivate a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 104.

      `# geopg update` *protection-group-name*

      protection-group-name Specifies the name of the protection group

      For more information about synchronizing protection groups, see "How to Resynchronize a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 107.

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

   `# geopg validate` *protection-group-name*

   protection-group-name Specifies a unique name that identifies a single protection group

   For more information, see "How to Validate a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 84.

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`.

   `# geopg start -e local -n` *protection-group-name*

   `-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.

---

protection-group-name — Specifies the name of the protection group

For more information, see "How to Activate a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 101.

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

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

```
# geopg takeover  [-f] protection-group-name
```

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

protection-group-name — Specifies the name of the protection group

For more information about the `geopg takeover` command, see "How to Force Immediate Takeover of Sun StorEdge Availability Suite 3.2.1 Services by a Secondary Cluster" on page 116.

The protection group on `cluster-paris` now has the `primary` role, and the protection group on `cluster-newyork` has the `secondary` role.

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

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

```
# geopg start -e local [-n] protection-group-name
```

`-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.

protection-group-name — Specifies the name of the protection group

For more information about the `geopg start` command, see "How to Activate a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 101.

6. **Start data replication.**

   To start data replication, activate the protection group on the primary cluster, `cluster-paris`.

   ```
   # geopg start -e local protection-group-name
   ```

   For more information about the `geopg start` command, see "How to Activate a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 101.

Recovering From a Sun StorEdge Availability Suite 3.2.1 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.

For example, suppose a device group controlled by Sun StorEdge Availability Suite 3.2.1 that is called `avsdg` changes to a `Volume failed` state, `VF`. This state is reflected in the following resource status:

```
Resource Status = "FAULTED"
Resource status message = "FAULTED : Volume failed"
```

Note – The `Resource State` remains `Online` because the probe is still running correctly.

Because the resource status has changed, the protection group status also changes. In this case, the local `Data Replication` state, the `Protection Group` state on the local cluster, and the overall `Protection Group` state become `Error`.

To recover from an error state, complete the relevant steps in the following procedure.

▼ How to Recover From a Data Replication Error

Steps

1. **Use the procedures in the Sun StorEdge Availability Suite 3.2.1 documentation to determine the causes of the `FAULTED` state. This state is indicated as `VF`.**

2. **Recover from the faulted state by using the Sun StorEdge Availability Suite 3.2.1 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 protection-group-name
   ```

   protection-group-name Specifies the name of the Sun StorEdge Availability Suite 3.2.1 protection group

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

   ```
   phys-paris-1# geopg list protection-group-name
   ```

   protection-group-name Specifies the name of the Sun StorEdge Availability Suite 3.2.1 protection group

CHAPTER 9

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 can replicate data with Hitachi TrueCopy software, you must be familiar with the Hitachi TrueCopy documentation and have the Hitachi TrueCopy product and 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 155. 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 175.

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 9–1 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 129.
Create a protection group that is configured for Hitachi TrueCopy data replication.	See "How to Create and Configure a Hitachi TrueCopy Protection Group" on page 144.
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 155.
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 152.
Replicate the protection group configuration to a secondary cluster.	See "How to Replicate the Hitachi TrueCopy Protection Group Configuration to a Partner Cluster" on page 163.
Test the configured partnership and protection groups to validate the setup.	Perform a trial a switchover or takeover and test some simple failure scenarios. See Chapter 11.
Activate the protection group.	See "How to Activate a Hitachi TrueCopy Protection Group" on page 167.
Check the runtime status of replication.	See "Checking the Runtime Status of Hitachi TrueCopy Data Replication" on page 175.
Detect failure.	See "Detecting Cluster Failure on a System That Uses Hitachi TrueCopy Data Replication" on page 179.
Migrate services by using a switchover.	See "Migrating Services That Use Hitachi TrueCopy Data Replication With a Switchover" on page 180.
Migrate services by using a takeover.	See "Forcing a Takeover on a System That Uses Hitachi TrueCopy Data Replication" on page 183.
Recover data after forcing a takeover.	See "Failback of Services to the Original Primary Cluster on a System That Uses Hitachi TrueCopy Replication" on page 187.
Detect and recover from a data replication error.	See "Recovering From a Hitachi TrueCopy Data Replication Error" on page 197.

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 current supported Sun Cluster configurations.

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

First, configure the Hitachi TrueCopy device groups on shared disks in the primary cluster. Disks `d1` and `d2` are configured to belong to a Hitachi TrueCopy device group that is called `devgroup1`. This configuration information is specified in the `/etc/horcm.conf` file on each of the cluster's nodes that has access to the Hitachi array. The application, `apprg1`, can run on these cluster nodes.

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 9–2 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.

Steps

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.**

   The output of this command should show `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

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.

Steps

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.

Steps

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 3.1 Data Service Planning and Administration Guide*.

Example 9–1 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).

```
# newfs dev/vx/dsk/oradg1/vol1
```

2. An entry in the `/etc/vfstab` file is created as follows:

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

3. Add the `HAStoragePlus` resource type.

```
# 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 just as you did on the primary cluster. Disks `d1` and `d2` are configured to belong to a Hitachi TrueCopy device group that is called `devgroup1`. This configuration information is specified in the `/etc/horcm.conf` file on each of the cluster's nodes that has access to the Hitachi array. The application, `apprg1`, can run on these cluster nodes.

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 9–3 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.

Steps

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`, using the VERITAS Volume Manager commands `vxdiskadm` and `vxassist` command.

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

Steps

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.**

   The output of this command should show `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 10

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 which of the following strategies is best for you:

- 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 to take for each strategy.

Creating a Protection Group While the Application is Offline

The steps to take to create a protection group while the application resource groups is offline follow:

- Create the protection group from a node on one cluster.

  For more information, see "How to Create and Configure a Hitachi TrueCopy Protection Group" on page 144.
- 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 155.
- 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 152.
- On the other cluster, retrieve the protection group configuration.

  For more information, see "How to Replicate the Hitachi TrueCopy Protection Group Configuration to a Partner Cluster" on page 163.
- From either cluster, start the protection group "globally".

  For more information, see "How to Activate a Hitachi TrueCopy Protection Group" on page 167.

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 node on one cluster.

  For more information, see "How to Create and Configure a Hitachi TrueCopy Protection Group" on page 144.

- 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 155.
- Start the protection group locally.

  For more information, see "How to Activate a Hitachi TrueCopy Protection Group" on page 167.
- 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 152.

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 Partner Cluster" on page 163.
- Activate the protection group locally.

  For more information, see "How to Activate a Hitachi TrueCopy Protection Group" on page 167.

EXAMPLE 10–1 Example of Creating a Hitachi TrueCopy Protection Group While the Application Remains Online

This example describes how to create 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 tc -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 an application resource group that is already online to the protection group.

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

   Verify that the application resource group was added successfully.

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

EXAMPLE 10–1 Example of Creating a Hitachi TrueCopy Protection Group While the Application Remains Online *(Continued)*

```
   Partnership "paris-newyork-ps"    : OK
      Partner clusters               : newyork
      Synchronization                : OK

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

5. On one node of the partner cluster, retrieve the protection group as follows:

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

6. 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.
```

7. Verify that the protection group was successfully created an 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

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

EXAMPLE 10–1 Example of 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" on page 144
- "How to Modify a Hitachi TrueCopy Protection Group" on page 146
- "How to Validate a Hitachi TrueCopy Protection Group" on page 147
- "How to Delete a Hitachi TrueCopy Protection Group" on page 148

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` *data-replication-type* 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" on page 150.

▼ How to Create and Configure a Hitachi TrueCopy Protection Group

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 Partner Cluster" on page 163.

Steps

1. **Log in to one of the cluster nodes.**

   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" on page 43.

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 partnership-name -o local-role -d truecopy [-p property-settings [-p...]] \
   protection-group-name
   ```

`-s` *partnership-name*	Specifies the name of the partnership
`-o` *local-role*	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 Hitachi TrueCopy
`-p` *property-setting*	Sets the properties of the protection group The properties you can set are the following: ■ `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.

protection-group-name 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.

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

Example 10–2 Creating and Configuring a Hitachi TrueCopy Protection Group

The following example illustrates how to create 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 10–3 Creating a Hitachi TrueCopy Protection Group for Application Resource Groups That Are Online

The following example illustrates how to create 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 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 of the nodes given 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 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 of 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-`*protection-group-name*.

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.

Steps

1. **Log in to one of the cluster nodes.**

   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" on page 43.

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-settings [-p...] \
   protection-group-name
   ```

   `-p` *property-setting* Sets the properties of the protection group.

For more information about the properties you can set, see Appendix A.

protection-group-name 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.

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

Example 10–4 Modifying the Configuration of a Protection Group

The following example illustrates how to modify the `Timeout` property of the protection group that was created in Example 10–2:

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

▼ How to Validate a Hitachi TrueCopy Protection Group

Before You Begin

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 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 issue the`geopg validate` command again.

Before validating the configuration of a protection group, 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.

Steps

1. **Log in to one of the cluster nodes.**

   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" on page 43.

2. **Validate the protection group's configuration.**

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

   ```
   # geopg validate protection-group-name
   ```

   protection-group-name Specifies a unique name that identifies a single protection group

Example 10–5 Validating the Configuration of a Protection Group

The following example illustrates how to validate a protection group.

```
# geopg validate tcpg
```

More Information Validations Made by the Data Replication Subsystem During Protection Group Validation

During protection group validation, the Hitachi TrueCopy data replication layer makes the following validations:

- The data replication layer confirms that the `horcmd` daemon is running on at least one of the nodes given 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*
- 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 by using the `scstat -D` command. The data replication layer also verifies that the device group is of a valid type.
- The data replication layer validates the properties of each Hitachi TrueCopy device group that has been added to the protection group.

▼ 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 10–7 and Example 10–10 for examples of this procedure.

Steps

1. **Log in to a nodes 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" on page 43.

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 protection-group-name
   ```

   protection-group-name Specifies the name of the protection group

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

Example 10–6 Deleting a Protection Group

The following example illustrates how to delete a protection group from both partner clusters.

The `cluster-paris` is the primary cluster. For a reminder of the sample cluster configuration, see Figure 2–1.

```
# rlogin cluster-paris -l root
cluster-paris# geopg delete tcpg
# rlogin cluster-newyork -l root
cluster-newyork# geopg delete tcpg
```

Example 10–7 Deleting a Hitachi TrueCopyProtection Group While Keeping Application Resource Groups Online

The following example illustrates how to keep two application resource groups, `apprg1` and `apprg2`, online 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
```

Creating a Protection Group That Does Not Require Data Replication

Some of the protection groups will not require data replication. If you are using the Sun Cluster Geographic Edition software to manage only resource groups and to handle data replication differently, you can create protection groups that do not replicate data. The `geoadm status` command shows a state for these protection groups of `Degraded`. This section describes how to configure your protection group to not use data replication.

For information about how to create a Hitachi TrueCopy protection group that uses data replication, see "How to Create and Configure a Hitachi TrueCopy Protection Group" on page 144.

Note – You cannot add device groups to a protection group that does not use data replication.

▼ How to Create a Protection Group That Does Not Require Data Replication

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 that 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.

Steps

1. **Log in to one of the cluster nodes.**

   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" on page 43.

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 partnership-name -o local-role \
   [-p property-settings [-p...]] \
   protection-group-name
   ```

`-s` *partnership-name*	Specifies the name of the partnership
`-o` *local-role*	Specifies the role of this protection group on the local cluster as either `Primary` or `Secondary`
`-p` *property-setting*	Sets the properties of the protection group The properties you can set are the following: ■ `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.
protection-group-name	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.

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

Example 10–8 Creating and Configuring a Protection Group That Is Not Replicated

The following example illustrates how to create an protection group that is not replicated:

```
# geopg create -s paris-newyork-ps -o primary example-pg
```

Next Steps See "Administering Hitachi TrueCopy Application Resource Groups" on page 152 for information on adding resource groups to a protection group.

Administering Hitachi TrueCopy Application Resource Groups

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

All of 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

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 to add already 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
  ```

  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 will fail.

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

If the resource group is `Offline` and the protection group is `Active` after the configuration of the protection group has changed, the protection group's local state 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 activate resource group.

Steps

1. **Log in to one of the cluster nodes.**

   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" on page 43.

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 resource-group-list protection-group
   ```

resource-group-list	Specifies the name of the application resource group You can specify more than one resource group in a comma-separated list.
protection-group	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.

   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 10–9 Adding an Application Resource Group to a Protection Group

The following example illustrates how to add 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

Before You Begin

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

Before you remove an application resource group from a protection group, 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 protection group's application resource groups. For example, you cannot remove a resource group that belongs to the data replication management entity.

Steps

1. **Log in to one of the cluster nodes.**

   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" on page 43.

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 resource-group-list protection-group
   ```

   resource-group-list — Specifies the name of the application resource group

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

   protection-group — 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 10–10 Deleting an Application Resource Group From a Protection Group

The following example illustrates how to remove 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:

- "How to Add a Data Replication Device Group to a Hitachi TrueCopy Protection Group" on page 155
- "Validations Made by the Data Replication Subsystem" on page 157
- "How the State of the Hitachi TrueCopy Device Group is Validated" on page 158
- "How to Modify a Hitachi TrueCopy Data Replication Device Group" on page 161
- "How to Delete a Data Replication Device Group From a Hitachi TrueCopy Protection Group" on page 162

For details about configuring a Hitachi TrueCopy data replication protection group, see "How to Create and Configure a Hitachi TrueCopy Protection Group" on page 144.

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

Steps

1. **Log in to one of the cluster nodes.**

   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" on page 43.

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-settings [-p...] device-group-name protection-group-name
   ```

   `-p` *property-settings* — Sets the properties of the data replication device group

   The Hitachi TrueCopy specific property that you can set is the following:

   - `Fence_level` – 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.

     This property can take the values of `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` to `data` or `status`, the application on the primary cluster might fail because the secondary cluster is not available for reasons such as:

     - 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 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 set depend upon the type of data replication you are using. For details about these properties, see Appendix A.

   device-group-name — Specifies the name of the new data replication device group

protection-group-name Specifies the name of the protection group which 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.

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

Example 10–11 Adding a Data Replication Device Group to a Hitachi TrueCopy Protection Group

The following example illustrates how to create 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 all of the nodes given 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 of the nodes given in the `Nodelist` property. The storage device controls the newHitachi TrueCopy device group.
- The Hitachi TrueCopy device group properties specified in the `geopg add-device-group` command are validated as described in the following table.

Hitachi TrueCopy Device Group Property	Validation
device-group-name	Checks that the specified Hitachi TrueCopy device group is configured on all of the cluster nodes that are specified in the `Nodelist` property.

Hitachi TrueCopy Device Group Property	Validation
`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-`*protection-group-name*`-`*device-group-name*. This resource is placed in the corresponding Sun Cluster resource group, which is named `rg-tc-`*protection-group-name*.

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 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 10–1 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`
	A state of `SSWS` can be seen when you use the `pairdisplay -g <DG> -fc` command.

`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`
	A state of `SSWS` can be seen 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 10–2 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`

For any other combination of individual device group states, the aggregate device group state cannot be obtainable and 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 10–3 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 10–12 Validating the Aggregate Device Group State

The following example illustrates how the state of a Hitachi TrueCopy device group is validated 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
```

EXAMPLE 10–12 Validating the Aggregate Device Group State *(Continued)*

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 given 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
```

The output of the `pairvolchk` command is `23`, which corresponds in Table 10–1 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 10–3.

▼ How to Modify a Hitachi TrueCopy Data Replication Device Group

Steps

1. **Log in to one of the cluster nodes.**

   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" on page 43.

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-settings [-p...] \
   TC-device-group-name  protection-group-name
   ```

   `-p` *property-settings* — Sets the properties of the data replication device group

   For more information about the properties you can set, see Appendix A.

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

Example 10–13 Modifying the Properties of a Hitachi TrueCopy Data Replication Device Group

The following example illustrates how to modify 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 148. 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 154.

Steps

1. **Log in to one of the cluster nodes.**

   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" on page 43.

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 device-group-name protection-group-name
   ```

device-group-name	Specifies the name of the data replication device group
protection-group-name	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-`*protection-group-name*`-`*device-group-name*, 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 10–14 Deleting a Replication Device Group From a Hitachi TrueCopy Protection Group

The following example illustrates how to remove a Hitachi TrueCopy data replication device group:

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

Replicating the Hitachi TrueCopy Protection Group Configuration to a Partner 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 Partner Cluster

Before You Begin Before you replicate the configuration of a Hitachi TrueCopy protection group to a partner 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
  ```

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

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

Steps

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" on page 43.

   `phys-newyork-1` is the only node on the secondary cluster. For a reminder of which node is `phys-newyork-1` , see Figure 2–1.

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 partnership-name [protection-group]
   ```

`-s` *partnership-name*	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.
protection-group	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 10–15 Replicating the Hitachi TrueCopy Protection Group to a Partner Cluster

The following example illustrates how to replicate 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 Hitachi TrueCopy Protection Group

When you activate a protection group, the protection group assumes the role that you assigned to it during configuration. You can activate a protection group in the following ways:

- Globally, meaning you activate a protection group on both clusters where the protection group is configured.
- On the primary cluster only, so that the secondary cluster remains inactive.
- On a secondary cluster only, when it remains inactive on the primary cluster.

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

- The data replication configuration of the protection group is validated. During validation, the protection group's current local role is compared with the aggregate device group state as described in Table 10–3. 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 protection group's local role is `primary` to the cluster on which the protection group's local role is `secondary`.

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

Activating a protection group has the following impact 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 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 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 configured for the device group.

TABLE 10–4 Commands Used to Start Hitachi TrueCopy Data Replication

Aggregate Device Group State	Valid Local Protection Group Role	Hitachi TrueCopyStart 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 up on the remote cluster.

`Regular Primary`	`primary`	If the local state code is 22, 23, 25, 26, 29, 42, 43, 45, 46, or 47, no command is issued because data is already being replicated. If the local state code is 24, 44, or 48, then the following command is issued: `pairresync -g dg [-l]` If the local state code is 11, then the following command is issued: `paircreate -vl -g dg -f fl` Both commands require that the `horcmd` process is up 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 issued because data is already being replicated. If the local state code is 34, 54, or 58, then the following command is issued: `pairresync -g dg` If the local state code is 11, the following command is issued: `paircreate -vr -g dg -f fl` Both commands require that the `horcmd` process is up on the remote cluster.

TABLE 10–4 Commands Used to Start Hitachi TrueCopy Data Replication *(Continued)*

Aggregate Device Group State	Valid Local Protection Group Role	Hitachi TrueCopyStart Command
`Takeover Primary`	`primary`	If the local state code is 34 or 54, the following command is issued: `pairresync -swaps -g` If the local state code is 11, then the following command is issued: `paircreate -vl -g dg -f fl` The `paircreate` command requires that the `horcmd` process is up on the remote cluster.
`Takeover Secondary`	`secondary`	If the local state code is 24, 44, 25, or 45, the following command is issued: `pairresync -swapp -g dg` If the local state code is 11, the following command is issued: `paircreate -vr -g dg -f fl` Both commands require that the `horcmd` process is up on the remote cluster.

▼ How to Activate a Hitachi TrueCopy Protection Group

Steps

1. **Log in to one of the cluster nodes.**

   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" on page 43.

2. **Activate the protection group.**

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

   ```
   # geopg start -e scope [-n] protection-group-name
   ```

   `-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.

protection-group-name — Specifies the name of the protection group

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

Example 10–16 How the Sun Cluster Geographic Edition Software Issues the Command to Start Replication

The following example illustrates how the Sun Cluster Geographic Edition determines the Hitachi TrueCopy command 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 TrueCopydevice group, `devgroup1`, is given 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..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 executes the `paircreate -g devgroup1 -vl -f async` command at the data replication level. If the command is successful, the state of `devgroup1` is given 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 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 10–17 Activating a Hitachi TrueCopy Protection Group Globally

The following example illustrates how to activate 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 10–18 Activating a Hitachi TrueCopy Protection Group Locally

The following example illustrates how to activate a protection group on a local cluster only. This local cluster might be a primary cluster or a secondary cluster, depending on the cluster's role.

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

Deactivating a Hitachi TrueCopy Protection Group

You can deactivate a protection group in the following ways:

- Globally, meaning you deactivate a protection group on both the primary and the secondary cluster where the protection group is configured
- On the secondary cluster only, so that the primary cluster remains active
- On a primary cluster, after the protection group was previously deactivated on the secondary cluster

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

- The data replication configuration of the protection group is validated. During validation, the protection group's current local role is compared with the aggregate device group state as described in Table 10–3. 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, no matter whether the deactivation occurs on a primary or secondary cluster.

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

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

The Hitachi TrueCopy command 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 10–5 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 issued because no data is being replicated.
`Regular Primary`	`primary`	If the local state code is 22, 23, 26, 29, 42, 43, 46, or 47, then the following command is issued: `pairsplit -g dg [-l]` If the local state code is 11, 24, 25, 44, 45, or 48, then no command is issue because no data is being replicated.

TABLE 10–5 Commands Used to Stop Hitachi TrueCopy Data Replication *(Continued)*

Aggregate Device Group State	Valid Local Protection Group Role	Hitachi TrueCopyStop Command
`Regular Secondary`	`secondary`	If the local state code is 32, 33, 35, 36, 39, 52, 53, 55, 56, or 57, the following command is issued: `pairsplit -g dg` If the local state code is 33 or 53 and the remote state is `PSUE`, no command is issued to stop replication. If the local state code is 11, 34, 54, or 58, then no command is issue because no data is being replicated.
`Takeover Primary`	`primary`	No command is issued because no data is being replicated.
`Takeover Secondary`	`secondary`	No command is issued because no data is being replicated.

▼ How to Deactivate a Hitachi TrueCopy Protection Group

Steps

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" on page 43.

2. **Deactivate the protection group.**

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

   ```
   # geopg stop -e scope [-D] protection-group-name
   ```

   `-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 while leaving the protection group online. If you omit this option, the data replication subsystem and the protection group are both stopped.
protection-group-name	Specifies the name of the protection group.

Example 10–19 How the Sun Cluster Geographic Edition Software Issues the Command to Stop Replication

The following example illustrates how the Sun Cluster Geographic Edition software determines the Hitachi TrueCopy command used to stop data replication.

The current state of the Hitachi TrueCopy device group, `devgroup1`, is given 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    -
```

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 executes 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 executes the `pairsplit -g devgroup1` command at the data replication level.

If the command is successful, the state of `devgroup1` is given 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 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 10–20 Deactivating a Protection Group on All Clusters

The following example illustrates how to deactivate a protection group on all clusters:

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

Example 10–21 Deactivating a Protection Group on a Local Cluster

The following example illustrates how to deactivate a protection group on the local cluster:

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

Example 10–22 Stopping Data Replication While Leaving the Protection Group Online

The following example illustrates how to stop 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 reissue the command without the `-D` option:

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

Example 10–23 Deactivating a Hitachi TrueCopy Protection Group While Keeping Application Resource Groups Online

The following example illustrates how to keep 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 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" on page 42.

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 on deactivating a protection group, see "Deactivating a Hitachi TrueCopy Protection Group" on page 169.

Steps

1. **Log in to one of the cluster nodes.**

   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" on page 43.

2. **Resynchronize the protection group.**

   ```
   # geopg update protection-group-name
   ```

   protection-group-name Specifies the name of the protection group

Example 10–24 Resynchronizing a Protection Group

The following example illustrates how to resynchronize 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.

Printing 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" on page 217.

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

Steps

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" on page 43.

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 printed 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 given for the data replication device group you want to check.

See Also For more information about these fields, see Table 10–6.

Printing 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_`*truecopy-protection-group-name*

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-`*truecopy-protection-group-name*`-`*truecopy-devicegroup-name*

You can monitor the status of replication of this device group by looking at the `Status` and `Status Message` of this resource. Resource status and the status message are shown by the `scstat -g` command.

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 10–6 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

TABLE 10–6 State and Status Messages of an Online Hitachi TrueCopy Replication Resource Group *(Continued)*

Status	Status Message
Degraded	P-Vol/S-Vol:COPY
Degraded	P-Vol/S-Vol:COPY:Remote horcmd not reachable
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 details about these values, refer to the Hitachi TrueCopy documentation.

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

CHAPTER 11

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 group 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" on page 224.

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 Figure 2–1.
- 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 cluster.

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 command given on the new primary cluster, `cluster-newyork`, are described in the following table.

TABLE 11–1 Hitachi TrueCopy Switchover Validations on the New Primary Cluster

Aggregate Device Group State	Valid Device Group State on Local Cluster	Hitachi TrueCopy Switchover Command Issued on `cluster-newyork`
SMPL	None	None
Regular primary	23, 43	No command is issued, 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 given when the `fence_level` of the Hitachi TrueCopy device group is `async`. The value is calculated as 80% of the protection group's `Timeout` property. For example, if the protection group has a `Timeout` of 200 seconds, the value of `-t` used in this command will be 80% off 200 seconds, or 160 seconds.

TABLE 11–1 Hitachi TrueCopy Switchover Validations on the New Primary Cluster *(Continued)*

Aggregate Device Group State	Valid Device Group State on Local Cluster	Hitachi TrueCopy Switchover Command Issued on `cluster-newyork`
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 pre-switchover `PVOL_PAIR` volumes become the `SVOL_PAIR` volumes. The pre-switchover `SVOL_PAIR` volumes 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 protection group's `Local-role` becomes `Primary`. On the cluster on which the protection group had a `Local-role` of `Primary`, the protection group's `Local-role` 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 up and 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.

Steps

1. **Log in to one of the cluster nodes.**

   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" on page 43.

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 new-primary-cluster protection-group-name
   ```

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

Example 11–1 Forcing a Switchover From Primary to Secondary

The following example illustrates how to perform 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

You 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 following steps occur after takeover is initiated:

- If the former primary cluster, `cluster-paris`, can be reached, the application services are taken offline on the former primary cluster.

  For a reminder of which cluster is `cluster-paris`, see Figure 2–1.

- 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.

This following sections describe the steps you must perform to force 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 step 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 will fail. The Hitachi TrueCopy commands used for the takeover are described in the following table.

TABLE 11–2 Hitachi TrueCopy Takeover Validations on the New Primary Cluster

Aggregate Device Group State	Valid Local State Device Group State	Hitachi TrueCopy Takeover Command Issued on `cluster-newyork`
SMPL	All	No command is issued.
Regular primary	All	No command is issued.

TABLE 11–2 Hitachi TrueCopy Takeover Validations on the New Primary Cluster *(Continued)*

Aggregate Device Group State	Valid Local State Device Group State	Hitachi TrueCopy Takeover Command Issued on `cluster-newyork`
Regular secondary	All Regular secondary states except 32 or 52 For a list of Regular secondary states, refer to Table 10–1 and Table 10–2.	`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 protection group's `Timeout` property. For example, if the protection group has a `Timeout` of 200 seconds, the value of `-t` used in this command will be 80% off 200 seconds, or 160 seconds.
Takeover primary	All	No command is issued.
Takeover secondary	All	`pairsplit -R-g dgpairsplit -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, regardless of 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 protection group's `Local-role` becomes `Primary`. On `cluster-paris`, where the protection group had a `Local-role` of `Primary`, the following may occur:

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

If the takeover is successful, the applications are brought online. You do not need to issue 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 issue 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 up and 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.

Steps

1. **Log in 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" on page 43.

2. **Initiate the takeover.**

   ```
   # geopg takeover [-f] protection-group-name
   ```

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

Example 11–2 Forcing a Takeover by a Secondary Cluster

The following example illustrates forcing the takeover of `tcpg` by the secondary cluster `cluster-newyork`.

`phys-newyork-1` 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" on page 33.

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

Failback of Services to the Original Primary 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 original primary cluster's data was resynchronized with the data on the secondary cluster, `cluster-newyork`.

  For a reminder of which clusters are `cluster-paris` and `cluster-newyork`, see Figure 2–1.
- 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 are primary, are discarded.

▼ 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 this cluster's data has been resynchronized with the data on the current primary cluster, `cluster-newyork`.

Before You Begin

Before you perform a failback-switchover, 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" on page 42.
- 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.

Steps

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` *partnership-name*

      partnership-name Specifies the name of the partnership

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

      For more information about synchronizing partnerships, see "Resynchronizing a Partnership" on page 57.

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

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

      `# geopg update` *protection-group-name*

      protection-group-name Specifies the name of the protection group

      For more information about synchronizing protection groups, see "Resynchronizing a Hitachi TrueCopy Protection Group" on page 174.

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

   `# geopg validate` *protection-group-name*

   protection-group-name Specifies a unique name that identifies a single protection group

   For more information, see "How to Validate a Hitachi TrueCopy Protection Group" on page 147.

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 protection-group-name
   ```

`-e local`	Specifies the scope of the command By specifying a `local` scope, the command operates on the local cluster only.
protection-group-name	Specifies the name of the protection group

   Note – The `-n` option must *not* be given when doing a failback-switchover 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 167.

4. **Before performing the switchover, wait for the data to be completely synchronized.**

   The data is completely synchronized when the state of the protection group on `cluster-newyork` is `OK`. 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`.

   To confirm that the state of the protection group on `cluster-newyork` is `OK` use the following command:

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

   Refer to the `Protection Group` section of the output

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

   ```
   # geopg switchover [-f] -m cluster-paris protection-group-name
   ```

   For more information, see "How to Switch Over a Hitachi TrueCopy Protection Group From Primary to Secondary" on page 182.

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

▼ 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.

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 have used the `-n` option any time you used the `geopg start` command.

Before You Begin

Before you begin the failover-takeover operation, 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.

Steps

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.**

      ```
      # geops update partnership-name
      ```

      partnership-name Specifies the name of the partnership

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

      For more information about synchronizing partnerships, see "Resynchronizing a Partnership" on page 57.

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

Use the `pairsplit` commands to put 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 given 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..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.

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

```
# geopg update protection-group-name
```

protection-group-name Specifies the name of the protection group

For more information about resynchronizing protection groups, see "How to Resynchronize a Protection Group" on page 174.

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

```
# geopg validate protection-group-name
```

protection-group-name Specifies a unique name that identifies a single protection group

For more information, see "How to Validate a Hitachi TrueCopy Protection Group" on page 147.

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`.

   `# geopg start -e local -n` *protection-group-name*

   `-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.

   ---

   protection-group-name Specifies the name of the protection group

   For more information, see "How to Activate a Hitachi TrueCopy Protection Group" on page 167.

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

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

   `# geopg takeover [-f]` *protection-group-name*

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

   protection-group-name 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 186.

   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`.

   `# geopg start -e local [-n]` *protection-group-name*

`-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.
protection-group-name	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 167.

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

When the `geopg switchover` command is executed, a `horctakeover` command is executed 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 may 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 the case of failure, the `horctakeover` command usually returns a value of `5`, which indicates a SVOL-SSUS-takeover.

One reason the `horctakeover` command might fail to do 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 bring up the applications on the new intended primary cluster in a switchover failure scenario.

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

Switchover Failure Conditions

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

A switchover is executed to switch 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 a 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 will bring the application online on the appropriate cluster.

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

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

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

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

   Make the original primary cluster, `cluster-paris`, `Primary` for the protection group if you intend to bring up the application on the original primary cluster. The application will use the current data on the original primary cluster.

   Make the original secondary cluster, `cluster-newyork`, `Primary` for the protection group if you intend to bring up the application on the original secondary cluster. The application will use 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` may not be synchronized with the data volumes on `cluster-paris`. If you intent to bring up the application with the same data as 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

Steps

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 has successfully executed, `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 brings up 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

Steps

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 has successfully executed, `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 brings up the application on `cluster-newyork`. Data replication starts from `cluster-newyork` to `cluster-paris`.

Caution – This command will overwrite 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

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.

For information about how different `Resource status` values map to actual replication pair states, see Table 10–6.

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 given 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 10–6 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 display 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

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

Steps

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 protection-group-name
   ```

protection-group-name Specifies the name of the Hitachi TrueCopy protection group

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

```
phys-paris-1# geopg list protection-group-name
```

protection-group-name Specifies the name of the Hitachi TrueCopy protection group

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

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

CHAPTER 12

Administering Heartbeats

Sun Cluster Geographic Edition software uses heartbeats over the public network as a way for the individual clusters participating in partnerships to detect cluster failures at partner sites. The heartbeat monitor uses plug-in modules to query the heartbeat status of its partners.

This chapter discusses the following topics:

Introduction to Heartbeats

A heartbeat in Sun Cluster Geographic Edition is a container for a collection of heartbeat plug-ins. A heartbeat has a name and one property that you can tune, `Query_interval`. The `Query_interval` property specifies the delay between heartbeat status requests.

The heartbeat plug-in facilitates the actual physical monitoring activity. The plug-in is defined by a required query command or query library, an optional requester and responder agent, a type, and a `Plugin_properties` string.

The Sun Cluster Geographic Edition product provides the following default plug-ins:

- `tcp_udp_plugin` — Performs a simple heartbeat check on the cluster logical host IP address. If `tcp_udp_plugin` cannot use UDP port 2084, the plug-in tries to use TCP port 2084.
- `icrm_plugin` — Calls the cluster management agent that is running on the remote cluster by using the logical hostname.
- `ping_plugin` — Pings the cluster logical hostname on the remote cluster.

A default heartbeat that uses the default heartbeat plug-ins is created every time you run `geops create` or `geops join` without specifying a custom heartbeat. The name of the default heartbeat is `hb_`*local-cluster-name~remote-cluster-name*. For more information about the `geops` command, refer to the `geops`(1M) man page.

You can create custom heartbeat plug-ins and associate them with existing default heartbeats or with new custom heartbeats.

Note – Custom heartbeats are provided for special circumstances and require careful configuration. Consult your Sun specialist for assistance if your system requires the use of custom heartbeats.

Creating a Heartbeat

This section describes procedures for creating heartbeats.

▼ How to Create a Heartbeat

Use this procedure to create a new heartbeat. If you're planning on using the heartbeat with a partnership, you must create the heartbeat before you create a partnership. If you create a partnership before you create the custom heartbeat, the default heartbeat that is used by the partnership will prevent the custom heartbeat from being created.

A custom heartbeat prevents the default heartbeat from being used during partnership creation. If you want to use the default heartbeat for your partnership, you must delete the custom heartbeat before running the `geops create` command.

Steps **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" on page 43.

2. **Create the heartbeat.**

```
# geohb create -r remote-clustername \
[-p property-setting [-p...]] heartbeat-name
```

`-r` *remote-clustername*	Specifies the name of the remote, secondary partner cluster
`-p` *property-setting*	Specifies a heartbeat property which is assigned a value by using a *name=statement* pair. Multiple properties might be set at one time by using multiple statements. For more information about the properties you can set, see Appendix A.
heartbeat-name	Specifies an identifier for the heartbeat.

Caution – The name of the custom heartbeat on each cluster in the same partnership must be different. Choose a name that identifies the heartbeat uniquely, such as `paris-to-newyork` on the cluster `cluster-paris` and `newyork-to-paris` on cluster `cluster-newyork`.

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

Example 12–1 Creating a Heartbeat

This example creates a heartbeat that is named `paris-to-newyork`.

```
# geohb create -r cluster-newyork paris-to-newyork
```

Creating a Heartbeat Plug-in

This section describes procedures for creating a heartbeat plug-in.

▼ How to Create Heartbeat Plug-in

Steps

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" on page 43.

2. **Add the heartbeat plug-in to an existing heartbeat.**

   ```
   # geohb add-plugin heartbeat-name plug-in-name \
   [-p property-setting [-p...]]
   ```

heartbeat-name	Specifies the identifier for an heartbeat on the local cluster.
plugin-name	Specifies the name of the heartbeat plug-in.
`-p`*property-setting*	Specifies a heartbeat plug-in property which is assigned a value by using a *name=statement* pair. Multiple properties might be set at one time by using multiple statements. For more information about the properties you can set, see Appendix A.

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

Example 12–2 Creating a Heartbeat Plug-in

This example creates a heartbeat plug-in that is named `command1`

```
# geohb add-plugin paris-to-newyork command1 -p Query_cmd=/usr/bin/hb/
```

Modifying a Heartbeat Plug-in Property

This section describes procedures for modifying heartbeat plug-in properties. When you modify a plug-in property, your changes take effect immediately.

▼ How to Modify the Properties of a Heartbeat Plug-in

Steps

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" on page 43.

2. **Modify the heartbeat plug-in properties.**

   ```
   # geohb modify-plugin -p property-setting \
   [-p...] plugin-name heartbeat-name
   ```

heartbeat-name	Specifies an identifier for the heartbeat.
plugin-name	Specifies the name of the heartbeat plug-in.
`-p` *property-setting*	Specifies a heartbeat plug-in property which is assigned a value by using a *name=statement* pair. Multiple properties might be set at one time by using multiple statements. For more information about the properties you can set, see Appendix A.

   Note – You cannot edit some properties of the default plug-ins.

   For information about the names and values that are supported by Sun Cluster Geographic Edition software, see Appendix B.

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

Example 12–3 Modifying the Properties of the Heartbeat Plug-in

This example modifies the settings of the default TCP/UDP plug-in, `tcp_udp_plugin`, to use only TCP.

```
# geohb modify-plugin -p Plugin_properties=paris-cluster/TCP/2084 \
tcp_udp_plugin hb_cluster-paris~cluster-newyork
```

Deleting Heartbeats and Heartbeat Plug-ins

This section describes procedures for deleting heartbeats and heartbeat plug-ins.

▼ How to Delete a Heartbeat

Steps

1. **Log in 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" on page 43.

2. **Delete the heartbeat.**

   ```
   # geohb delete heartbeat-name
   ```

   heartbeat-name Specifies an identifier for the heartbeat settings

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

Caution – You must not delete the default heartbeats `tcp_upd_plugin`, `icrm_plugin`, `ping_plugin`.

Example 12–4 Deleting a Heartbeat

This example deletes a heartbeat that is named `paris-to-newyork`.

```
# geohb delete paris-to-newyork
```

▼ How to Delete a Plug-in From a Heartbeat

Steps

1. **Log in 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" on page 43.

2. **Remove the plug-in from the heartbeat.**

```
# geohb remove-plugin plugin-name heartbeat-name
```

plugin-name Specifies the name of the custom heartbeat plug-in.

heartbeat-name Specifies an identifier for the heartbeat that contains this plug-in.

Note – You cannot delete the default plug-in `icrm_plugin`.

For information about the names and values that are supported by Sun Cluster Geographic Edition software, see Appendix B.

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

Example 12–5 Deleting a Plug-in From a Heartbeat

This example removes the plug-in that is named `command1`, from the heartbeat that is named `paris-to-newyork`.

```
# geohb remove-plugin command1 paris-to-newyork
```

Printing Heartbeat Configuration Information

This section describes procedures for printing heartbeat configuration information.

▼ How to Print Heartbeat Configuration Information

Steps

1. **Log in 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" on page 43.

2. **Display the current configuration information for a specific heartbeat or the whole heartbeat subsystem.**

```
# geohb list [heartbeat-name-list]
```

heartbeat-name-list Specifies the names of the specifies heartbeats on the local cluster for which configuration information should be printed

If you do not specify a list of heartbeat names, this command displays information about all the configured heartbeats.

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

Example 12–6 Displaying Heartbeat Configuration Information

This example displays information about the `paris-to-newyork` heartbeat.

```
# geohb list paris-to-newyork
```

Tuning the Heartbeat Properties

Default heartbeats are created as part of partnership creation. If you use a custom heartbeat, the custom heartbeat should be created before you create a partnership. You can modify the properties of the default and custom heartbeats by using the `geohb set-prop` command. For more information about this command, refer to the `geohb`(1M) man page.

Note – Custom heartbeats are provided for special circumstances and require careful configuration. Consult your Sun specialist for assistance if your system requires the use of custom heartbeats.

If you modify the default value of the `Query_interval` property, ensure that the interval is sufficiently long. An interval that is too short causes a timeout and heartbeat-loss event before the logical hostname resource is available. This failover should result in no more than two unanswered heartbeat requests. Setting a default `query_interval` value of 120 seconds with the default `heartbeat.retries` parameter of 3 enables the peer cluster to be unresponsive for 6 minutes (120 * 3) without having a false failure declared.

The `heartbeat.retries` parameter is specified in the `com.sun.cluster.agent.geocontol.xml` file.

If you adjust the delay setting of the `Query_interval` property, the following should be true:

```
Query_interval > worst-case logical-host failover time / 2
```

You must empirically determine the logical-host failover time for the cluster in question.

The following must be true to avoid false failures:

```
Query_interval > worst-case logical-host failover time / 3
```

You should not need to change the `heartbeat.retries` value. If you want to change the default value of the `heartbeat.retries` property, contact a Sun specialist.

▼ How to Modify the Heartbeat Properties

Steps

1. **Log in to one of the cluster nodes.**

   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" on page 43.

2. **Modify the heartbeat properties.**

   ```
   # geohb set-prop -p property-setting \
   [-p...] heartbeat-name
   ```

   `-p` *property-setting* — Sets the default properties of the heartbeat

   A heartbeat property is assigned a value by a *name=statement* pair. Multiple properties can be set at one time by using multiple statements.

   For more information about the properties you can set, see Appendix A.

   heartbeat-name — Specifies an identifier for the heartbeat settings

   For information about the names and values that are supported by Sun Cluster Geographic Edition software, see Appendix B.

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

Example 12–7 Modifying the Properties of the Default Heartbeat

The following example illustrates how to modify the settings for the default heartbeat between `cluster-paris` and `cluster-newyork`:

```
# geohb set-prop -p Query_interval=60 hb_cluster-paris~cluster-newyork
```

Creating a Heartbeat That Uses a Custom Heartbeat Plug-in

You can create a custom heartbeat plug-in and configure an existing default heartbeat or a new custom heartbeat to use this custom heartbeat plug-in.

Custom heartbeats are provided for special circumstances and require careful configuration. Consult your Sun specialist for assistance if your system requires the use of custom heartbeats.

Note – If you configure a custom heartbeat, ensure that the name of your custom heartbeat is different from the name of the custom heartbeat on the partner cluster.

Caution – The presence of a custom heartbeat prevents the default heartbeat from being used during partnership creation. If you want to use the default heartbeat for your partnership, you must delete the custom heartbeat before running the `geops create` command.

Creating a Custom Heartbeat Plug-in

When a heartbeat is created, your custom heartbeat plug-in is passed the following arguments by the Sun Cluster Geographic Edition software:

query-interval	The value of the `Query-interval` property, which defines the delay in seconds after which a heartbeat status request is declared a failure
mode	The mode for the plug-in startup, either `Normal` or `Emergency`
plugin-property-values	The value of the `Plugin-properties` property that is configured for the heartbeat plug-in, if any For more information about the properties you can set, see Appendix A.

Your custom heartbeat plug-in is expected to check the heartbeat on the secondary cluster and return one of the following exit values:

- Zero if successful, meaning that the secondary cluster is alive

- Nonzero on failure, meaning that the secondary cluster did not respond to the heartbeat check

▼ How to Add a Custom Heartbeat Plug-in to an Existing Default Heartbeat

Steps

1. **Log in a node in 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" on page 43.

2. **Add the custom heartbeat plug-in to the default heartbeat.**

   ```
   # geohb add-plugin -p property-setting [-p...] \
   plugin-name hb_local-cluster-name-remote-cluster-name
   ```

`-p` *property-setting*	Sets the properties of the heartbeat plug-in by using a *name=statement* pair Specify the path to your custom heartbeat plug-in by using the `Query_cmd` property. For more information about the properties you can set, see Appendix A.
plugin-name	Specifies the name of the custom heartbeat plug-in
hb_*local-cluster-name-remote-cluster-name*	Specifies the name of the default heartbeat to which you want to add the custom heartbeat plug-in

3. **Verify that your changes were made correctly.**

   ```
   # geoadm status
   ```

4. **Repeat the previous steps on one node of the secondary cluster.**

Example 12–8 Adding a Custom Heartbeat Plug-in to the Default Heartbeat

The following example illustrates how to add the custom heartbeat plug-in, `command1`, to the default heartbeat, `hb_cluster-paris~cluster-newyork`:

```
# geohb add-plugin -p query_cmd=/usr/bin/hb command1 \
hb_cluster-paris~cluster-newyork
# geoadm status
```

▼ How to Create a Custom Heartbeat Plug-in and Add It to a Custom Heartbeat

Steps

1. **Log 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" on page 43.

2. **Create the new custom heartbeat.**

   ```
   # geohb create -r remote-clustername \
   [-p property-setting [-p...]] heartbeat-name
   ```

`-r` *remote-clustername*	Specifies the name of the remote, secondary partner cluster
`-p` *property-setting*	Sets the default properties of the heartbeat A heartbeat property is assigned a value by a *name=statement* pair. For more information about the properties you can set, see Appendix A.
heartbeat-name	Specifies an identifier for the heartbeat settings

   

   Caution – The name of the custom heartbeat on each cluster in the same partnership must be different. Choose a name that identifies the heartbeat uniquely, such as `paris-to-newyork` on the cluster `cluster-paris` and `newyork-to-paris` on cluster `cluster-newyork`.

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

3. **Add the custom heartbeat plug-in to the heartbeat.**

   ```
   # geohb add-plugin -p property-setting [-p...] \
   plugin-name heartbeat-name
   ```

`-p` *property-setting*	Sets the properties of the heartbeat plug-in by using a *name=statement* pair Specify the path to your custom heartbeat plug-in by using the `Query_cmd` property.

For more information about the properties you can set, see Appendix A.

plugin-name — Specifies the name of the custom heartbeat plug-in

heartbeat-name — Specifies an identifier for the heartbeat settings

4. **Create the partnership that will use the heartbeat that you created in the previous step.**

```
# geops create -c remote-cluster-name -h heartbeat-name \
[-p property-setting [-p...]] partnership-name
```

`-c` *remote-cluster-name* — Specifies the name of remote cluster that will participate in the partnership

This name matches the logical hostname used by the Sun Cluster Geographic Edition infrastructure on the remote cluster.

`-h` *heartbeat-name* — Specifies the custom heartbeat to be used in the partnership to monitor the partner cluster's availability

`-p` *property-setting* — Sets the value of partnership properties with a string of *name=value* pair statements

For more information about the properties you can set, see Appendix A.

partnership-name — Specifies the name of the partnership

Note – A default plug-in that is named `icrm_plugin` is created when the custom heartbeat is added to the partnership.

For more information about using `geops create` command to create a partnership, see "How to Create a Partnership" on page 50.

5. **Verify that your changes were made correctly.**

```
# geoadm status
```

Example 12–9 Adding a Custom Heartbeat Plug-in to a New Custom Heartbeat

This example creates the heartbeat `paris-to-newyork`, which uses a custom heartbeat plug-in, and associate the heartbeat with a new partnership.

```
# geohb create -r cluster-newyork paris-to-newyork
# geohb add-plugin -p query_cmd=/usr/bin/hb/ command1 paris-to-newyork
# geops create -c cluster-newyork -h paris-to-newyork paris-newyork-ps
# geoadm status
```

Configuring Loss of Heartbeat Notification

When a heartbeat is lost, you can configure the Sun Cluster Geographic Edition software to send email notification or to execute an action script. You configure loss of heartbeat notification by using the optional `Notification_emailaddrs` and `Notification_actioncmd` properties.

Heartbeat-loss notification occurs if the heartbeat still fails after the interval you configure with the `Query_interval` property of the heartbeat. The heartbeat monitor sends out a heartbeat request to the responder on the logical host every `Query_interval` period. If no response is received within the `Query_interval`, an internal count is incremented. If the recount reaches the number that is specified in the `heartbeat.retries` property, the heartbeat is deemed to have failed.

For example, you can use sing the default `Query_interval` of 120 seconds and the default `heartbeat.retries` of 3. The heartbeat-lost event will be sent a maximum of 10 minutes after the last heartbeat response from the partner cluster.

```
120sec (delay since last query) + 3*120sec (wait for normal response)
+ 120 sec (wait for retry response)
```

Additional delays can occur between the generation of the heartbeat-lost system event and the triggering of the heartbeat-loss notification. You might experience some further delay of the delivery of email if you configure email notification.

Note – A heartbeat loss event does not necessarily indicate that the remote cluster has crashed.

The following sections describe how to configure the heartbeat-loss notification properties and how to create a custom action script that the Sun Cluster Geographic Edition software executes after a heartbeat-loss event.

Configuring the Loss of Heartbeat Notification Properties

You can configure loss of heartbeat notification by using two partnership properties, `Notification_emailaddrs` and `Notification_actioncmd`. You set these properties by using the `geops` command.

You can set these properties on the default heartbeat during partnership creation. For more information, see "How to Create a Partnership" on page 50. You can also modify these properties by using the procedure that is described in "How to Modify the Heartbeat Properties" on page 209.

If you want to be notified of heartbeat loss by email, set the `Notification_emailaddrs` property. You can specify a list of email addresses, separated by commas. If you want to use email notification, the cluster nodes must be configured as email clients. For more information about configuring mail services, see the *Solaris System Administration Guide: Network Services*.

If you want a command to be executed in response to heartbeat loss, set the `Notification_actioncmd` property.

EXAMPLE 12–10 Configuring Loss of Heartbeat Notification for an Existing Partnership

A notification email address and a custom notification script are specified for the existing partnership, `paris-newyork-ps`, as follows:

```
phys-paris-1# geops set-prop \
-p Notification_emailaddrs=ops@paris.com,ops@newyork.com \
-p Notification_actioncmd=/opt/hb_action.sh paris-newyork-ps
```

Creating an Action Shell Script for Loss of Heartbeat

You can create an action shell script that is executed when the local cluster detects a loss of heartbeat with the partner cluster. The script is executed with root permissions, so the file must have root ownership and execution permissions.

If you have configured the `Notification_actioncmd` property, the action command is executed with arguments that provide information about the event in the following command line:

```
# custom-action-command-path -c local-cluster-name -r remote-cluster-name -e 1 \
-n node-name -t time
```

custom-action-command-path	Specifies a path to the action command you have created
`-c` *local-cluster-name*	Specifies the name of the local cluster
`-p` *remote-cluster-name*	Specifies the name of the remote partner cluster
`-e1`	Specifies that `HBLOST=1`, meaning that a heartbeat-loss event has occurred
`-n`*node-name*	Specifies name of the cluster node that sent the heartbeat-loss event notification

`-t` *timestamp*	Specifies the time of the heartbeat-loss event as the number of milliseconds since January 1, 1970, 00:00:00 GMT

Caution – You can use this script to perform an automatic takeover on the secondary cluster. However, such an automated action is risky. If the heartbeat loss notification is caused by a total loss of all heartbeat connectivity on both the primary and secondary clusters, such an automated action could lead to a situation where two primary clusters exist.

EXAMPLE 12–11 How a Notification Action Script Parses the Command-Line Information Provided by the Sun Cluster Geographic Edition Software

This example shows the event information that is provided in the command-line being parsed in a notification action shell script.

```
#!/bin/sh

set -- `getopt abo: $*`
if [ $? != 0]
then
      echo $USAGE
      exit 2

fi
for i in $*
do

      case $i in
      -p)      PARTNER_CLUSTER=$1; shift;;
      -e)      HB_EVENT=$2; shift;;
      -c)      LOCAL_CLUSTER=$3; shift;;
      -n)      EVENT_NODE=$4; shift;;
      esac
done
```

CHAPTER 13

Monitoring and Validating the Sun Cluster Geographic Edition Software

This chapter describes the files and tools that you can use to monitor and validate the Sun Cluster Geographic Edition software.

This chapter discusses the following topics:

Monitoring the Runtime Status of the Sun Cluster Geographic Edition Software

You can print the runtime status of the local Sun Cluster Geographic Edition enabled cluster by using the `geoadm status` command. When you run this command, it displays output that is organized in the following sections:

- Cluster – This section provides the name of the local cluster.
- Partnership – This section provides information about the partnership, including the name of the partner cluster, the synchronization state, the local heartbeats, and the local heartbeat plug-in.
- Protection group – This section provides information about the status of the defined protection groups, including information about the local cluster and the remote cluster.
- Pending operations – This section provides status information about the current ongoing transaction processes.

You must be assigned the Basic Solaris User RBAC rights profile to run the `geoadm status` command. For more information about RBAC, see "Sun Cluster Geographic Edition Software and RBAC" on page 43.

For example, an administrator runs the `geoadm status` command on `cluster-paris` and the following information is printed:

```
phys-paris-1# geoadm status

Cluster: cluster-paris

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

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

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

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

  Cluster cluster-newyork   : OK
     Role                   : Secondary
     PG activation state    : Activated
     Configuration          : OK
     Data replication       : OK
     Resource groups        : OK

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

The information printed shows that the protection group, `tcpg`, is activated on both the primary cluster, `cluster-paris`, and the secondary cluster, `cluster-newyork`. Data is replicating between the partner clusters and both partners are synchronized.

The following table describes the meaning of the status values.

TABLE 13–1 Status Value Descriptions

Field	Value Descriptions
`Partnership`	`OK` – The partners are connected. `Error` – The connection between the partner clusters is lost. `Degraded` – The partnership has been successfully created but a connection with the partner cluster has not yet been establish. This status value occurs when the partnership has been created and the partner cluster has not been configured.
`Synchronization`	`OK` – The configuration information is synchronized between partner clusters. `Error` – The configuration information differs between the partner clusters. You need to resynchronize the partnership, for a partnership synchronization error, or resynchronize the protection group, for a protection group synchronization error. For information about resynchronizing a partnership, see "Resynchronizing a Partnership" on page 57. For information about resynchronizing a Sun StorEdge Availability Suite 3.2.1 protection group, see "Resynchronizing a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 106 For information about resynchronizing a Hitachi TrueCopy protection group, see "Resynchronizing a Hitachi TrueCopy Protection Group" on page 174. `Mismatch` – Configuration information has been created separately on the clusters. The configuration information must be deleted from one cluster and replaced by a copy of the partner's configuration information. `Unknown` – Information is not accessible because the partners are disconnected or because some of the protection group's components cannot be reached.

TABLE 13–1 Status Value Descriptions *(Continued)*

Field	Value Descriptions
`Heartbeat`	`OK` – Heartbeat checks are running and the partner cluster responds within the specified timeout and retry periods. `Offline` – Heartbeat checks are not running. `Error` – Heartbeat checks are running but the partner is not responding and retries have timed out. `Degraded` – Heartbeat checks are running but one of the primary plug-ins is degraded or not running.
`Heartbeat plug-in`	`OK` – Responses are being received from the partner. `Inactive` – Plug-in is not in use but is a standby for retrying to contact the partner if the other plug-ins obtain no response. `No-Response` – Partner cluster is not responding.
`Protection group` (overall protection group state)	`OK` – The synchronization state is `OK` and the state of the protection group on each cluster is `OK`. `Degraded` – The synchronization state is `OK`. The state of the protection group is `Degraded` on either one or both clusters in the partnership. `Unknown` – The synchronization state or the state of the protection group on one or both clusters is unavailable. The protection group can be online or offline. `Error` – The synchronization state or the state of the protection group on one or both clusters is in `Error`. The protection group can be online or offline.

TABLE 13–1 Status Value Descriptions *(Continued)*

Field	Value Descriptions
`Protection group > Cluster` (state of protection group on each cluster)	`OK` – The state of all the protection group components, such as configuration data, data replication, or resource groups, is `OK`, `NONE`, or `N/A` on the cluster. `Degraded` – The state of one or more of the protection group components is in the `Degraded` state on the cluster. `Unknown` – The state of some of the protection group's components, such as configuration data, data replication, or resource groups, is unavailable. `Error` – The state of some of the protection group's components, such as configuration data, data replication, or resource groups, is in `Error`.
`Protection group > Cluster > Role`	`Primary` – The cluster is the `Primary` for this protection group. `Secondary` – The cluster is the `Secondary` for this protection group. `Unknown` – Information is not accessible because the partners are disconnected or because some of the protection group's components cannot be reached.
`Protection group > Cluster > PG activation state`	`Activated` – The protection group is activated. `Deactivated` – The protection group is deactivated. `Unknown` – Information is not accessible because the partners are disconnected or because some of the protection group's components cannot be reached.

TABLE 13–1 Status Value Descriptions *(Continued)*

Field	Value Descriptions
`Protection group > Cluster > Configuration`	`OK` – Protection group configuration has been validated without errors on the cluster. `Error` – Protection group configuration validation resulted in errors on the cluster. You need to revalidate the protection group. For information about validating a Sun StorEdge Availability Suite 3.2.1 protection group, see "How to Validate a Sun StorEdge Availability Suite 3.2.1 Protection Group" on page 84. For information about validating a Hitachi TrueCopy protection group, see "How to Validate a Hitachi TrueCopy Protection Group" on page 147. `Unknown` – Information is not accessible because the partners are disconnected or because some of the protection group's components cannot be reached.

TABLE 13–1 Status Value Descriptions *(Continued)*

Field	Value Descriptions
`Protection group` > `Cluster` > `Data replication`	`None` – Data replication is not configured. `OK` – Data replication is running and data is synchronized with the partner cluster when the protection group is activated. Replication is suspended when the protection group is deactivated. This state represents data replication on this cluster and does not reflect the overall state of data replication. This state is mapped from the corresponding state in the data replication subsystem. `Degraded` – Data is not replicated and not synchronized with the partner cluster when the protection group is activated. New writes will succeed but not be replicated. This state represents data replication on this cluster and does not reflect the overall state of data replication. This state is mapped from the corresponding state in the data replication subsystem. `Error` – Data replication from the primary cluster to the secondary cluster is in error if the data replication subsystem reports an error or if data replication is not suspended when the protection group is deactivated. This state represents data replication on this cluster and does not reflect the overall state of data replication. This state is mapped from the corresponding state in the data replication subsystem. `Unknown` – Information is not accessible because the partners are disconnected or because some of the protection group's components cannot be reached. `N/A` – The data replication state of the protection group could not be mapped. Data replication is in a valid state on its own but in an `Error` state for the protection group. This state is available only if you are using Sun StorEdge Availability Suite 3.2.1 data replication.

TABLE 13–1 Status Value Descriptions *(Continued)*

Field	Value Descriptions
`Protection group` > `Cluster` > `Resource groups`	`None` – No resource groups is protected by this protection group. `OK` – If the cluster has the `Primary` role, all resource groups are online when the protection group is activated or unmanaged when the protection group is deactivated. If the cluster is has the `Secondary` role, all resource groups are unmanaged. `Error` – If the cluster has the `Primary` role, not all resource groups are online when the protection group is activated or unmanaged when the protection group is deactivated. If the cluster is has the `Secondary` role, not all resource groups are unmanaged. `Unknown` – Information is not accessible because the partners are disconnected or because some of the protection group's components cannot be reached.

For more specific information about checking the runtime status of replication, see "Checking the Runtime Status of Sun StorEdge Availability Suite 3.2.1 Data Replication" on page 107 or "Checking the Runtime Status of Hitachi TrueCopy Data Replication" on page 175.

Viewing the Sun Cluster Geographic Edition Log Messages

All of the Sun Cluster Geographic Edition components produce messages that are stored in log files.

Information about the loading, running, and stopping Sun Cluster Geographic Edition components into the common agent container is recorded in the following log files. The most recently logged messages are in file 0, then 1, and 2.

- `/var/opt/SUNWcacao/logs/cacao.0`
- `/var/opt/SUNWcacao/logs/cacao.1`
- `/var/opt/SUNWcacao/logs/cacao.2`

System log messages are stored in the `/var/adm/messages` log file.

Each cluster node keeps separate copies of the previous log files. The combined log files on all cluster nodes form a complete snapshot of the currently logged information. The Sun Cluster Geographic Edition module's log messages are updated on the node where the Sun Cluster Geographic Edition software is currently active. The data replication control-log messages are updated on the node where the data replication resource is currently `Online`.

Printing Configuration Information About Partnerships and Protection Groups

You can print the current local cluster partnership configuration. The information that you print includes a listing of all partnerships that are defined between the local cluster and remote clusters.

You can also print the current configuration of a specific protection group or of all the protection groups that are defined on a cluster.

▼ How to Display Configuration Information About Partnerships

Steps

1. **Log in a cluster node.**

   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" on page 43.

2. **Display information about the partnership.**

   ```
   # geops list partnership-name
   ```

   partnership-name Specifies the name of the partnership. If you do not specify a partnership, then the `geops list` command displays information on all the partnerships.

   For information about the names and values that are supported by Sun Cluster Geographic Edition software, see Appendix B.

Example 13–1 Displaying Partnership Configuration Information

This example displays configuration information about the partnership between local `cluster-paris` and remote `cluster-newyork`.

```
# geops list paris-newyork-ps
```

▼ How to Display Configuration Information About Protection Groups

Steps

1. **Log in to a cluster node.**

   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" on page 43.

2. **Display information about one or all protection groups.**

   ```
   # geopg list [protection-group-name]
   ```

   protection-group-name — Specifies the name of a protection group

   If you do not specify a protection group, then the command lists information about all of the protection groups configured on your system.

Example 13–2 Displaying Configuration Information About a Protection Group

This example displays configuration information for `avspg`, which is configured on `cluster-paris`.

```
# geopg list avspg
```

CHAPTER 14

Customizing Switchover and Takeover Actions

This chapter describes how you can create a script that is executed when a protection group's role changes from secondary to primary. The chapter contains the following topics:

Creating a Role-Change Action Script

You can configure the Sun Cluster Geographic Edition software to run a command when a cluster within a protection group changes from the `secondary` to the `primary` role. This change can happen as a result of either a switchover or takeover operation.

The action command is executed on the cluster where the role is changing from secondary to primary, with arguments that provide information about the event. The script is executed before the Sun Cluster Geographic Edition software bring the resource groups online. The following command-line runs the script:

```
# custom-action-command-path -o primary -c cluster-name \
 -s partnership-name protection-group-name user-arguments
```

custom-action-command-path	Specifies a path to the action command you have created.
`-o primary`	Specifies that the role being assumed by the cluster is `primary`.
`-c` *cluster-name*	Specifies the name of the secondary cluster that is assuming the new role of primary cluster.

`-s` *partnership-name*	Specifies the name of the partnership that hosts the protection group.
protection-group-name	Specifies the name of the protection group that is undergoing the role change.
user-arguments	Specifies static arguments that are passed after all of the Sun Cluster Geographic Edition supplied options. This free-form string can be parsed by the script as required. For example, you could specify a list of `key=value` pairs, such as `name=sun.com,ip=10.1.2.3`. You could also specify a sequence of options, such as `-n sun.com -a 10.1.2.3.4`. The format of these arguments is not restricted by the Sun Cluster Geographic Edition software.

The exit status of the role-change action script is reported as part of the result of the `geopg switchover` or `geopg takeover` command. The exit status is zero if the action script was launched successfully. A nonzero exit status indicates an error or failure. The value of the exit status does not affect other aspects of the role-change actions. The switchover or takeover proceeds to bring the application resource groups in the protection group online, regardless of the exit status of the action script.

You should wait for the script to return before proceeding with other operation. Consider the time required to run the script when creating the action script.

EXAMPLE 14–1 Switchover Action Script for Updating the DNS

This example creates a script that uses the `nsupdate` command to reconfigure the host name to point to a new cluster. For more information about the `nsupdate` command, refer to the `nsupdate`(1M) man page.

Clients that try to connect to `companyX.com` are referred by the name service to the address of the primary cluster for a protection group, `cluster-paris`. When the primary cluster fails to respond, the administrator performs a switchover of the protection group to the alternative cluster, `cluster-newyork`.

```
#!/bin/ksh
# script to update dns
# Assumes each cluster has an entry with name "lh-paris-1" in /etc/hosts
# but different value for the IP in each cluster
# for forward DNS (A) entry: will delete old entry for "lh-paris-1"
# and add one that is correct for "this cluster"
#
# For reverse (PTR) DNS entry, will just add one for this cluster.
# Will NOT delete PTR record left over from old cluster. So
# eventually you will just have reverse lookup for the IP for both clusters
# doing reverse resolution to the same name (lh-paris-1.odyssey.com)
# This should be fine, as long as the forward resolution stays "correct"
```

EXAMPLE 14–1 Switchover Action Script for Updating the DNS *(Continued)*

```
#
# The blank line of input at the end of nsupdate is REQUIRED
#
# A short TTL is put on the new records (600 = 10 minutes)
# but you can't really control what kind of caching goes on on
# the client side


# get IP corresponding to name "lh-paris-1" on THIS Cluster
NEWIP=$(getent hosts lh-paris-1|cut -f1)

# this bit splits out the octets in order to add the reverse PTR entry
IFS=.
set $NEWIP
unset IFS

/usr/sbin/nsupdate <<ENDNSUPDATE
update delete ora-lh.odyssey.com A
update add ora-lh.odyssey.com 600 A $NEWIP
update add $4.$3.$2.$1.in-addr.arpa 600 PTR ora-lh.odyssey.com.

ENDNSUPDATE
```

Configuring a Protection Group to Execute a Script at Switchover or Takeover

After you have created a script, you must configure the protection group to execute the script if a switchover or takeover occurs. If a switchover or takeover occurs, the script is executed on the cluster that is becoming primary.

▼ How to Configure a Protection Group to Execute a Script at Switchover or Takeover

Steps

1. **Log in to one of the cluster nodes.**

   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" on page 43.

2. **Configure the `RoleChange_ActionCmd` and `RoleChange_ActionArgs` properties of the protection group.**

```
# geopg set-prop -p RoleChange_ActionCmd=fully-qualified-script -p RoleChange_ActionArgs=script-arguments
```

`-p` *property-setting*	Sets the properties of the protection group Specify the path to the command by using the `RoleChange_ActionCmd` property. This path should be valid on all nodes of all partner clusters that can host the protection group. Define the arguments that you want to append to the command line when the action command is run by using the `RoleChange_ActionArgs` property. For more information about the properties you can set, see Appendix A.
protection-group-name	Specifies the name of the protection group

Example 14–2 Configuring a Protection Group to Execute a Command at Cluster Switchover or Takeover

The following command configures a protection group to execute a custom command called `newDNS`:

```
# geopg set-prop -p RoleChange_ActionCmd=/usr/bin/newDNS \
-p RoleChange_ActionArgs=domain=companyx.com,ip=1.2.3.4 avspg
```

APPENDIX **A**

Standard Sun Cluster Geographic Edition Properties

This appendix provides the standard properties of Sun Cluster Geographic Edition heartbeats, heartbeat plug-in, partnerships, protection groups, and data replication device groups.

This appendix contains the following topics:

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

General Heartbeat Properties

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

TABLE A–1 General Heartbeat Properties

Property Name	Description
`Query_interval` (integer)	Specifies the delay in seconds between heartbeat status requests. Tuning recommendations: The value of this property is assigned at creation and can be tuned at runtime. Category: Optional Default: 120 seconds

General Heartbeat Plug-in Properties

The following table describes the general heartbeat plug-in properties that the Sun Cluster Geographic Edition software defines.

TABLE A–2 General Heartbeat Plug-in Properties

Property	Description
`Plugin_properties` (string)	Specifies a property string specific to the plug-in. Tuning recommendations: The value of this property is assigned at creation and can be tuned at runtime. Category: Optional Default: None except for heartbeats that use the default heartbeat plug-ins, `tcp_udp_plugin` and `ping-plugin`. For the `tcp_udp_plugin` plug-in, the format of this string is predefined as *remote-IP-address*`/UDP/2084/ipsec`, *remote-IP-address*`/TCP/2084/ipsec`. The *remote_IP_address* argument specifies the IP address of the partner cluster. The optional `ipsec` argument specifies if the plug-in uses IPsec with a Boolean value of `true` or `false`. For the `ping-plugin`, the format of this string is predefined as *remote_IP_address*, where *remote_IP_address* specifies the IP address of the partner cluster.

TABLE A–2 General Heartbeat Plug-in Properties *(Continued)*

Property	Description
`Query_cmd` (string)	Specifies the path to the heartbeat status request command. Tuning recommendations: The value of this property is assigned at creation and can be tuned at runtime. Category: Required property if the plug-in does not specify a predefined plug-in. Default: None
`Requester_agent` (string)	Specifies the absolute path to the requester agent. Tuning recommendations: The value of this property is assigned at creation and can be tuned at runtime. However, the `Requester_agent` property of the default plug-in should never need to be tuned except for testing purposes. Category: Optional Default: None
`Responder_agent` (string)	Specifies the absolute path to the responder agent. Tuning recommendations: The value is assigned at creation and can be tuned at runtime. However, the `Responder_agent` property of the default plug-in should never need to be tuned except for testing purposes. Category: Optional Default: None
`Type` (enum)	Designates the type of plug-in. Set to either `primary` or `backup`. Tuning recommendations: The value of this property is assigned at creation and can be tuned at runtime. Category: Required Default: None, except for the default heartbeat that is named `ping_plugin`. If using this plug-in, the default value is `backup`.

Partnership Properties

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

TABLE A–3 Partnership Properties

Property	Description
`Description` (string)	Describes the partnership. Tuning recommendations: The value of this property is assigned at creation and can be tuned at runtime. Category: Optional Default: Empty string
`Notification_ActionCmd` (string)	Provides the path to the action script that is triggered when heartbeat-loss notification is issued. Tuning recommendations: The value of this property is assigned at creation and can be tuned at runtime. Category: Optional Default: Empty string
`Notification_EmailAddrs` (string array)	Lists the email addresses that are sent email when heartbeat-loss notification is issued. The list is comma delimited. Tuning recommendations: The value of this property is assigned at creation and can be tuned at runtime. Category: Optional Default: Empty string

General Properties of a Protection Group

The following table describes the protection group properties that the Sun Cluster Geographic Edition software defines.

TABLE A–4 General Properties of a Protection Group

Property	Description
`Description` (string)	Describes the protection group. Tuning recommendations: This property can be tuned at any time. Category: Optional Default: Empty string

TABLE A–4 General Properties of a Protection Group *(Continued)*

Property	Description
`RoleChange_ActionArgs` (string)	Defines a string of arguments that are appended to the end of the command line when the role-change action command, `RoleChange_ActionCmd`, is run. Tuning recommendations: This property can be tuned at any time. Category: Optional Default: Empty string
`RoleChange_ActionCmd` (string)	Specifies the path to an executable command. This command is run when the cluster designated as primary for a protection group changes. This path should be valid on all nodes of all partner clusters that can host the protection group. Tuning recommendations: This property can be tuned at any time. Category: Optional Default: Empty string
`Timeout` (integer)	Specifies the timeout period for the protection group in seconds. The timeout period is the longest time Sun Cluster Geographic Edition waits for a response after a `geopg` command is executed, such as `start`, `stop`, `switchover`, and `takeover`. If the command does not respond within the timeout period, Sun Cluster Geographic Edition reports the operation as timed out, even if the underlying command that was executed eventually completes successfully. The timeout period applies to operations on a per-cluster basis. An operation with a local scope times out if the operation does not complete after the specified timeout period. An operation with a global scope consists of an action on the local cluster and an action on the remote cluster. The local and remote action are timed separately. So, an operation with a global scope times out if the local operation does not complete after the specified timeout period or if the remote operation does not complete after the specified timeout period. Tuning recommendations: This property can be tuned only when the protection group is offline. Category: Optional Range: 20-1000000 seconds Default: 200

Sun StorEdge Availability Suite 3.2.1 Properties

The following table describes the Sun StorEdge Availability Suite 3.2.1 properties that the Sun Cluster Geographic Edition software defines.

TABLE A–5 Sun StorEdge Availability Suite 3.2.1 Properties

Property	Description
Data Replication Property: `Nodelist` (string array)	Lists the host names of the machines that can be primary for the device group in the protection group. Device groups in the protection group must share the same ordered nodelist. This list is comma delimited. Tuning recommendations: This property can be tuned only when the protection group is offline. Category: Optional Default: All the nodes in the cluster
Device Group Property: `Enable_volume_set` (Boolean)	Defines whether the volume sets that are defined in the file (`/var/cluster/geo/avs/`*<AVS-device-group-name>*`-volset.ini`) are enabled when the device group is added. Set to either `true` or `false`. Tuning recommendations: This property cannot be tuned after it has been successfully validated during creation, replication, or synchronization. Category: Optional Default: `false`
Device Group Property: `Local_logical_host` (string)	Defines the local logical hostname that is used for the replication of the device group. Do not use an underscore (_) character in the logical hostname. Tuning recommendations: This property cannot be tuned after it has been successfully validated during creation, replication, or synchronization. Category: Required Default: None

TABLE A–5 Sun StorEdge Availability Suite 3.2.1 Properties *(Continued)*

Property	Description
Device Group Property: `Remote_logical_host` (string)	Defines the remote logical hostname that is used for the replication of the device group. Do not use an underscore (_) character in the logical hostname. Tuning recommendations: This property cannot be tuned after it has been successfully validated during creation, replication, or synchronization. Category: Required Default: None

Sun StorEdge Availability Suite 3.2.1 Properties That Should Not Be Changed

Some data replication properties that Sun Cluster Geographic Edition software changes internally must not be edited manually.

For Sun StorEdge Availability Suite 3.2.1, do not edit the following properties:

- `Network_resources_used`
- `Device_group`
- `Remote_logical_host`
- `Role`

Hitachi TrueCopy Properties

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

TABLE A–6 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
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 the 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. For more information about setting this property, see "How to Add a Data Replication Device Group to a Hitachi TrueCopy Protection Group" on page 155 Tuning recommendations: This property can only be tuned when the protection group is offline. Category: Required Default: None

Hitachi TrueCopy Properties That Should Not Be Changed

Some data replication properties that Sun Cluster Geographic Edition software changes internally must not be edited manually.

For Hitachi TrueCopy, do not edit the following properties:

- `dev_group`
- `replication_role`

APPENDIX **B**

Legal Names and Values of Sun Cluster Geographic Edition Entities

This appendix lists the requirements for legal characters for the names and values of Sun Cluster Geographic Edition entities.

This appendix contains the following topics:

Legal Names for Sun Cluster Geographic Edition Entities

Sun Cluster Geographic Edition entity names consist of the following:

- Host names
- Partnership names
- Protection group names
- Custom heartbeat names

All names must comply with the following rules:

- Must be in ASCII.
- Must start with a letter.
- Can contain upper and lowercase letters, digits, dashes (-), and underscores (_).
- Must not exceed 255 characters.

Legal Values for Sun Cluster Geographic Edition Entities

The values of Sun Cluster Geographic Edition entities fall into two categories: property values and description values, both of which share the same rules, as follows:

- Values must be in ASCII.
- The maximum length of a value is 4 megabytes minus 1, that is, 4,194,303 bytes.
- Values cannot contain any of the following characters: newline or semicolon.

APPENDIX C

Takeover Postconditions

This appendix provides details about the state of the primary and secondary clusters after a `geopg takeover` command is executed.

This appendix contains the following topics:

Results of a Takeover When the Partner Cluster Can Be Reached

This section describes the activation state of the primary and secondary clusters before and after a `geopg takeover` command is executed. The results described in this section assume that the partner cluster can be reached

The following table describes the states when the `geopg takeover` command is executed on the secondary cluster, `cluster-newyork`.

TABLE C–1 Takeover Results When `geopg takeover` Is Executed on the Secondary Cluster

Cluster Role and State Before Takeover	Cluster Role and State After Takeover
`cluster-paris`: primary, deactivated `cluster-newyork`: secondary, deactivated	`cluster-paris`: secondary, deactivated `cluster-newyork`: primary, deactivated

TABLE C–1 Takeover Results When `geopg takeover` Is Executed on the Secondary Cluster *(Continued)*

Cluster Role and State Before Takeover	Cluster Role and State After Takeover
`cluster-paris`: primary, activated `cluster-newyork`: secondary, deactivated	`cluster-paris`: secondary, deactivated `cluster-newyork`: primary, deactivated
`cluster-paris`: primary, deactivated `cluster-newyork`: secondary, activated	`cluster-paris`: secondary, deactivated `cluster-newyork`: primary, activated, with data replication stopped
`cluster-paris`: primary, activated `cluster-newyork`: secondary, activated	`cluster-paris`: secondary, deactivated `cluster-newyork`: primary, activated, with data replication stopped

The following table describes the states when the `geopg takeover` command is executed on the primary cluster, `cluster-paris`.

TABLE C–2 Takeover Results When `geopg takeover` Is Executed on the Primary Cluster

Cluster Role and State Before Takeover	Cluster Role and State After Takeover
`cluster-paris`: primary, deactivated `cluster-newyork`: secondary, deactivated	`cluster-paris`: primary, deactivated `cluster-newyork`: secondary, deactivated
`cluster-paris`: primary, activated `cluster-newyork`: secondary, deactivated	`cluster-paris`: primary, activated, with data replication stopped `cluster-newyork`: secondary, deactivated
`cluster-paris`: primary, deactivated `cluster-newyork`: secondary, activated	`cluster-paris`: primary, deactivated `cluster-newyork`: secondary, deactivated
`cluster-paris`: primary, activated `cluster-newyork`: secondary, activated	`cluster-paris`: primary, activated, with data replication stopped `cluster-newyork`: secondary, deactivated

Results of a Takeover When the Partner Cluster Cannot Be Reached

This section describes the activation state of the primary and secondary clusters before and after a `geopg takeover` command is issued when the partner cluster cannot be reached or when the protection group on the partner cluster is busy.

The following table describes the states when the `geopg takeover` command is issued on the secondary cluster, `cluster-newyork`, and the primary cluster cannot be reached or the protection group on the primary cluster is busy.

Note – The cluster role and state after the takeover, which is given in the table, is only available when the partner cluster can once more be reached.

TABLE C–3 Takeover Results When `geopg takeover` Is Executed on the Secondary Cluster and the Primary Cluster Cannot Be Reached

Cluster Role and State Before Takeover	Cluster Role and State After Takeover
`cluster-paris`: primary, deactivated, synchronization status `Unknown` `cluster-newyork`: secondary, deactivated, synchronization status `Unknown`	`cluster-paris`: primary, deactivated, synchronization status `Error` `cluster-newyork`: primary, deactivated, synchronization status `Error`
`cluster-paris`: primary, activated, synchronization status `Unknown` `cluster-newyork`: secondary, deactivated, synchronization status `Unknown`	`cluster-paris`: primary, activated, synchronization status `Error` `cluster-newyork`: primary, deactivated, synchronization status `Error`
`cluster-paris`: primary, deactivated, synchronization status `Unknown` `cluster-newyork`: secondary, activated, synchronization status `Unknown`	`cluster-paris`: primary, deactivated, synchronization status `Error` `cluster-newyork`: primary, activated, with data replication stopped, synchronization status `Error`
`cluster-paris`: primary, activated, synchronization status `Unknown` `cluster-newyork`: secondary, activated, synchronization status `Unknown`	`cluster-paris`: primary, activated, synchronization status `Error` `cluster-newyork`: primary, activated, with data replication stopped, synchronization status `Error`

The following table describes the states when the `geopg takeover` command is executed on the primary cluster, `cluster-paris`, and the secondary cluster cannot be reached or the protection group on the secondary cluster is buy.

TABLE C–4 Takeover Results When `geopg takeover` Is Executed on the Primary Cluster and the Secondary Cluster Cannot Be Reached

Cluster Role and State Before Takeover	Cluster Role and State After Takeover
`cluster-paris`: primary, deactivated, synchronization status `Unknown` `cluster-newyork`: secondary, deactivated, synchronization status `Unknown`	`cluster-paris`: primary, deactivated, synchronization status `OK`, `Error`, or `Mismatch` `cluster-newyork`: secondary, deactivated, synchronization status `OK`, `Error`, or `Mismatch`
`cluster-paris`: primary, activated, synchronization status `Unknown` `cluster-newyork`: secondary, deactivated, synchronization status `Unknown`	`cluster-paris`: primary, activated, with data replication stopped, synchronization status `OK`, `Error`, or `Mismatch` `cluster-newyork`: secondary, deactivated, synchronization status `OK`, `Error`, or `Mismatch`
`cluster-paris`: primary, deactivated, synchronization status `Unknown` `cluster-newyork`: secondary, activated, synchronization status `Unknown`	`cluster-paris`: primary, deactivated, synchronization status `OK`, `Error`, or `Mismatch` `cluster-newyork`: secondary, activated, synchronization status `OK`, `Error`, or `Mismatch`
`cluster-paris`: primary, activated, synchronization status `Unknown` `cluster-newyork`: secondary, activated, synchronization status `Unknown`	`cluster-paris`: primary, activated, with data replication stopped, synchronization status `OK`, `Error`, or `Mismatch` `cluster-newyork`: secondary, activated, synchronization status `OK`, `Error`, or `Mismatch`

Index

D

I

J

L

M

N

O

P

R

S

T

V